

TWO LOGICAL FACES OF BELIEF REVISION ¹

Johan van Benthem, Amsterdam & Stanford

Draft January 2012

Abstract This piece proposes a style of thinking using modal frame correspondence that puts Segerberg's dynamic doxastic logic and 'Dutch' dynamic-epistemic logic for belief change in one setting. While our technical results are elementary, they do suggest new lines of thought.

1 Two modal logics for belief change

Belief revision theory is a small corner of the world of philosophy and computer science, and modal logic is a small corner of the world of logic. When two specialized topics come together, surely, there can be only one way of doing that? The dynamic-doxastic logic *DDL* of Segerberg 1995, 1999 ² has abstract modal operators describing transitions in abstract universes of models to describe changes in belief, and then encodes basic postulates on belief change in modal axioms that can be studied by familiar techniques. But there is also another line in the logical literature, started in van Benthem 2007, Baltag & Smets 2008 ³ that works differently. Here belief changes are modeled in the framework of dynamic-epistemic logic (*DEL*) as acts of changing a plausibility ordering in a current model, and the update rule for doing that is made explicit, while its properties are axiomatized completely in modal terms. The contrast may be stated as follows. Segerberg follows *AGM* belief revision theory (Gaerdenfors 1988) in its *postulational* approach constraining spaces of all possible belief changes, while the *DEL* approach is *constructive*, studying specific update rules and the complete logics of their corresponding dynamic model-changing modalities.

Stated this way, there need not be any conflict between the two approaches – and in fact, there is not. Still, there are many differences in their subsequent technical agenda.

¹ Krister Segerberg's seminal work has been a beacon in modal logic ever since the late 1960s. Add the attractive personality to the deep intellect, and one understands why my writing in this volume is a case of duty coinciding, not just with Kantian inclination, but with active desire.

² See also Leitgeb & Segerberg 2007 for extensive discussion of the research program.

³ Relevant predecessors to this work are Aucher 2004, van Ditmarsch 2005.

⁴ One could spend much time analyzing these differences, but my aim in this paper is modest. I want to suggest that, for colleagues from modal logic, *DDL* and *DEL* fit very well, if we use the method of *frame correspondence*. This suggestion occurs in van Benthem 2011, but I will pursue it more systematically here. My results are simple technically, but they suggest new perspectives. I start with knowledge in Section 2, exploring frame correspondences for ‘public announcement logic’ *PAL*. Many general methodological points can be made at this level, as they are not specific to belief. Next, I give modal correspondence for logics of belief change in Section 3. In Section 4, I discuss two generalizations: full dynamic-epistemic logic with product update over event models, and an extension of correspondence analysis to neighborhood models, using the *DEL* treatment in van Benthem & Pacuit 2011. Section 5 lists new general issues coming to light in my analysis, all of them ‘to be explored’. Section 6 states the conclusion of this paper, though it will already be clear right here at the start: the two existing styles of modal logic for belief revision live well together, and analyzing their connections actually reveals some interesting issues that will unfold in due course.

2 Correspondence for information update and knowledge

We start with a phenomenon that is not very interesting in the *AGM* style, though it becomes wildly exciting when we study it in a constructive setting: update with new hard information that shrinks agents’ current ranges of options for the actual situation.

2.1 Hard information, knowledge, and public announcement logic

Basic epistemic logic We start by recalling some basics. Standard epistemic logic *EL* describes semantic information encoded in agents’ ranges of uncertainty. The language extends propositional logic with modal operators $K_i\varphi$ (i knows that φ), for agents i , and $C_G\varphi$ (φ is common knowledge in group G). Epistemic models $\mathbf{M} = (W, \{\sim_i\}_{i \in I}, V)$

⁴ *DEL*-style logics of belief revision depart from the *AGM*-format in a number of ways. (i) The content of new beliefs need not be factual, but it can itself consist of complex statements about beliefs. (ii) What changes in acts of revision is not just beliefs, but crucially also conditional beliefs. (iii) Infinitely many types of triggering event can be analyzed structurally in the logic by mechanisms like ‘event models’ or ‘model-change programs’. (iv) The setting is essentially multi-agent, making, in principle, social acts of belief merge as crucial to the logical system as individual acts of revision (cf. the logics for merging in Girard 2008, Liu 2011).

have a set of worlds W , accessibility relations \sim_i for agents i in some total group I , and a valuation V for proposition letters. Pointed models (\mathbf{M}, s) mark a actual world s .⁵ The key truth condition is that $\mathbf{M}, s \models K_i\varphi$ iff for all worlds t with $s \sim_i t$: $\mathbf{M}, t \models \varphi$.^{6,7} Complete logics capturing epistemic reasoning about oneself and others are known (Fagin et al. 1995). The base system is a minimal modal logic. A restriction to equivalence relations add $S5$ axioms of positive and negative introspection, while the complete logic of common knowledge can be axiomatized with *PDL*-techniques.

Information update by elimination Now for the logical dynamics of information flow. An event $!\varphi$ yielding the information that φ is true shrinks the current model to just those worlds that satisfy φ . This is the well-known notion of *public hard information*. More precisely, for any epistemic model \mathbf{M} , world s , and formula φ true at s , the new $(\mathbf{M}|\varphi, s)$ (*M relativized to φ at s*) is the sub-model of \mathbf{M} whose domain is the set $\{t \in \mathbf{M} \mid \mathbf{M}, t \models \varphi\}$. This mechanism models public communication, but also public observation. There is much more to this dynamics than meets the eye in standard views of ‘mere update’ with factual formulas. For instance, crucially, truth values of complex epistemic formulas may change after update: agents who did not know that φ now do. Therefore, it makes sense to get clear on the exact dynamic logic behind this.

Public announcement logic The language of *public announcement logic* *PAL* adds action expressions to *EL*, plus matching modalities, defined by the syntax rules:

Formulas	$F:$	$p \mid \neg\varphi \mid \varphi \vee \psi \mid K_i\varphi \mid C_G\varphi \mid \langle A \rangle \varphi$
Action expressions	$A:$	$!F$

The semantic clause for the dynamic action modality looks ahead between models:

$$\mathbf{M}, s \models \langle !\varphi \rangle \psi \text{ iff } \mathbf{M}, s \models \varphi \text{ and } \mathbf{M}|\varphi, s \models \psi$$

⁵ Further relational conditions on \sim_i encode special assumptions about agents’ powers of observation and introspection: very common is the special case of equivalence relations.

⁶ As for common knowledge, $\mathbf{M}, s \models C_G\varphi$ iff for all worlds t that are reachable from s by some finite sequence of arbitrary \sim_i steps ($i \in G$): $\mathbf{M}, t \models \varphi$.

⁷ In what follows, for convenience, we mostly suppress agent indices, and use standard modal notation for the epistemic modality of one accessibility relation R . Also for convenience, we will work mostly with existential modalities \diamond instead of universal boxes \square .

PAL is axiomatized by any complete logic over static models plus *recursion axioms*

$$\begin{aligned}
\langle !\varphi \rangle q &\leftrightarrow (\varphi \wedge q) && \text{for proposition letters } q \\
\langle !\varphi \rangle (\psi \vee \chi) &\leftrightarrow (\langle !\varphi \rangle \psi \vee \langle !\varphi \rangle \chi) \\
\langle !\varphi \rangle \neg \psi &\leftrightarrow (\varphi \wedge \neg \langle !\varphi \rangle \psi) \\
\langle !\varphi \rangle \diamond \psi &\leftrightarrow (\varphi \wedge \diamond \langle !\varphi \rangle \psi)
\end{aligned}$$

Intuitively, the final recursion axiom for knowledge captures the essence of getting hard information. We will see in just which sense this is true in our further analysis. For further theory and applications of *PAL* and related systems, cf. Baltag, Moss and Solecki 1998, van Ditmarsch, van der Hoek and Kooi 2007, van Benthem 2011.

2.2 Switching directions: from valid axioms to constraints

PAL is about one constructive way of taking incoming hard information: elimination of incompatible worlds. Now we reverse the perspective. Let us ask which postulates look plausible for hard update, of course, always keeping in mind that our intuitions need to be valid for arbitrary propositions, bringing the logic in harmony.⁸ Having done that, we can see which transformations of models validate them. This sounds grand. In what follows, however, I take a simple approach, investigating the recursion axioms of *PAL* themselves as postulates, since they have a lot of general appeal. To make this work, we need a suitably abstract setting – close to the models of *DDL*.⁹

Update universe and update relations Consider any family \mathbf{M} of pointed epistemic models (M, s) , viewed as an ‘update universe’ where model changes can take place. Possible changes are given as a family of update relations $R_p(\mathbf{M}, s)(\mathbf{N}, t)$ relating pointed models, where the index set P is a subset of \mathbf{M} : intuitively, the proposition triggering the update. One can think of the R as recording the action of some update

⁸ It is a curiously overlooked mismatch that modal logics for philosophical notions are often based on philosophers’ intuitions about factual statements only, whereas the logic itself also deals with complex assertions that make good sense, for which the philosophers’ intuitions might have to be different. Other imbalances of this sort occur in logics for non-standard consequence relations, and accounts of knowledge proposed in formal epistemology.

⁹ The setting chosen here is more abstract and flexible than that used in the correspondence analysis of van Benthem 2011, and it removes some infelicities in that earlier treatment.

operation \heartsuit occurring in the syntax of our language that depends on the proposition P . Here different operations can have different effects: from our hard updates $!\varphi$ to the soft updates $\uparrow\varphi$ to be discussed below. As just said, this is essentially the semantic setting of Krister Segerberg's dynamic doxastic logic, where each transition relation has a matching modality.¹⁰ Now, for each formula φ , let $[[\varphi]]$ be the set of worlds in \mathbf{M} satisfying φ . We set, for the update modality matching R :

$$\mathbf{M}, s \models \langle \heartsuit\varphi \rangle \psi \quad \text{iff} \quad \text{there exists a model } (\mathbf{N}, t) \text{ in } \mathbf{M} \text{ with} \\ SR_{[[\varphi]]}(\mathbf{M}, s)(\mathbf{N}, t) \text{ and } (\mathbf{N}, t) \models \psi$$

Remark To be yet more precise, we are really interpreting our language in a *three-index format* $\mathbf{M}, \mathbf{M}, s$, and for the accessibility relations R in this update universe \mathbf{M} , we have that $(\mathbf{M}, s) R (\mathbf{M}, t)$ iff Rst in \mathbf{M} , without any jumps out of the model \mathbf{M} . This precision can be ignored for most of what follows, but it will come up occasionally.

2.3 A correspondence theorem for eliminative update

In what follows, the reader is supposed to know how modal frame correspondence works: cf. the textbooks van Blackburn, de Rijke & Venema 2000, van Benthem 2010. We will analyze the *PAL* recursion axioms one by one in this style to see what they say, as a way of determining their total content as a correspondence constraint on update operations. But before doing so, we need to address a subtlety.

Substitution closure Correspondence arguments use frame truth of modal formulas, i.e., truth under all possible valuations for the proposition letters. Thus, if a formula is true, so are all its substitution instances: proposition letters are schematic variables for arbitrary propositions. But this sits badly with the system *PAL*, whose valid principles are not closed under substitution. In particular, the base axiom $\langle !\varphi \rangle q \leftrightarrow (\varphi \wedge q)$ is only valid for proposition letters q . Substituting to the general form $\langle !\varphi \rangle \psi \leftrightarrow (\varphi \wedge \psi)$ yields obviously invalid instances for epistemic assertions ψ . Much can be said about

¹⁰ This is not the only possible format, and one can experiment with others. In particular, making the relational transition depend on just an extensional set of worlds reflects the valid *PAL* rule of *Replacement of Provable Equivalents*. Stated as one axiom in a language extended with a universal modality U ranging over the whole universe, this is the following implication making announced propositions 'extensional': $U(\varphi \leftrightarrow \psi) \rightarrow (\langle !\varphi \rangle \alpha \leftrightarrow \langle !\psi \rangle \alpha)$.

this phenomenon (cf. Holliday, Hoshi & Icard 2011), but in this paper, we take a simple line. We will first analyze the substitution-closed principles of *PAL*, and then return to the correspondence status of the base axiom. Thus, for the moment, we only look at the following obviously substitution-closed special case:

$$\langle !\varphi \rangle T \Leftrightarrow \varphi$$

In our correspondence setting, substitution failures relate to the semantics of atomic propositions p . Inside one epistemic model \mathbf{M} , the obvious choice seems to be sets of worlds. But in an update universe \mathbf{M} as above, propositions range over all *pairs* (\mathbf{M}, s) , and hence one p could have different truth values at pairs (\mathbf{M}, s) , (\mathbf{N}, s) . We will view Greek letters in axioms as standing for such general context-dependent propositions in what follows, returning to the original view of *PAL*-atoms as sets of worlds later on. Finally, here is one more important convention in what follows:

Remark Throughout, we will fix announced formulas φ in contexts $\langle !\varphi \rangle \psi$, refraining from varying these in correspondence. Think of distinguished fixed propositions.

Now we are ready to go through the axioms:

Base axiom The axiom $\langle !\varphi \rangle T \Leftrightarrow \varphi$ says that, given any model \mathbf{M} , the domain of the transition relation $R_{[! \varphi]}$ is the set of worlds satisfying φ in \mathbf{M} . In other words, our abstract update action has the truth of φ as a necessary and sufficient precondition.

Disjunction axiom There is no special constraint expressed by the modal formula $\langle !\varphi \rangle (\psi \vee \chi) \Leftrightarrow \langle !\varphi \rangle \psi \vee \langle !\varphi \rangle \chi$, since this law holds for any transition relation.

Negation axiom One direction of this axiom expresses no constraint on the update operation: $(\varphi \wedge \neg \langle !\varphi \rangle \psi) \rightarrow \langle !\varphi \rangle \neg \psi$ is valid, given that is equivalent to $\langle !\varphi \rangle T$. But the converse $\langle !\varphi \rangle \neg \psi \rightarrow (\varphi \wedge \neg \langle !\varphi \rangle \psi)$, even just $\langle !\varphi \rangle \neg \psi \rightarrow \neg \langle !\varphi \rangle \psi$, says by a standard correspondence argument that the transition relation is a *partial function*:

$$\text{if } (\mathbf{M}, s) R_{[! \varphi]} (\mathbf{N}, t) \text{ and } (\mathbf{M}, s) R_{[! \varphi]} (\mathbf{K}, u), \text{ then } (\mathbf{N}, t) = (\mathbf{K}, u).^{11}$$

¹¹ The above comment on interpreting propositions is crucial here: in the argument, we use the singleton set of the pointed model (\mathbf{N}, t) as the denotation of ψ in the update universe \mathbf{M} .

Using this observation, we now simplify the original transition relations R_p in the update universe to *partial functions* F_p on pointed models. In particular, given any model \mathbf{M} with a subset P , we can meaningfully talk about its image $F_p[\mathbf{M}]$.

Knowledge axiom So far, we were just doing preliminaries. The heart of the matter is evidently the recursion axiom for knowledge: $\langle !\varphi \rangle \diamond \psi \leftrightarrow (\varphi \wedge \diamond \langle !\varphi \rangle \psi)$. The two directions of this clearly express two constraints on the update function – and together, they enforce a well-known notion from modal logic (Seegerberg 1971):

Fact The update function satisfies frame truth of $\langle !\varphi \rangle \diamond \psi \leftrightarrow (\varphi \wedge \diamond \langle !\varphi \rangle \psi)$
iff every map F_p is a *p-morphism* between \mathbf{M} and $F_p[\mathbf{M}]$.

Proof We do this first proof in a bit of detail, mainly to show how simple correspondence arguments for update functions are. Consider any model \mathbf{M} , with $[[\varphi]] = P$. First we show that F_p is a homomorphism. Suppose that Rst in \mathbf{M} , with s, t both in the domain of F_p . Now set $V(\psi) = \{F_p(t)\}$. Then $(\mathbf{M}, s) \models \varphi \wedge \diamond \langle !\varphi \rangle \psi$, and therefore also, $(\mathbf{M}, s) \models \langle !\varphi \rangle \diamond \psi$. By the definition of $V(\psi)$, this implies that $R F_p(s) F_p(t)$. Next, for the backward clause of being a *p-morphism*, suppose that $R F_p(s) u$, and now set $V(\psi) = \{u\}$. Then we have $(\mathbf{M}, s) \models \langle !\varphi \rangle \diamond \psi$. It follows from the truth of our axiom that $(\mathbf{M}, s) \models \varphi \wedge \diamond \langle !\varphi \rangle \psi$, and hence there exists a t in \mathbf{M} with Rst and $F_p(t) = u$. ■

Collecting all our observations so far, we have the following result:

Theorem An update universe satisfies the substitution-closed principles of *PAL*
iff its transition relations F_p are partial *p-morphisms* defined on the sets P .

Discussion This is not quite the formation of *submodels* in standard elimination. Here is why. First, having a *p-morphism* is enough for validity of the *PAL* axioms, so we found a generalization of the standard semantics that may be of independent interest. Also, contracting several worlds into one during update occurs naturally in the setting of *PAL*: cf. van Benthem 2011 on the use of bisimulation contractions in updating.¹²

¹² If one insists on making the maps one-to-one, this can be done by enriching the modal language, and enforcing one more reduction axiom for public announcement, namely, for the *difference modality* $D\psi$ saying that ψ holds in a least one different world.

The base axiom once more Still, the above outputs enforced by our update mechanism are relational subframes, rather than submodels. What about the atomic propositions? *PAL* update assumes that these stay the same when a world does not change. Here is how we can think of this. Consider the usual proposition letters of epistemic logic as distinguished atomic propositions. The base axiom tells us that these special propositions have a special behavior: if they hold for an pointed model (\mathbf{M}, s) , they also hold for any of its update images under a map F_p , and vice versa:

$$(\mathbf{M}, s) \models p \text{ iff } F_p(\mathbf{M}, s) \models p$$

This might be the only content to the base axiom: update maps respect distinguished atomic propositions. But we can say a bit more in correspondence style. We assumed that proposition letters ranged over all sets of pointed models in the update universe. Now introduce special ‘context-independent’ proposition letters q ranging only over special sets of pointed models, with the property that they only depend on worlds:

$$(\mathbf{M}, s) \models q \text{ iff } (\mathbf{N}, s) \models q, \quad \text{for all models } \mathbf{M}, \mathbf{N} \text{ in } \mathbf{M}$$

Fact An update universe satisfies the base axiom $\langle !\varphi \rangle q \leftrightarrow (\varphi \wedge q)$ for all context-independent q iff the update maps are the identity on worlds:

$$F_p(\mathbf{M}, s) = (\mathbf{N}, s) \text{ for some model } \mathbf{N}.$$

Proof Consider a pointed model (\mathbf{M}, s) in the domain of F_p . Now set $V(q) = \{(\mathbf{N}, s) \mid (\mathbf{N}, s) \text{ is in } \mathbf{M}\}$. This is clearly a context-independent predicate. Taking this as $V(q)$, the true implication $(\varphi \wedge q) \rightarrow \langle !\varphi \rangle q$ says that $F_p(\mathbf{M}, s) = (\mathbf{N}, s)$ for some model \mathbf{N} . ■

Even so, models \mathbf{N} occurring in F_p -values for pointed models (\mathbf{M}, t) with the same \mathbf{M} could still differ. We will soon see a further recursion law making this uniform.¹³

This concludes our discussion of the correspondence content of the *PAL* axioms.¹⁴

¹³ For an analogy, think of correspondence theory for intuitionistic logic (Rodenburg 1986), where axioms are only valid for all ‘hereditary propositions’.

¹⁴ Readers who like open problems may ponder this: how should the above analysis be modified to allow *factual change*, as in van Benthem, van Eijck & Kooi 2006?

2.3 Variations, extensions, and a provocation

Recursion axioms as general postulates We have determined the update content of one specific axiom for update. But there is more to this. Dynamic-epistemic recursion axioms are not just ‘any sort of principle’. They have several features that make them candidates for general postulates on information update.¹⁵ In particular, our analysis says that the *PAL* recursion axiom for knowledge expresses a sort of *partial bisimulation* between the original model and the output of an update rule applied to it. I find abstract simulation behavior very appealing as a general semantic constraint on update functions, though I am not sure how to define it in its proper generality.¹⁶

Protocols Update universes also suggest a different setting, that has been proposed in dynamic-epistemic logic for independent reasons. So far, we had that $\langle !\varphi \rangle T \leftrightarrow \varphi$. This says that executing an action $!\varphi$ requires truth of the precondition φ , but also, whenever φ is true, $!\varphi$ can be executed. But in civilized conversation or regimented inquiry, the latter assumption is often untenable. To represent this, ‘protocol models’ make restrictions on propositions that can be announced or observed. Hoshi 2009 shows how *PAL* changes in this setting, since the earlier recursion axioms will now be valid only with $\langle !\varphi \rangle T$ in the place of φ on their right-hand sides. This move has many technical repercussions, though the system remains axiomatizable and decidable. From our correspondence perspective, nothing much changes: the only new thing is that the domain of an update map F_p will now be a subset of P , but not necessarily all of P . Our analysis of the modified recursion axioms remains essentially as before.

Language extensions We analyzed update axioms for the epistemic base language. But *PAL* also has a complete version for the full epistemic language with common knowledge. The recursion axiom then requires a new notion of ‘conditional common knowledge’ (van Benthem, van Eijck & Kooi 2006). Since the axiom for single-agent

¹⁵ The commutation of action and knowledge in the key *PAL* recursion axiom has an appealing interpretation in terms of desirable features of logically well-endowed agents. It expresses notions of *Perfect Recall* and *No Miracles* in the sense of Halpern & Vardi 1989.

¹⁶ A relevant analogy here may be with the modal logic of a *bisimulation* Z itself, viewed as a relation on a universe whose worlds are models. The key back-and-forth clause of bisimulation is precisely a commutation axiom $\langle Z \rangle \diamond \psi \leftrightarrow \diamond \langle Z \rangle \psi$.

knowledge already fixed the *PAL* update rule, as we have seen, no further constraints arise. We will return later to what this ‘passive behavior’ of common knowledge vis-à-vis single-agent knowledge means in terms of definability or derivability.¹⁷ A useful language extension whose recursion axiom does add to our correspondence analysis introduces an *existential modality* $E\psi$ saying that ψ is true in some world in the current model, accessible or not. In update universes \mathbf{M} , we interpret this as saying, at a pointed model (\mathbf{M}, s) , that there is some t in \mathbf{M} with ψ true at (\mathbf{M}, t) .

Fact On update universes \mathbf{M} satisfying the earlier *PAL* update conditions, the axiom $\langle !\varphi \rangle E\psi \leftrightarrow (\varphi \wedge E\langle !\varphi \rangle \psi)$ is frame-true iff, for every model \mathbf{M} , the update images of worlds in \mathbf{M} have the same model \mathbf{N} throughout.

Proof First, the axiom is clearly valid in the intended update universes. Conversely, its right to left direction implies the stated property. Consider any two worlds (\mathbf{N}, t) , (\mathbf{K}, u) in the image $F_p(\mathbf{M})$. Set $V(\psi) = \{(\mathbf{K}, u)\}$. Then the F_p -original of (\mathbf{N}, t) in \mathbf{M} satisfies $\varphi \wedge E\langle !\varphi \rangle \psi$. It follows that $\langle !\varphi \rangle E\psi$, and by the preceding definition, this only happens when (\mathbf{N}, t) and (\mathbf{K}, u) share the same model component. ■

Finally, update universes suggest yet further language extensions. For instance, there is also a natural relation $(\mathbf{M}, s) \sim (\mathbf{N}, s)$ holding between different models sharing the same distinguished world. Its modality would make sense, even though it does not make sense inside single epistemic models, the way basic epistemic logic works.

What is the right version of PAL? We conclude with a more provocative feature of our analysis. We started by analyzing what standard public announcement says about update, and then determined its force in update universes. But doing so involved a natural distinction between the substitution-closed principles of *PAL* and the more ‘accidental’ base axiom holding only for a restricted class of valuations. So, what is ‘public announcement logic’ after all? Is its base semantics perhaps the one on update universes with context-dependent propositions and substitution-closed validities? And if so, is what we call the ‘standard version’ perhaps an accident of formulation?

¹⁷ There is also the question whether the recursion axiom for conditional common knowledge by itself fixes world elimination as the update rule – but we will consider this issue only with an analogous case in the dynamic logic of belief change.

2.5 Other natural operations: link cutting

Update with hard information that φ does show variety beyond the above elimination. In a well-known *link-cutting* variant, the operation $\mid\varphi$ performed announces *whether* φ is the case. This means that the domain of worlds stays the same, but all epistemic links get cut between φ -worlds and $\neg\varphi$ -worlds in the current model – an operation used by many authors. The changes induced in the *PAL* axioms are mainly these:

$$\begin{aligned} \langle\mid\varphi\rangle q &\leftrightarrow q \text{ (this implies the substitution-closed instance } \langle\mid\varphi\rangle T) \\ \langle\mid\varphi\rangle\Diamond\psi &\leftrightarrow ((\varphi \wedge \Diamond(\varphi \wedge \langle\mid\varphi\rangle\psi)) \vee (\neg\varphi \wedge \Diamond(\neg\varphi \wedge \langle\mid\varphi\rangle\psi))) \end{aligned}$$

The following result can be proved in the same correspondence style as before:

Fact Link cutting is the only model-changing operation that satisfies the reduction axioms for the dynamic modality $\langle\mid\varphi\rangle$.

Proof We merely give a sketch of the substitution-closed part. Start from any pointed model \mathcal{M}, s . The modified base axiom tells us that the update map is now total on the whole domain of \mathcal{M} . Next, the recursion axiom for knowledge, read from left to right, says that the only links in the image come from already existing links between either φ -worlds, or $\neg\varphi$ -worlds. Finally, from right to left, the axiom says that all links of the two mentioned types existing in \mathcal{M} get preserved into the image. ■

3 Correspondence analysis of modal logics for belief change

Now that we have seen how to analyze principles of knowledge update by changing domains or accessibility relations, an extension to belief revision is straightforward. We mainly need to decide what models we will be working with.

3.1 Soft information and belief

Doxastic models are structures $\mathcal{M} = (W, \{\leq_i\}_{i \in I}, V)$ where the \leq_i are binary comparison relations $\leq_i xy$ saying that agent i considers x at least as plausible as y . As before, for convenience, we drop agent indices henceforth. These plausibility relations are usually taken to be reflexive and transitive, making the modal base logic *S4* – or also connected, like the ‘Grove models’ of belief revision theory, making the logic *S4.3*. Such options are important in practice, but they do not affect the analysis to follow.

These models encode varieties of information. While the whole domain represents our current hard information in the earlier sense, the most plausible worlds in the ordering \leq represent our *soft information* about the actual world. This soft information is the basis of our beliefs and actions based on these, but it is defeasible: the actual world may lie outside of the most plausible area, and we may learn this as a scenario unfolds. In this setting, belief is commonly interpreted as truth in all most plausible worlds:

$$\mathbf{M}, s \models B\varphi \text{ iff } \mathbf{M}, t \models \varphi \text{ for all worlds } t \text{ that are minimal in the ordering } \leq \quad ^{18}$$

But absolute belief does not suffice for most purposes. We need *conditional belief*:

$$\mathbf{M}, s \models B^*\varphi \text{ iff } \mathbf{M}, t \models \varphi \text{ for all } \leq\text{-minimal worlds in } \{u \mid \mathbf{M}, u \models \psi\} \quad ^{19}$$

This point returns with recursion axioms for belief change. From a systematic logical perspective, we should not analyze changes in beliefs only (the usual practice in belief revision theory), but also changes in conditional belief.

Conditional logic Complete logics for conditional belief can be found in close analogy with *conditional logic* based on similarity semantics (Lewis 1973). One difference is that conditional models usually involve a *ternary* comparison ordering $\leq_z xy$: world x is closer to world z than world y . A generalization from binary to ternary relation also makes sense for plausibility semantics of belief, but we forego this here. ²⁰

Safe belief While the preceding belief modalities are interesting, it has become clear recently that the plain base modality of plausibility models has independent interest.

$$\mathbf{M}, s \models \langle\langle\leq\rangle\rangle\varphi \text{ iff there exists a } t \geq s \text{ with } \mathbf{M}, t \models \varphi$$

The corresponding universal modality offers an interesting doxastic notion in between knowledge and belief. Consider this picture with the actual world s in the middle:



¹⁸ We disregard some modifications of truth clauses needed with infinite models.

¹⁹ Absolute belief can be retrieved as the special case of $\psi = T$.

²⁰ Another natural generalization are *epistemic-doxastic models* $\mathbf{M} = (W, \{\sim_i\}_{i \in I}, \{\leq_{i,s}\}_{i \in I}, V)$ allowing for both knowledge update and belief revision. Our methods also work there.

$K\varphi$ describes what we know: φ must be true in all three worlds in the range, less or more plausible than the current one. $B\varphi$ describes beliefs, which have to be true in the right-most world only. Now $[\leq]\varphi$ describes our *safe beliefs*, referring to the actual s plus the right-most world. These cannot be refuted by any future correct observations. Technically, safe belief can also define the other kinds of belief (Boutilier 1994):

on finite pre-orders, $B^*\varphi$ is defined by $U(\psi \rightarrow \langle \leq \rangle (\psi \wedge [\leq](\psi \rightarrow \varphi)))$

with U the universal modality, or in epistemic-doxastic models, a knowledge modality. Thus, an analysis of belief change might focus on safe belief without losing much.

3.2 Dynamic logics of belief change

Now we can write complete logics for belief change. Indeed, there are several systems for this, depending on what kind of new information triggers the change.²¹

Hard information For hard information, the complete dynamic logic is as follows:

Theorem The logic of conditional belief under public announcements is axiomatized completely by (a) any complete static logic for the model class chosen, (b) the *PAL* recursion axioms for atomic facts and Boolean operations, (c) an axiom for conditional belief: $\langle !\varphi \rangle B^*\psi \leftrightarrow (\varphi \wedge B^{\langle !\varphi \rangle} \langle !\varphi \rangle \psi)$.

A similar analysis can be given for safe belief, with a simpler key recursion axiom

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

Formally, this is just the earlier recursion axiom for a modality \diamond .

Soft information and plausibility change Now comes a major further step. Triggers for belief change can be of many kinds, and we do not always expect the same model changes. In particular, incoming new information may be soft rather than hard, which means that it does not eliminate worlds, but merely *rearranges the current plausibility order*. A common example is a *radical upgrade* $\uparrow\varphi$ changing the current ordering \leq between worlds in a model (\mathbf{M}, s) to a new model $(\mathbf{M} \uparrow\varphi, s)$ as follows:

all φ -worlds in the current model become better than all $\neg\varphi$ -worlds,
while, within those two zones, the old plausibility ordering remains.

²¹ The results cited in this subsection and the next are from van Benthem 2007.

Like for public announcement, we introduce an upgrade modality into our language:

$$M, s \models \langle \uparrow\varphi \rangle \psi \quad \text{iff} \quad M \uparrow\varphi, s \models \psi$$

The earlier techniques extend. Again there is a complete set of recursion axioms:

Theorem The dynamic logic of lexicographic upgrade is axiomatized by

- (a) any complete static logic for the model class chosen,
- (b) the following recursion axioms:

$$\begin{aligned} \langle \uparrow\varphi \rangle q &\Leftrightarrow q && \text{for all atomic proposition letters } q \\ \langle \uparrow\varphi \rangle \neg\psi &\Leftrightarrow \neg \langle \uparrow\varphi \rangle \psi \\ \langle \uparrow\varphi \rangle (\psi \vee \chi) &\Leftrightarrow \langle \uparrow\varphi \rangle \psi \vee \langle \uparrow\varphi \rangle \chi \\ \langle \uparrow\varphi \rangle B^a\psi &\Leftrightarrow (E(\varphi \wedge \langle \uparrow\varphi \rangle \alpha) \wedge B^{a \wedge \langle \uparrow\varphi \rangle \alpha} \langle \uparrow\varphi \rangle \psi) \\ &\quad \vee (\neg(E(\varphi \wedge \langle \uparrow\varphi \rangle \alpha) \wedge B^{a \wedge \langle \uparrow\varphi \rangle \alpha} \langle \uparrow\varphi \rangle \psi)) \end{aligned}$$

Again, there is also an evident valid recursion axiom for changes in safe belief:

$$\langle \uparrow\varphi \rangle \langle \leq \rangle \psi \Leftrightarrow E(\varphi \wedge \langle \uparrow\varphi \rangle \psi) \vee (\neg\varphi \wedge \langle \leq \rangle \langle \uparrow\varphi \rangle \psi)$$

Given the earlier modal definition of absolute and conditional belief in terms of safe belief, one can even derive the preceding recursion axioms from this one. Other belief change policies can be treated in the same style, using the relation transformers of van Benthem & Liu 2007 or the priority product update of Baltag & Smets 2008.

3.3 Correspondence for axioms of belief change

As before with knowledge, we can now invert the preceding results and use the key recursion axioms as constraints to determine the space of possible update operations. For update operations transforming plausibility relations only, leaving domains of models the same, a more complex correspondence proof than earlier ones shows:

Theorem The recursion axioms of the dynamic logic of radical upgrade hold universally for an update operation on a universe of pointed plausibility models iff that operation is in fact radical upgrade.²²

It is important to realize what is going on here. *AGM*-style postulates on changes in beliefs will not fix the relational transformation: we need to constrain the changes in

²² Here as before, we work with the substitution-closed version of the logic. In particular, the atomic case simplifies to just $\langle \uparrow\varphi \rangle T$: radical upgrade is defined everywhere.

conditional beliefs, since the new plausibility order encodes all of these. A similar analysis works for other revision policies, such as ‘conservative’ belief change. But actually, there is an easier road to such results, closer to earlier arguments.

Theorem Radical upgrade is the only update operation validating
the given recursion axioms for atoms, Booleans plus safe belief.

Proof Suppose that the axiom is valid on a universe of plausibility models. The axiom for atoms tells us in particular that our update function is defined everywhere. Now consider any model (M, s) . From left to right, taking ψ to denote just one world (N, t) with $F_p(M, s) \leq (N, t)$, it follows that (N, t) was either the image of some φ -world in M , or $s \leq u$ in M for some world u mapped to (N, t) , i.e., the new \leq -link came from an old one originating in a $\neg\varphi$ -world. This means that each new relational link comes from the set defined by radical upgrade. That in fact all such links occur in the F_p -image of M follows by similar unpacking of the reverse implication of the axiom. ■

Given this last correspondence result, the earlier more complex ones seem less urgent, since safe belief defines absolute and conditional belief. Indeed, *AGM*-style postulates on ‘safe-belief change’ might be easier conceptually than those for regular belief.²³

3.4 Discussion: generality of the analysis

We have seen how recursion laws in constructive logics of belief change can serve as general postulates to constrain, and almost uniquely fix, possible updates. As before, this relates the *DDL* and *DEL* approaches to modal logics of belief change, softening a contrast that we started out with. Also as before, issues of generality arise. Are the recursion axioms too specific for belief change postulates? Here we repeat our earlier intuition of ‘simulation’ between input and output models of the transformation. One might add that a recursive postulate may itself be philosophically attractive as providing the core ‘dynamic equation’ driving the process of update or revision. Finally, here is an issue more specific to belief. Given the overwhelming variety of belief

²³ Still, it is interesting that recursion axioms for conditional belief fix radical upgrade, too. This might imply further definability and proof-theoretic connections between the various doxastic notions mentioned. If one recursion axiom fixes update, it looks as if others should be derivable in some way. We cannot explore this technical line here.

revision policies, what is the general thrust of correspondence results like ours? We will return to this issue in Section 5, when discussing product update and other general mechanisms replacing separate revision rules by one master rule plus richer input.²⁴

4 Richer formats as a test case

The style of analysis proposed here works on richer semantic formats for update than modal relational models. In this brief digression, we sketch two examples. These will also raise some issues about the scope and limitations of our earlier analysis.

4.1 Event models and product update

While public announcement logic *PAL* is a good pilot system, its restriction to public information makes it unsuitable for analyzing individual differences in observation and communication. A much richer dynamic-epistemic logic for the latter tasks is true *DEL* (Gerbrandy 1999, Baltag, Moss & Solecki 1998). It uses *action models* E that collect events with attached ‘preconditions’, with epistemic uncertainty links between events representing agents’ observational access to what actually happens. Action models have been used to represent a wide variety of triggers for information change. Next, by performing *product update* of an action model E with the current epistemic or doxastic model M one obtains a new updated information model $M \times E$ displaying the right information for all agents involved after the event has taken place.

We assume that the reader knows how *DEL* update works, including its complete set of recursion axioms. We display two of these for later reference – suppressing agent indices as before, and using the letter R to denote the agent’s accessibility relation:

$$\begin{aligned} \langle E, e \rangle T &\leftrightarrow Pre_e \\ \langle E, e \rangle \diamond \psi &\leftrightarrow (Pre_e \wedge \bigvee_{e R f \text{ in } E} \langle E, f \rangle \psi) \end{aligned}$$

This mechanism changes epistemic or doxastic models much more drastically than the earlier world elimination or relation change. In particular, the set $\{(s, e) \mid s \in M, e \in E, M, s \models Pre_e\}$, of worlds in $M \times E$ may grow beyond the size of the initial model M .

²⁴ Here is a more technical issue. We have only analyzed single update mechanisms so far. But some *AGM*-postulates mix update and revision. Can we use modal versions of such postulates to get correspondence results for axioms with two update modalities simultaneously?

Theorem The recursion axioms for the dynamic modality $\langle E, e \rangle \varphi$ of *DEL* determine product update uniquely modulo p -morphism.

The precise sense in which this fact is true will emerge from the following discussion.

Proof sketch As in our study of PAL, we analyze the impact of the DEL recursion axioms on an update universe of epistemic models with an abstract transition relation for the update for the pointed event model (E, e) . The negation axiom of DEL tells us that this is a partial function $F_{E, e}$. This functionality means that we can think of values $F_{E, e}(M, s)$ as pairs (s, e) without loss of information. Next, the substitution-closed base axiom tells us that $F_{E, e}$ is defined on those models (M, s) whose s satisfies the precondition of e in E . Finally, also as before, the recursion axiom for individual knowledge puts constraints on the function $F_{E, e}$. First, if $s R t$ in M , and $e R f$ in E , while $F_{E, e}(M, s), F_{E, e}(M, t)$ are both defined, then $(s, e) R (t, f)$ holds by the direction from right to left in the axiom. Vice versa, any link in the image of the model M must also arise in this way, if we unpack the left-to-right direction of the axiom.²⁵ ■

One update logic to bind them all? The preceding analysis may still be too piecemeal, ignoring a key innovation of *DEL* in the area of constructive update logics. An earlier trend had been to define specific model changes for particular kinds of informational event: ‘announcements that’, link cutting ‘announcements whether’, or more complex types of private information flow, such as sending a *bcc* message over email. One gets different complete logics for each case. But *DEL* changed the game. All relevant structure triggering different updates is put in matching event models E , and the logic for the special case is then a direct instance of the above ‘mother logic’ of $\langle E, e \rangle \varphi$. In this light, characterizing specific update functions may have some value, but the real logical insight is the general product update mechanism. Is this, then, the best constructive counterpart to a postulational approach to update?

Belief and priority update Similar points can be made about belief revision. One can capture complete logics for specific revision policies, as we have shown. But one can also work at the level of product update with ‘plausibility event models’, where agents

²⁵ This argument still ignores some key features of product update, like its use of ordered pairs (s, e) of worlds and events by themselves without marking the context s in M, e in E .

now may think it more plausible that one event occurred rather than another. Update works with the priority rule that strict event plausibility overrides prior plausibility:

$$(s, e) \leq (t, f) \text{ iff } (s \leq t \wedge e \leq f) \vee e < f \text{ }^{26}$$

The key recursion axiom for the ‘mother logic’ is given in Baltag & Smets 2006:

$$\langle E, e \rangle \langle \leq \rangle \varphi \leftrightarrow (Pre_e \wedge (\bigvee_{e \leq f \text{ in } E} \langle \leq \rangle \langle E, f \rangle \varphi \vee (\bigvee_{e < f \text{ in } E} E \langle E, f \rangle \varphi)) \text{ }^{27}$$

We will not analyze this approach further, but this seems the most general dynamic-epistemic counterpart to the postulational approach of dynamic doxastic logic.²⁸

4.2 Updating neighborhood models for evidence

It is hard to roam for long in modal logic without finding Krister Segerberg’s traces. Another long-standing interest of his are *neighborhood models* (Segerberg 1971) that have been used recently as a model for the epistemological notion of *evidence* and its dynamics (cf. van Benthem & Pacuit 2011 for technical details of what follows).

Static neighborhood logic An epistemic accessibility relation encodes an agent’s current range of worlds after some history of informational events. If we want to retain some of the latter ‘evidence’, a set of neighborhoods (sets of worlds) does well – where we think of the current range as the intersection of all evidence sets.²⁹ The simplest neighborhood models, and all that we consider here, have just one family \mathbf{N} of sets on a domain of worlds. We then interpret an evidence modality as follows:

$$\mathbf{M}, s \models \Box \varphi \text{ iff there is a set } X \text{ in } \mathbf{N} \text{ with } \mathbf{M}, t \models \varphi \text{ for all } t \in X$$

The base logic of this notion is that of a monotone modality that does not necessarily

²⁶ As an illustration, an event model with two signals $!\varphi, !\neg\varphi$, with the first more plausible than the second, generalizes the above radical upgrade $\uparrow\varphi$, that typically also had this overruling character for worlds that satisfied the distinguished triggering proposition φ .

²⁷ Here E is the earlier existential modality over all worlds in the model, accessible or not.

²⁸ Other ways of achieving generality in constructive update logics include the *PDL*-style *program format* of van Benthem & Liu 2007, specifying intended relation changes in models. Girard, Liu & Seligman 2011 define a merge of action models and programs that represents realistic social scenarios. We leave a correspondence analysis to another occasion.

²⁹ If not all given sets overlap, we need more subtle views of conflicting evidence.

distributive over either disjunction or conjunction. This generalization of modal logic supports correspondence analysis.³⁰ Neighborhood models support many epistemic notions. At least in finite models, one can define (cautious evidence-based) *belief* as what is true in all intersections of maximally overlapping families of evidence.³¹

Evidence dynamics: two samples In this setting, our pilot system *PAL* for information update can be seen as mixing different update actions into its public announcements $!\varphi$. The first is *evidence addition* $+\varphi$, adding the denotation $[[\varphi]]$ in the current model as one more piece of evidence to the current evidence family \mathbf{N} . The dynamic logic of this action can be determined completely. Here is one key recursion axiom:

$$\langle +\varphi \rangle \Box \psi \leftrightarrow (\Box \langle +\varphi \rangle \varphi \vee U(\varphi \rightarrow \psi))$$

Again, the content of this principle can be determined by a correspondence argument:

Fact An abstract update function on a universe of neighborhood models satisfies the recursion axiom for evidence addition iff each new evidence set is a superset of either some old evidence set or of the set $[[\varphi]]$.³²

A second aspect of a public announcement $!\varphi$ that now gets into its own is *removal of the evidence* for $\neg\varphi$. The general new operation $-\psi$ removes all evidence sets from the current family \mathbf{N} that are included in $[[\psi]]$. Complete recursion axioms are known for removal and the evidence modality, as well as belief, though a considerable extension of the standard static modal base languages over evidence models is required. Here is one such principle, using a notion of evidence conditional on $\neg\varphi$ being true:

$$\langle -\varphi \rangle \Box \psi \leftrightarrow (E\neg\varphi \rightarrow \Box_{\neg\varphi} \langle -\varphi \rangle \psi)$$
³³

³⁰ For instance, the *K*-axiom $\Box \bigwedge_i \psi_i \leftrightarrow \bigwedge_i \Box \psi_i$ forces \mathbf{N} to be generated from a binary accessibility relation - provided we read it with an infinitary conjunction.

³¹ There are links with modeling beliefs in relational plausibility models here that we ignore.

³² Recursion axioms for new beliefs under evidence addition extend the base language for evidence models to *conditional belief* in two basic varieties that had not surfaced so far.

³³ This is remarkable, since dealing with operations of contraction or removal has long been considered a stumbling block to constructive update logics. The reason why it works in the neighborhood setting after all is the richer model structure one is working on.

We leave a correspondence analysis of recursion axioms for removal to future work. Clearly, we have only scratched the surface here, but hopefully, the reader has seen that our analysis still makes sense when the semantic modeling of dynamic epistemic logic undergoes a drastic neighborhood extension of a sort that Krister Segerberg has long ago proposed for dynamic doxastic logic (Segerberg 1995, Girard 2008).

5 Further directions

We have shown how modal correspondence brings together the postulational format of *AGM* theory and dynamic doxastic logic with the constructive model transformation style of dynamic-epistemic logic. Our technical illustrations were very simple, and we opened up more new problems than closing old ones. Several technical and conceptual issues were already raised in the text. In this section we briefly mention a few more.

Extended semantic formats We have worked with binary accessibility relations for knowledge and belief. This analysis should be extended to ternary relational models, where plausibility can be world-dependent. Likewise, the analysis needs to be taken to the realm of neighborhood models, a natural finer modeling for belief and evidence.

Group knowledge and belief At the start of this paper, we said that a multi-agent perspective is crucial to *DEL*-style logics, but soon this social aspect vanished. One should also analyze update postulates for common knowledge or belief in our style.³⁴

‘Dancing with the stars’: propositional dynamic logic Common knowledge or belief go beyond the modal base language, being iterated modalities as found in dynamic logic *PDL*: another lifelong interest of Krister Segerberg. Iteration occurs naturally in dynamic-epistemic logic, also in the dynamic action component, as with repeated announcement or measurement. The resulting logical systems can be highly complex: cf. Miller & Moss 2004 on *PAL* with iteration, and Baltag & Smets 2009 on limit phenomena with iterated radical update. Still *PDL* is no obstacle to our analysis. There have been some striking advances in the treatment of modal frame correspondence for

³⁴ No complete dynamic logic has been given yet for changes in common belief produced by radical upgrade. Technical difficulties here might require a redesign of the base language to an analogue of the ‘epistemic *PDL*’ of van Benthem, van Eijck & Kooi 2006, a system defined for the purpose of stating recursion axioms for common knowledge with product update.

non-first-order principles like Löb's Axiom for provability logic of Segerberg's Axiom for dynamic logic, making them fall under an extended Sahlqvist syntax matching the system $LFP+FO$, first-order logic with added fixed-point operators. New results and references are found in van Benthem, Bezhanishvili & Hodkinson 2011.

Temporal setting and procedural information Both dynamic doxastic logic and *DEL* focus on single update steps. But equally essential is the temporal horizon. We make sense of local event in terms of global scenarios: a conversation, a process of inquiry, or a game. This 'procedural information' (Hoshi 2009) suggests interfacing dynamic logics with temporal logics of knowledge and belief (Parikh & Ramanujam 2003, Belnap et al. 2001, Bonanno 2007). Existing results at interface take the form of *representation theorems* for 'update evolution': cf. van Benthem, Gerbrandy, Hoshi & Pacuit 2009. One obvious question is how our correspondence results relate to representation theorems in the area of logics of belief. (cf. Dégrémont 2009).

General model theory The proofs in this paper were very simple. The recursion axioms all had Sahlqvist syntax (Blackburn, de Rijke & Venema 2000). One would like a correspondence analysis of axioms for belief change at the latter level of generality. Moreover, correspondence is not the only abstract analysis of concrete modal logics. The mechanism of model change behind the dynamic-epistemic logics in this paper invites reflection on their general features as modal logics. In an earlier book for Krister Segerberg, I gave a Lindström Theorem capturing basic modal logic in terms of bisimulation invariance and compactness. It would be of interest to take this further to capture the essentials of dynamic modal logics of model change.

Coda: have we really dealt with all logics of belief change? Do our two protagonists of dynamic-doxastic and dynamic-epistemic logic exhaust the field? My first attempt at doing modal logic of belief revision in van Benthem 1988 worked over a universe of information stages in the style of Beth or Kripke models for intuitionistic logic. An update with hard information was defined as a *minimal upward move* to a stage where the new information holds, while revision involved backtracking to the past and then going forward again to incorporate new information in conflict with what we thought so far. I am not sure how this third view relates to either *DDL* or *DEL*, though it, too, offers abstract spaces for a wide array of update actions.

6 Conclusion

We have shown how the two main logic approaches to belief change, Segerberg's dynamic doxastic logic and the *DEL* tradition, co-exist in the perspective of modal frame correspondence. Indeed, 'modal logic of belief revision' has two dual aspects that belong together. This much was our contribution to translatability and interaction between frameworks. Our evidence was a set of very simple technical observations – but around these, many new problems came to light. To me, this agenda of unknowns seems a virtue of the proposed analysis. Krister and I have our work cut out for us. Finally, a confession is in order. In starting this study, I thought the main beneficiary would be *DDL*, as it could now import new ideas from the pressure-cooker of *DEL*. But as will be clear at various places in the paper, I now feel that a correspondence perspective also raises serious issues about best design for dynamic-epistemic logics, rethinking their striking deviant feature of being non-substitution-closed. And so, I submit that both sides benefit from the style of analysis presented here.

7 References

- G. Aucher, 2004, *A Combined System for Update Logic and Belief Revision*, Master of Logic Thesis, ILLC, University of Amsterdam.
- 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.
- A. Baltag & S. Smets, 2008, 'A Qualitative Theory of Dynamic Interactive Belief Revision', in G. Bonanno, W. van der Hoek, M. Woolridge, eds., *Texts in Logic and Games Vol. 3*, Amsterdam University Press, 9–58.
- A. Baltag & S. Smets, 2009, 'Group Belief Dynamics under Iterated Revision: Fixed Points and Cycles of Joint Upgrades', *Proceedings TARK XII*, Stanford, 41–50.
- 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., *Logic Colloquium. Granada 1987*, North-Holland, Amsterdam, 331–375.

- J. van Benthem, 2007, ‘Dynamic Logic of Belief Revision’, *Journal of Applied Non-Classical Logics* 17, 129–155.
- J. van Benthem, 2010, *Modal Logic for Open Minds*, CSLI Publications, Stanford.
- J. van Benthem, 2011, *Logical Dynamics of Information and Interaction*, Cambridge University Press, Cambridge.
- J. van Benthem, G. Bezhanishvili & I. Hodkinson, 2011, ‘Sahlqvist Correspondence for Modal μ -Calculus’, *Studia Logica*, to appear.
- J. van Benthem, J. van Eijck & B. Kooi, 2006, ‘Logics of Communication and Change’, *Information and Computation* 204, 1620–1662.
- J. van Benthem, J. Gerbrandy, T. Hoshi & E. Pacuit, 2009, ‘Merging Frameworks for Interaction’, *Journal of Philosophical Logic* 38, 2009, 491–526.
- J. van Benthem & F. Liu, 2007, ‘Dynamic Logic of Preference Upgrade’, *Journal of Applied Non-Classical Logics* 17, 157–182.
- J. van Benthem & E. Pacuit, 2011, ‘Dynamic Logic of Evidence-Based Beliefs’, *Studia Logica* 99:1, 61–92,
- P. Blackburn, M. de Rijke & Y. Venema, 2000, *Modal Logic*, Cambridge University Press, Cambridge.
- G. Bonanno, 2007, ‘Axiomatic Characterization of the AGM theory of Belief Revision in a Temporal Logic’, *Artificial Intelligence* 171, 144–160.
- C. Boutilier, 1994, ‘Conditional Logics of Normality; A Modal Approach’, *Artificial Intelligence* 68, 87–154.
- C. Dégrémont, 2010, *The Temporal Mind: Observations on Belief Change in Temporal Systems*, Dissertation, ILLC, University of Amsterdam.
- H. van Ditmarsch, 2005, ‘Prolegomena to Dynamic Logic for Belief Revision’, *Synthese* 147, 229–275.
- H. van Ditmarsch, W. van der Hoek & B. Kooi, 2007, *Dynamic Epistemic Logic*, Cambridge University Press, Cambridge.
- R. Fagin, J. Halpern, Y. Moses & M. Vardi, 1995, *Reasoning about Knowledge*, The MIT Press, Cambridge (Mass.).
- P. Gärdenfors, 1988, *Knowledge in Flux*, The MIT Press, Cambridge, Mass.
- J. Gerbrandy, 1999, *Bisimulations on Planet Kripke*, Dissertation, ILLC, University of Amsterdam.

- P. Girard, 2008, *Modal Logics for Belief and Preference Change*, dissertation, Department of Philosophy, Stanford University (ILLC-DS-20008-04).
- P. Girard, F. Liu & J. Seligman, 2011, ‘A Product Model Construction for *PDL*’, Departments of Philosophy, Auckland University, and Tsinghua University.
- J. Halpern & M. Vardi, 1989, ‘The Complexity of Reasoning about Knowledge and Time, I: Lower Bounds’, *Journal of Computer and System Sciences* 38, 195–237.
- W. Holiday, T. Hoshi & Th. Icard, 2011, ‘Schematic Validity in Dynamic Epistemic Logic: Decidability’, in H. van Ditmarsch, J. Lang & S. Ju, eds., *Proc’s LORI-III, Guangzhou*, Springer Lecture Notes in Computer Science 6953, 87–96.
- T. Hoshi, 2009, *Epistemic Dynamics and Protocol Information*, Ph.D. thesis, Department of Philosophy, Stanford University (ILLC-DS-2009-08).
- H. Leitgeb & K. Segerberg, 2007, ‘Dynamic Doxastic Logic: why, how, and where to?’, *Synthese* 155, 167–190.
- D. Lewis, 1973, *Counterfactuals*, Blackwell, Oxford.
- F. Liu, 2011, *Reasoning about Preference Dynamics*, Synthese Librray, Springer, Heidelberg.
- J. Miller & L. Moss, 2005, ‘The Undecidability of Iterated Modal Relativization’, *Studia Logica* 97, 373–407.
- R. Parikh & R. Ramanujam, 2003, ‘A Knowledge-Based Semantics of Messages’, *Journal of Logic, Language and Information* 12, 453–467.
- P. Rodenburg, 1986, *Intuitionistic Correspondence Theory*, dissertation, Mathematical Institute, University of Amsterdam.
- H. Rott, 2007, ‘Information Structures in Belief Revision’, in P. Adriaans & J. van Benthem, eds., *Handbook of the Philosophy of Information*, Elsevier Science Publishers, Amsterdam, 457–482.
- K. Segerberg, 1971, *An Essay in Classical Modal Logic*, Philosophical Institute, University of Uppsala.
- K. Segerberg, 1995, ‘Belief Revision from the Point of View of Doxastic Logic’, *Bulletin of the IGPL* 3, 534–553.
- K. Segerberg, 1999, ‘Default Logic as Dynamic Doxastic Logic’, *Erkenntnis* 50, 333–352.