
Dynamic logic for belief revision

Johan van Benthem

Institute for Logic, Language and Computation (ILLC), University of Amsterdam, Netherlands and Department of Philosophy, Stanford University, USA.

E-mail: johan@science.uva.nl

ABSTRACT. We show how belief revision can be treated systematically in the format of dynamic-epistemic logic, when operators of conditional belief are added. The core engine consists of definable update rules for changing plausibility relations between worlds, which have been proposed independently in the dynamic-epistemic literature on preference change. Our analysis yields two new types of modal result. First, we obtain complete logics for concrete mechanisms of belief revision, based on compositional reduction axioms. Next, we show how various abstract postulates for belief revision can be analyzed by standard modal frame correspondences for model-changing operations.

KEYWORDS: dynamic epistemic logic, belief revision, conditional belief, compositionality

1. Information Update and Belief Revision

Belief revision theory in *AGM* style ([GAR 87], [GAR 95]) and dynamic-epistemic logic of information change (*DEL*; [BAL 98], [BEN 06d]) are two major manifestations of the ‘Dynamic Turn’ in logic, making a wide range of informational processes an explicit part of logical systems. There is an obvious issue of comparison between the two paradigms - and the aim of this paper is to provide a joint perspective.

But there are obstacles to such a merge. First, *AGM* analyzes belief change without committing to any fixed mechanism, providing just abstract postulates on the process. By contrast, *DEL* deals with concrete update procedures that change models, and finds complete logics encoding their particular properties. Also, *AGM* deals with single agents and factual information only, while *DEL* is about interaction between many agents, and it typically includes higher-order information about what others know, believe, or what not. And finally, *DEL* explicitly analyzes the ‘triggers’ for information change, from simple announcements of facts to complex information-carrying events. By contrast, *AGM* and its follow-up logics do not explicitly analyze the events that produce belief changes, focusing instead on three particular operations $+A$ (*update*), $*A$ (*revision*), $-A$ (*contraction*) whose completeness as a repertoire of epistemic actions is left open.

Despite these *prima facie* differences, the two styles of logical dynamics, information update and belief revision, can interact and even be integrated - by exploiting the generic character of the *DEL* methodology. To make this point, I will mainly look at two very simple triggers that change agents' information and beliefs, viz. 'hard facts' and 'soft facts'. Before introducing these, however, first consider the following intuitive distinction in the usual models for epistemic and doxastic languages.

Some propositions may be called *knowledge* in the sense that an agent considers them well-established truths. Other propositions represent more volatile *beliefs* that can change as new information comes in. One need not view this in heavy philosophical terms. Rather think of simple scenarios like this. The cards have been dealt. I know that there are 52 of them, and I know their colors. But I have only ephemeral beliefs about who holds which card, or about how the other agents will play. Of course, I could even be wrong about the cards (perhaps someone replaced the King of Hearts by Bill Clinton's visiting card), but this worry seems morbid, and not very useful in understanding normal information flow. Corresponding to this distinction, different events can trigger changes in my models. An incoming public announcement $!P$ of a fact P is a case of *hard information*, which changes what I know. If I see that the Ace of Spades is played on the table, I come to know that no one holds it any more. This is the sort of trigger that drives current logics of information update and knowledge change - as explained in Section 2 below, which outlines the basics of *DEL*. In addition, of course, hard information may also change current beliefs - and Section 3 provides a complete logical system for this.

But next, there is also *soft information*, which affects my beliefs without affecting my knowledge about the cards. I see you smile. This makes it more likely that you hold a trump card, but it does not rule out that you have not got one. Section 4 is about such soft informational actions $*P$ and the resulting belief changes for agents. These effects are produced by changing the 'plausibility relations' between worlds in the relevant static models, which support standard operators of absolute and conditional belief. Again we provide complete dynamic logics, this time for several revision policies. Taken together, these results show that particular belief revision policies can be axiomatized completely in the *DEL* style. Section 5 then reverses the perspective from 'below' to 'above'. We look at abstract postulates for belief revision, and show how these can be analyzed by the standard technique of 'modal frame correspondences', constraining possible model-changing operations. In particular, we show how *DEL* itself gives rise to such correspondence analysis, providing a new look at what its axioms say precisely about models and agents. Many further issues arise once these links have been established. In particular, there are also more complex events involving mixtures of hard and soft information, but to do justice to these, we need to get into the fine-structure of event models and *DEL*-style update. Section 6 provides an outline, while also discussing several further issues. Finally, Section 7 states our conclusions and concerns.

2. Dynamic Logic of Public Announcements

2.1. Standard epistemic logic

The syntax of epistemic logic has a classical propositional base with modal operators $K_i\varphi$ (' i knows that φ ') and $C_G\varphi$ (' φ is common knowledge in group G ')

$$p \mid \neg\varphi \mid \varphi \vee \psi \mid K_i\varphi \mid C_G\varphi.$$

We write $\langle i \rangle\varphi$ for the dual modality $K_i\varphi$: 'agent i considers φ possible'. The dual of $C_G\varphi$ is written $\langle C_G \rangle\varphi$. Models \mathcal{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. One often takes these relations to be equivalence relations, though this is optional here. The epistemic truth conditions are as follows:

$$\mathcal{M}, s \models K_i\varphi \text{ iff for all } t \text{ with } s \sim_i t : \mathcal{M}, t \models \varphi;$$

$$\mathcal{M}, s \models C_G\varphi \text{ iff for all } t \text{ that are reachable from } s \text{ by some finite sequence of } \sim_i \text{ steps } (i \in G): \mathcal{M}, t \models \varphi.$$

For complete epistemic logics over various model classes, see the standard literature (cf. [FAG 95]).

2.2. Public announcement as world elimination

Public announcements of true propositions P change the current model as follows:

For any model \mathcal{M} , world s , and formula P true at s , $(\mathcal{M}|P, s)$ (\mathcal{M} relativized to P at s) is the submodel of \mathcal{M} whose domain is the set $\{t \in \mathcal{M} \mid \mathcal{M}, t \models P\}$.

As shown in Figure 1, one goes from \mathcal{M} to $\mathcal{M}|P$:




Figure 1.

Crucially, truth values of formulas may change in such an update step: most notably, because agents who did not know that P now do after the announcement. This truth value change can be quite subtle over time. Therefore, it is useful to keep track of it systematically in a suitable logical formalism. The language of *public announcement logic* *PAL* is the epistemic language with added action expressions:

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

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

$$\mathcal{M}, s \models [!P]\varphi \quad \text{iff} \quad \text{if } \mathcal{M}, s \models P, \text{ then } \mathcal{M}|P, s \models \varphi.$$

Here is the complete logical calculus of information flow under public announcement (cf. [GER 99], [PLA 89]):

THEOREM 1. — *PAL without common knowledge is axiomatized completely by the usual laws of epistemic logic plus the following reduction axioms:*

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

$$[!P]\neg\varphi \leftrightarrow (P \rightarrow \neg[!P]\varphi)$$

$$[!P](\varphi \wedge \psi) \leftrightarrow ([!P]\varphi \wedge [!P]\psi)$$

$$[!P]K_i\varphi \leftrightarrow (P \rightarrow K_i[!P]\varphi)$$

EXAMPLE 2 (SOUNDNESS OF REDUCTION AXIOMS). — We do the crucial final case of knowledge after announcement. This compares two models: (\mathcal{M}, s) and $(\mathcal{M}|P, s)$ before and after the update. It helps to draw pictures relating these to understand the following proof. The formula $[!P]K_i\varphi$ says that, in $\mathcal{M}|P$, all worlds \sim_i -accessible from s satisfy φ . The corresponding worlds in \mathcal{M} are those worlds which are \sim_i -accessible from s and which satisfy P . Moreover, given that truth values of formulas may change in an update step, the correct description of these worlds in \mathcal{M} is not that they satisfy φ (which they do in $\mathcal{M}|P$), but rather $[!P]\varphi$: they become φ after the update. Finally, $!P$ is a partial operation, as P has to be true for its public announcement. Thus, we need to make our assertion on the right conditional on $!P$ being executable, i.e., P being true. Putting all this together, $[!P]K_i\varphi$ says the same as $P \rightarrow K_i(P \rightarrow [!P]\varphi)$. But given the effect of the operator $[!P]$ for a partial operation, we can simplify this final formula to the equivalent $P \rightarrow K_i[!P]\varphi$. \square

This type of argument is at the same time a heuristic analysis of a reductive situation, and it explains all further reduction axioms that we will find in what follows.¹

These elegant axioms analyze reasoning about effects of getting hard information, through observation, communication, or other reliable means. There are two major features to this approach. First, the analysis is *compositional*, breaking down the ‘postconditions’ behind the dynamic modalities $[!P]$ recursively. Next, the dynamic ‘reduction axioms’ take every formula of our dynamic-epistemic language eventually

1. *PAL* was designed to reason about what people tell each other, and it is quite successful in that. Yet it has no explicit axioms relating different agents. The ‘social’ character only shows in its syntax of complex formulas with iterations. It is highlighted much more by new notions of group knowledge: cf. the discussion of common knowledge later on.

to an equivalent formula inside the static pure epistemic language. In terms of models, this means that the current static model already contains all information about what might happen when agents communicate what they know. This feature places a constraint on the static base language: it has to be rich enough to allow for *pre-encoding* - just as, e.g., conditionals pre-encode tendencies toward future actions of belief revision. Phrased in a slogan: ‘The epistemic/doxastic present already contains the epistemic future’. In terms of the logic, the reduction procedure means that *PAL* is *decidable*, since this is true for its static epistemic base language. There is much more to *PAL*, including a bisimulation-based model theory, and issues of expressive power and computational complexity. Some of this will be relevant below. Cf. [BEN 06b] for a survey of open problems.

2.3. The *DEL* methodology

Theorem 1 demonstrates the general *DEL* methodology in a nut-shell, as it can be used, in principle, to ‘dynamify’ any given logical system. First, one chooses a static language and matching models that represent information states for groups of agents. Next one analyzes the relevant information-carrying events as updates changing these models. These events are then described explicitly in a dynamic extension of the language, which can also state the effects of events in terms of propositions that hold after their occurrence. This adds a dynamic superstructure over a more traditional substructure. The resulting logics have a two-tier set-up:

static base logic—dynamic extension

At the static level, one gets a complete axiom system for whatever models one has chosen. But on top of that, there is a set of dynamic *reduction axioms* for effects of events. In cases where this works, every formula is equivalent to a static one - and hence, if the static base logic is decidable, so is its dynamic extension.²

In principle, this design of dynamic epistemic logics is modular, and independent from any specific properties of the static models and their language. In particular, the reduction axioms for *PAL* do not depend on any assumption about epistemic accessibility relations. Hence Theorem 1 holds just as well if the underlying models are arbitrary, validating the minimal logic *K*, serving as some minimal logic of *belief*. Indeed, this is how some core texts on *DEL* set up their logic systems from the start - cf. Section 2.5 for further details. Thus, in what follows, we shall concentrate mostly on the dynamic superstructure.

Nevertheless, some interplay between static and dynamic structure does occur.

EXAMPLE 3. — *Preserving frame conditions*

Suppose we impose relational conditions on base models, like epistemic accessibilities

2. This reduction does not settle computational complexity: translation via the axioms may increase formula length exponentially. But in fact (cf. [LUT 06]), for *PAL*, the complexity of satisfiability remains that of epistemic logic, viz. Pspace-complete.

being equivalence relations. This gives a matching constraint on the update mechanism: it should *preserve these frame conditions*. For equivalence relations, and other conditions defined by universal first-order formulas, this is guaranteed by passing to submodels as above. The more general ‘product update’ of Section 6, however, outputs submodels of direct products of given models. In that case, the only first-order frame conditions that are guaranteed to be preserved are those definable as *universal Horn sentences*. Reflexivity, symmetry, and transitivity are indeed of the latter special form, but a non-Horn universal frame condition like *linearity* of an accessibility ordering is not, and may be lost. [KOO 05] has some further technical investigation. \square

EXAMPLE 4. — *Enriching the base language to obtain reduction axioms*

Suppose that we add a new epistemic operator to our base language, say common knowledge. In this case, it turns out that there is no reduction axiom for formulas $[!P]C_G\varphi$. To get one, we need to enrich the standard language of epistemic logic with a new notion of *conditional common knowledge*

$$C_G(P, \varphi),$$

stating that φ is true in all worlds reachable via some finite path of accessibilities running entirely through worlds satisfying P . Plain $C_G\varphi$ is a special case of this, viz. $C_G(\text{True}, \varphi)$. Once we have this new operator, we can formulate the following valid reduction axiom for common knowledge in *PAL*:

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

Note the role of the $[!P]\varphi$ in the consequent here. On the left, we are looking at worlds which satisfy the formula φ in $\mathcal{M}|P$, after the update for $!P$. But these correspond to worlds satisfying the ‘look-ahead formula’ $[!P]\varphi$ in the original model \mathcal{M} . Conditional common knowledge is not definable in the basic epistemic language - but it is bisimulation-invariant, and existing completeness proofs are easily adapted. On this extended base, we have a valid general reduction axiom extending *PAL*. Note that we have a richer base language now, so we must have reduction axioms that work for the new stronger form of common knowledge, not just $C_G\varphi$. The next axiom shows that the hierarchy stops here ([BEN 06d]): \square

THEOREM 5. — *PAL with conditional common knowledge is axiomatized completely by adding the reduction law $[!P]C_G(\varphi, \psi) \leftrightarrow C_G(P \wedge [!P]\varphi, [!P]\psi)$.*

Conditional common knowledge $C_G(P, \varphi)$ is again a way of *pre-encoding*, in the current model, common knowledge that would obtain after the fact P is learnt.

2.4. Some semantic core facts about *PAL*

We now briefly review a few semantic peculiarities of *PAL*, which are essential to understanding later dynamic logics in the same mold.

Changes in truth value and persistence Typically, incoming information does not change atomic facts, but it does change knowledge or ignorance of agents. The result-

ing truth value changes in assertions of the form $K_i\varphi$ or $C_G\varphi$ can be subtle, and the point of the dynamic-epistemic language is precisely to keep track of these. Some true statements even have the perverse feature of becoming false upon their announcement. An example is the Moore-type assertion:

$$\neg K_i p \wedge p \quad \text{'you don't know it, but } p\text{'}$$

Upon its announcement, the fact p becomes common knowledge, thus falsifying the first conjunct. But other assertions have the property that, when truly announced, they do become common knowledge. This holds, e.g., for all formulas of the form

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

In particular, *PAL* has an obvious sublogic where one only announces factual assertions without any epistemic operators. In that case, every announcement $!P$ produces common knowledge of P , and the reduction process gets much simpler.

Conditionals as dynamic modalities The reduction axiom

$$[!P]K_i\varphi \leftrightarrow P \rightarrow K_i[!P]\varphi$$

is actually equivalent to another form, as we saw in our Soundness argument:

$$[!P]K_i\varphi \rightarrow P \rightarrow K_i(P \rightarrow [!P]\varphi)$$

The antecedent P on the right just states the precondition for a true announcement of P . The rest of the axiom then says that the following two perspectives are equivalent:

- (a) knowing that φ once we have added the information that P ,
- (b) knowing the conditional that P implies ' φ ', where again, $[!P]\varphi$ describes φ 's truth after the update.

Here, and later on, the distinction between φ and $[!P]\varphi$ is negligible as long as φ is a non-epistemic factual statement. In that case, axioms reduce to simpler versions.³

The idea that a conditional $A \Rightarrow B$ resembles a dynamic modality $[!A]B$ is old folklore. In the *PAL* setting, two obvious principles then come straight from the minimal modal axioms:

FACT 6. — The following laws are valid for the dynamic announcement conditional:

- (a) Conjunction of Consequents ($A \Rightarrow B, A \Rightarrow C$ imply $A \Rightarrow B \wedge C$)
- (b) Upward Monotonicity in the Consequent ($A \Rightarrow B$ implies $A \Rightarrow B \vee C$).

But counter-examples exist with concrete dynamic update models for

- (c) Reflexivity, Downward Monotonicity in Antecedents, Disjunction of Antecedents, Cautious Transitivity, and Cautious Monotonicity.

3. There is an analogy here with the 'Ramsey Test' for conditionals in terms of belief revision. But, our axioms rather ask to which extent a conditional assertion in the changed model can still be defined in terms of some related conditional true in the original model.

Even so, there are also some new conditional validities in the system, witness the non-standard but intriguing conditional law $(A \Rightarrow (B \Rightarrow C)) \leftrightarrow ((A \Rightarrow B) \Rightarrow C)$ which would correspond to the iteration principle in the following subsection. As mentioned before, with just non-epistemic antecedents and consequents, update changes no truth values, and the conditional is an ordinary modal implication. We leave the complete logic of these dynamic update conditionals as an open question.⁴

Iteration Assertions can be iterated to form longer conversations, games, etcetera. The language of *PAL* describes this by stacking modal operators, as in

$$[!A][!B]\varphi$$

But the logic has an interesting valid principle saying that the effect of two consecutive assertions can also be achieved by making just one:

FACT 7. — $[!A][!B]\varphi \leftrightarrow [!(A \wedge [A!]B)]\varphi$ is a valid principle of *PAL*.

Indeed, this principle is *schematically valid*, in that each of its substitution instances is also a validity. Schematic validity is not a feature of all *PAL*-axioms, however, witness the earlier reduction axiom for atomic propositions

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

which definitely does not work when we replace q by an arbitrary epistemic formula. But the reduction axioms for the logical operations are schematically valid. It is not known if the schematic validities of *PAL* are axiomatizable, let alone decidable.

The natural next step in studying iteration would be to allow complex instructions for conversation, using three well-known operations on computer programs:

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

We crucially say things in a certain order, what we say may depend on circumstances, and we may have to keep repeating assertions until some intended effect obtains, as in flattery or threats. This richer language of conversation has a simple syntax and semantics, resembling that of propositional dynamic logic *PDL* - and it is still like *PAL* in crucial ways. E.g., its formulas are all invariant for epistemic bisimulation. But there is a surprise in terms of the complexity of validity ([MIL 05]):

THEOREM 8. — *PAL* with all *PDL* program operations added to the action part of the language is undecidable, and even non-axiomatizable.

4. With the full language of *PAL* again, [BEN 03] provides a complete description of the abstract structural rules that are valid for dynamic inference.

2.5. From knowledge to belief

We have described the logic of public announcement in terms of *knowledge*. While this is convenient for some examples, it also has a disadvantage, as it may suggest that the approach is peculiar to knowledge. This is not so at all. Everything we have said about *PAL* works just as well when we read the $K_i\varphi$ as operators of *belief*:

$$[!P]B_i\varphi \leftrightarrow P \rightarrow B_i[!P]\varphi$$

Indeed, for most applications of the framework, as we have noted before, the best reading of the relevant epistemic operator may be something like this:

“to the best of my information ...”

In this case, we simply drop the requirement that accessibility for agents should be an equivalence relation. For instance, the following generalized model shows how a fact p can be true while the agent believes mistakenly that $\neg p$:




Figure 2.


With this view of doxastic modalities, the whole machinery of *DEL* works exactly as before. In the next Section, we analyze the belief version a bit more in detail, adding some fine-structure. Later on, we will also look at update systems where we have *two modal operators*: one stricter for knowledge, one more easy to satisfy for belief. This set-up does not seem to add deep new issues (though compare the discussion in [DIT 06]), but it is a great convenience in practice.

We have discussed the simple dynamic-epistemic logic of public announcements in great detail. This is not so much for its intrinsic importance, but as an illustration of our general methodology, and the logical issues that it raises. In Section 6, we will take a brief look at more complex *DEL*-style systems with arbitrary informational events and the more complex mechanism of ‘*product update*’ (cf. [BAL 98], [DIT 07], [BEN 06d]). But the simple background of public announcement suffices for our main goal in this paper: the intended extensions to belief revision, which we will now develop at much greater speed.

3. Belief Change under Hard Information

3.1. A problem with eliminative belief revision

Redraw the preceding belief model, now with knowledge and belief combined. In the actual world x , p is the case, and I do not know if p , but I believe that $\neg p$:

**Figure 3.**

Now here is a problem with eliminative update: *PAL* as it stands does not do true belief revision. A ‘hard announcement’ $!p$ of the real situation would turn this initial situation into the one-world model $\{x\}$ with an *empty* doxastic accessibility relation - where I believe that p , but even $B\perp$. . . But that is not what we want: I should just come to (know and) believe that p ! Here is a solution.

3.2. World comparison and conditional belief

Models A richer view of belief follows the intuition that we believe those things that hold in the ‘best’ or ‘most relevant’ worlds epistemically accessible to us. I believe that this train will take me home on time, even though I do not know *stricto sensu* that it will not suddenly fly away from the tracks as in “Back to the Future, Part III”. But the worlds where it stays on track are more plausible than those where it flies, and among the latter, those where it arrives on time are more plausible than those where it does not. Static models for this situation are of the form

$$\mathcal{M} = (W, \{\leq_{i,s}\}_{i \in I, V})$$

where the $\leq_{i,s}$ are ternary comparison relations for agents, read as follows,

$\leq_{i,s} xy$: in world s , agent i considers x at least as plausible as y

Models like this have been proposed by many authors, starting with the work of Lewis in conditional logic, all the way to the ‘graded models’ of [SPO 88], and [SHO 88] on generalized preference relations in AI. One can impose several mathematical conditions on the relations $\lambda xy. \leq_{i,s} xy$, depending on their intuitive reading. The minimum found with [BUR 84], and [VEL 85] is *reflexivity* and *transitivity*. [LEW 73] also imposes *connectedness*: worlds either precede each other, or they have the same predecessors and successors. The latter condition yields the well-known geometrical systems of ‘nested spheres’. As before with epistemic models, our dynamic analysis works largely independently from such formal design decisions, important though they may be when fine-tuning to specific applications.

Languages and logics One can interpret many new logical operators in this comparative order structure. In what follows, we choose the intuitive ‘minimality’ formulations, even though these must (and can) be modified somewhat in models allowing infinite descent in the ordering. First of all, there is plain belief:

$\mathcal{M}, s \models B_i \varphi$ iff $\mathcal{M}, t \models \varphi$ for all worlds t which are minimal for the ordering $\lambda xy. \leq_{i,s} xy$.

But the more general notion is that of a *conditional belief*:

$\mathcal{M}, s \models B_i(\varphi|\psi)$ iff $\mathcal{M}, t \models \varphi$ for all worlds t which are minimal for $\lambda xy. \leq_{i,s} xy$ in the set $\{u \mid \mathcal{M}, u \models \psi\}$.

Conditional beliefs *pre-encode* beliefs that we would have *if* we learnt certain things. The formal analogy with conditionals is this. A conditional $C \Rightarrow D$ says that D is true in the minimal worlds where C is true (as measured by some comparison order on worlds). This is exactly the above $B_i(D|C)$. Indeed, on the reflexive transitive models for the conditional language, $B_i(\varphi|\psi)$ satisfies just the axioms of the minimal conditional logic listed in the note below.⁵

REMARK 9. — *Pre-encoding once more*

This is a good moment to take the technical side of ‘pre-encoding’ a bit further. A conditional belief $B_i(\varphi|\psi)$ does not quite tell us what we would believe if we learnt the antecedent. For, the action of learning the antecedent ψ *changes the current model* \mathcal{M} , and hence the truth value of the consequent φ might change. The reason is that the modalities occurring in φ may range over different worlds in the models \mathcal{M} and $\mathcal{M}|\psi$. This is a well-known phenomenon in many areas of logic. E.g., the relativized quantifier in “All mothers have daughters” does not say that, if we relativize to the subset of mothers, all of them have daughters who are mothers themselves.⁶ □

REMARK 10. — *Richer modal languages*

Next, one can also interpret richer modal languages on these models. E.g., the idea of a ‘best’ world really induces a *binary* relation ‘ $best_i$ ’ between worlds s and t :

t is minimal in $\lambda xy. \leq_{i,s} xy$

One could introduce a modality for this in the style of propositional dynamic logic (in conditional logic, this is like having a world dependent ‘selection function’), and the above belief modality $B_i\varphi$ would then be read as follows:

$[best_i]B_i\varphi$

Even more powerful modal preference languages are under development today.⁷ □

5. (a) $A \Rightarrow A$, (b) $A \Rightarrow B$ implies $A \Rightarrow B \vee C$, (c) $A \Rightarrow B, A \Rightarrow C$ imply $A \Rightarrow B \wedge C$, (d) $A \Rightarrow B, C \Rightarrow B$ imply $(A \vee C) \Rightarrow B$, (e) $A \Rightarrow B, A \Rightarrow C$ imply $(A \wedge B) \Rightarrow C$.

6. Standard modal languages can talk about a definable submodel $\{s \text{ in } \mathcal{M} \mid \mathcal{M}, s \models P\}$ by means of syntactic relativization, replacing modalities $\langle \rangle, \square$ by $\langle \rangle(P \wedge \dots), \square(P \rightarrow \dots)$, respectively. But this trick does not always work. E.g., the above unary belief modality defined as truth in all best worlds cannot be relativized using any such simple trick (to prove this impossibility requires some delicate model-theoretic reasoning). In such a setting, one really has to introduce an explicit new operator of conditional belief to help with this purpose.

7. One can be more radical than these still rather traditional approaches. The ‘parametrized’ binary ordering in our models supports a two-dimensional format $\mathcal{M}, s, x \models \langle pref \rangle_i \varphi$ iff for some $y \leq_{i,s} x : \mathcal{M}, y \models \varphi$. [BOU 94], [BEN 06e], [BEN 07], [BEN 06c] show how the resulting modal preference languages extend conditional logic and make its properties much more perspicuous. They also add the right expressive power needed to deal with more complex uses of preference relations, like those found in theories of action and games.

3.3. Dynamic logic of belief change under hard information

Now we are in a position to present our first new dynamic logic of belief revision. It arises from putting together the earlier logic *PAL* with static models for conditional belief. Our first language in this setting again takes public announcements $!P$ of true propositions P . It is a simple exercise to check one immediate reduction axiom:

FACT 11. — The following formula is valid in our semantics for beliefs that are acquired after some hard fact has been announced:

$$[!P]B_i\varphi \leftrightarrow (P \rightarrow B_i([!P]\varphi|P))$$

This is much like the *PAL* reduction law for knowledge under public announcement. This is correct, as the formal structure of the recursion is the same in both cases. Again, to keep the complete dynamic language in harmony, we also need a reduction axiom for conditional beliefs which stays inside the language as we have it now:

THEOREM 12. — *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* reduction axioms for atomic facts and Boolean operations,
- (c) the following reduction axiom for conditional beliefs:

$$[!P]B_i(\varphi|\psi) \leftrightarrow P \rightarrow B_i([!P]\varphi|P \wedge [!P]\psi)$$

PROOF. — This result is obvious, once we see the soundness of the new axiom. On the left hand side, it says that in a model $(\mathcal{M}|P, s)$, φ is true in the best ψ -worlds. With the usual precondition for the announcement, on the right-hand side, it says that in (\mathcal{M}, s) , the best worlds that are P now and will become ψ after announcing that P , will also become φ after announcing P . This is indeed equivalent. ■

3.4. Combining knowledge and belief

It is easy to combine the preceding systems. We take epistemic models as before for the K_i -operators, and think of the plausibility orders \leq_i for absolute or conditional beliefs $B_i\varphi, B_i(\varphi|\psi)$ as ordering the epistemic equivalence classes of \sim_i for each agent i . The extended dynamic language with announcements $!P$ will then record how both knowledge and belief of agents change as hard facts come in.

4. Belief change under soft information

4.1. Revision as relation change

Even in the above version with conditional beliefs, systems of public announcement do not perform genuine belief revision in its most general form. Consider the earlier model in whose actual world x , I believe that $\neg p$, though p is in fact the case:




Figure 4.

An announcement $!p$ turns this into the one-world model $\{x\}$ where I believe that p , but there is no way back. Nothing can *undo* my belief in p , since the $\neg p$ -world has disappeared. Thus, we need a new idea for scenarios where beliefs can change back and forth. Indeed, often, the trigger for belief revision is ‘softer’ than a call for world elimination, introducing just a greater ‘preference’ for p -worlds, without totally abandoning the others. A typical example is the treatment of conditional assertions $A \Rightarrow B$ as *default rules* (cf. [VEL 96]). Accepting a default rule does not say that all A -worlds must now be B -worlds. It rather says that the ‘exceptional’ $A \wedge \neg B$ -worlds are to be considered less plausible, or less relevant. This ‘soft information’ does not eliminate worlds, it rather *changes the ordering* of the existing worlds. For instance, one of Veltman’s rules does the following - stated in our notation, and suppressing agent indices which do not really matter for the idea:

(\uparrow) Lexicographic upgrade

$\uparrow P$ is an instruction for replacing the current ordering relation \leq between worlds by the following: *all P -worlds become better than all $\neg P$ -worlds, and within those two zones, the old ordering remains.*

This is just the lexicographic policy for relational belief revision of [ROT 06]. This move is like a social revolution where some underclass P now becomes the upper class. The outcomes of policy $\uparrow P$ are easy to picture - and lexicographic ordering is indeed a key notion in many theories of preference along dimensions.

Policies and uptake scenarios Belief revision theory allows for different policies of change, as agents may differ in the entrenchment of their old beliefs versus the new ones. Another re-ordering of world plausibility is this. Macchiavellistically, one just co-opts the leaders of the underclass, leaving the further social order unchanged:

(\uparrow) Elite change

$\uparrow P$ replaces the current ordering relation \leq by the following: *the best P -worlds come on top, but apart from that, the old ordering remains.*

One can think of revision policies as persistent habits of an agent over time. But they can also be more local ‘types of response’ to particular inputs. Speech act theories distinguish between incoming information per se (what is said) and the ‘uptake’, the way in which the recipient reacts to them. In that sense, the ‘softness’ of our scenarios for $*P$ might be in the response, rather than in the signal itself.⁸

8. Diversity of policies is sometimes seen as a hallmark of the ‘non-logicality’ - and messiness - of belief revision. But legitimate diversity occurs even with inference and information update, once we consider agents with bounded rationality: cf. [BEN 04], [LIU 06a].

Many sorts of relation change were described in [BEN 93] as a topic for systematic analysis, calling for a ‘dynamification’ of preference logic. [BEN 07] present a *DEL* treatment of general ‘upgrade’ mechanisms, showing how reduction axioms can be read off *automatically* under mild assumptions on definability of the relation change. Many such variants make sense for belief revision, witness the ‘27’ of [ROT 06]. The same is true for defaults, commands ([YAM 06]), and other areas where plausibility or preference can change. [LIU 06a] has a systematic discussion of this diversity, showing how it also arises in other components of dynamic-epistemic logics, reflecting agents’ powers of inference, introspection, and memory.

The next main point of this paper is how the *DEL* format applies to various scenarios of belief change under soft information. To show this, we will axiomatize the dynamic logic of the two key operators $\uparrow P$ and $\uparrow P$. For convenience, we will assume that all our orderings are total, leaving a generalization for later investigation.

4.2. Two complete dynamic logics of belief upgrade

We keep the same base language with conditional belief as before, and merely show what the extra dynamic superstructure for belief revision will look like. Again, in writing principles, we suppress agent subscripts, as this does not affect the key ideas. But we do need one more harmless operator for convenience, viz.

an *existential modality* $E\varphi$ saying that φ holds in some world.

This common device expresses ‘consistency’ of φ as far as relevant to the current model. In a combined epistemic-doxastic language, the role of $E\varphi$ would be played most naturally by the existential epistemic modality $\langle i \rangle \varphi$ for the relevant agent i .

This time, we have the following operations of model change:

- (a) $\mathcal{M}, s \models [\uparrow P]\varphi$ iff $\mathcal{M} \uparrow P, s \models \varphi$
with $\mathcal{M} \uparrow P$ the model \mathcal{M} with its order \leq changed as in (\uparrow) above;
- (b) and completely analogously for a dynamic modality $[\uparrow P]\varphi$.

We think of these operations as functions without any preconditions. Our first new result shows how *DEL* methods axiomatize this kind of belief revision completely:

THEOREM 13. — *The dynamic logic of lexicographic upgrade is axiomatized by*

- (a) *a complete axiom system for conditional belief on the static models, and*
- (b) *the following reduction axioms:*

$$\begin{aligned}
 &[\uparrow P]q \leftrightarrow q, \quad \text{for all atomic proposition letters } q \\
 &[\uparrow P]\neg\varphi \leftrightarrow \neg[\uparrow P]\varphi \\
 &[\uparrow P](\varphi \wedge \psi) \leftrightarrow ([\uparrow P]\varphi \wedge [\uparrow P]\psi) \\
 &[\uparrow P]B(\varphi|\psi) \leftrightarrow (E(P \wedge [\uparrow P]\psi) \wedge B([\uparrow P]\varphi \mid P \wedge [\uparrow P]\psi)) \vee \\
 &\quad (\neg E(P \wedge [\uparrow P]\psi) \wedge B([\uparrow P]\varphi \mid [\uparrow P]\psi))
 \end{aligned}$$

PROOF. — We only comment on the reduction axioms. Generally, the first three look slightly simpler than those for *PAL*, because there is no precondition for $\uparrow P$ as there was for $!P$. The first axiom expresses the fact that upgrade does not change truth values of atomic facts. The second says that the operation of model change is a function. The third is a general feature of any modal-type change operator.

The fourth axiom is the locus where the information has to show about the specific relation change that has been used. It looks forbidding, but it is really easy to grasp. On the left-hand side, we see that after the P -upgrade, all best ψ -worlds satisfy φ . On the right-hand side, there is a case distinction. Case (1): there are P -worlds in the original model that become ψ after the upgrade. In this case, the lexicographic reordering $\uparrow P$ makes the best of these worlds in \mathcal{M} the best ones over-all in $\mathcal{M} \uparrow P$ to satisfy ψ . Now, in the original model \mathcal{M} - viewed from some current s - the worlds of Case 1 are exactly those satisfying the formula $P \wedge [\uparrow P]\psi$. The formula $B([\uparrow P]\varphi \mid (P \wedge [\uparrow P]\psi))$ then says that the best among these in \mathcal{M} will indeed satisfy φ after the upgrade. And these best worlds are the same as those described earlier, as lexicographic reordering does not change the ordering of worlds inside the P -area. Case (2): no P -worlds in the original model become ψ after upgrade. What matters in this case is that the lexicographic reordering $\uparrow P$ makes the best worlds satisfying ψ after the upgrade just the same best worlds over-all as before that satisfied $[\uparrow P]\psi$. Thus, the formula $B([\uparrow P]\varphi \mid [\uparrow P]\psi)$ in the reduction axiom says that the best worlds become φ after upgrade.⁹ ■

Again, the axioms provide a reduction procedure for arbitrary dynamic formulas into pure formulas of the base language - and hence the logic is decidable. Moreover, in an epistemic/doxastic version, it is easy to add a valid reduction axiom for a knowledge operator, as a change in plausibility does not affect epistemic accessibility:

$$[\uparrow P]K_i\varphi \leftrightarrow K_i[\uparrow P]\varphi$$

Simplified versions For someone new to the recursive thinking of *DEL*, it may be instructive to look at some special cases. First, consider unconditional beliefs $B_i\varphi$. In that case, the above reduction axiom simplifies to the following equivalence (setting ψ equal to True):

$$[\uparrow P]B\varphi \leftrightarrow (EP \wedge B([\uparrow P]\varphi \mid P)) \vee (\neg EP \wedge B[\uparrow P]\varphi)$$

This looks much more like the *PAL*-style reduction axiom we had before. On the other hand, it is important to realize that providing the above reduction axioms for *conditional beliefs* is essential if we are to have the system closed under iteration.

Indeed, this is also a moral for *AGM*. Its basic postulates seem to concern acquiring non-conditional absolute beliefs only, while an interaction-closed system should also have intuitive postulates concerning the acquisition of conditional beliefs.¹⁰ In any case, even with the unary standard format, some useful comparisons can be made:

9. An alternative derivation of the axiom is found in [BEN 07], using preference logic.

10. But one might say that the *AGM*-postulate for conjunctions does have an iteration flavor.

REMARK 14. — *Connections with AGM postulates*

How to recognize standard *AGM* postulates in this *DEL* setting? Well, they might be derivable as theorems, insofar as expressible in this language. Consider Success:

$[\uparrow P]BP$, provided that P is consistent

Standard *DEL* wisdom is that this should fail, as the earlier ‘tricky’ epistemic updates, or upgrades, invalidate this intuition. But for factual upgrade with P atomic, Success *should* hold. Here is how, using our first reduction axiom, $[\uparrow P]P \leftrightarrow P$:

$$\begin{aligned} [\uparrow P]BP &\leftrightarrow (EP \wedge B([\uparrow P]P|P)) \vee (\neg EP \wedge B[\uparrow P]P) \leftrightarrow \\ &(EP \wedge B(P|P)) \vee (\neg EP \wedge BP) \leftrightarrow EP \vee BP. \end{aligned}$$

Thus $EP \rightarrow [\uparrow P]BP$ evaluates to *True*.

More generally, for all ‘factual’ statements φ without doxastic or epistemic operators, the difference between $[\uparrow P]\varphi$ and φ disappears. Thus, for such factual statements φ, ψ , the crucial fourth reduction axiom simplifies to:

$$[\uparrow P]B(\varphi|\psi) \leftrightarrow (E(P \wedge \psi) \wedge B(\varphi|(P \wedge \psi))) \vee (\neg E(P \wedge \psi) \wedge B(\varphi|\psi))$$

This is the precise sense in which a *Ramsey Test* as in conditional logic holds for our *DEL*-style logic of upgrade - but we will not pursue this matter here. \square

Theorem 14 is no accident. A similar analysis works for the other upgrade operation $\uparrow P$. This time, we leave it to the reader to verify the crucial reduction axiom.

THEOREM 15. — *The dynamic logic of conservative upgrade is axiomatized completely by (a) a complete axiom system for conditional belief on the static models, and (b) the following reduction axioms:*

$$\begin{aligned} [\uparrow P]q &\leftrightarrow q, \quad \text{for all atomic proposition letters } q \\ [\uparrow P]\neg\varphi &\leftrightarrow \neg[\uparrow P]\varphi \\ [\uparrow P](\varphi \wedge \psi) &\leftrightarrow ([\uparrow P]\varphi \wedge [\uparrow P]\psi) \\ [\uparrow P]B(\varphi|\psi) &\leftrightarrow (B(\neg[\uparrow P]\psi|P) \wedge B([\uparrow P]\varphi|[\uparrow P]\psi)) \vee \\ &(\neg B(\neg[\uparrow P]\psi|P) \wedge B([\uparrow P]\varphi|(P \wedge [\uparrow P]\psi))) \end{aligned}$$

4.3. Discussion

We now discuss a few repercussions of our logical stance on belief revision.

Modular architecture Our two theorems show that it is quite easy to provide complete dynamic logics of relation upgrade, and hence of concrete belief revision policies, in the compositional format of *DEL*. Moreover, results are not ad-hoc but systematic. *DEL*-style logics of belief revision have a modular architecture with

(a) a base logic for a static language of the right expressive power, perhaps ‘engineered’ for the purpose of reduction, (b) general reduction axioms reflecting how the operation is supposed to work: a (partial) function, perhaps even a relation, (c) a special reduction axiom for beliefs after upgrade, which encodes the particular upgrade mechanism being used.

Thus, the freedom in choosing ‘policies’ for belief revision is made visible in the key axiom for belief change. Perhaps this is still too implicit, and one might wish for some locus in the formal language where one can explicitly insert an agent-dependent ‘policy’. While this seems feasible, *AGM* and *DEL* as they stand offer no such systematic facility (cf. [BEN 04], [LIU 06a] for some proposals). [BEN 07] also give a general method for deriving reduction axioms like the above from relation-changing definitions, providing the latter are given in some *PDL*-style format. This seems applicable to many proposals in the literature on belief revision.¹¹

Static pre-encoding again Our compositional reduction says that any statement about effects of later information updates or belief revisions is already ‘encoded’ in the initial model, before any actions has taken place. We phrased this before as: ‘the epistemic present contains the epistemic future’. Well-understood, this is much like phenomena in classical logic. E.g., world elimination for $\uparrow P$ passes to a definable submodel of some model \mathcal{M} . But any model \mathcal{M} already encodes all formula φ true in its *definable submodels* $\mathcal{M}|A$, by means of syntactically relativized formulas $(\varphi)^P$, as

$$\mathcal{M} \models (\varphi)^P \text{ iff } \mathcal{M}|P \models \varphi$$

Likewise, relation changes like $\uparrow P$ can be pre-computed by a syntactic *translation* substituting new relation expressions for old ones. Thus, one who merely knows the ‘plain’ consequences of some formula, implicitly also knows a lot about what it entails in other models (cf. [BAR 99]).

Iteration once more Finally, consider the earlier issue of iteration. In *PAL*, successive announcements of hard facts could be compressed into one, using the law

$$[!A][!B]\varphi \leftrightarrow [!(A \wedge [A!]B)]\varphi$$

Is there a similar ‘compression law’ for relation change and belief revision? We do not have a general answer, but here is a partial one. Let the propositions A, B be *factual*. Then the above axiom for information update reduces to the following:

$$[!A][!B]\varphi \leftrightarrow [!(A \wedge B)]\varphi$$

Something similar occurs with beliefs following revision steps. E.g., for the above mechanism of minimal reordering, the following can happen. If we first apply $\uparrow A$, then the best A -worlds come on top, leaving the remaining world order the same. Then applying $\uparrow B$ leads to two cases. If there are $A \wedge B$ -worlds, the new topmost

11. Hannes Leitgeb has suggested that the methodology of this paper may place some constraints on upgrade policies that change agents’ plausibility orderings. But reduction axioms also depend on language design. More policies are definable in a richer *modal preference language* for comparative models: cf. Section 3.1.2.

worlds are the best $A \wedge B$ -worlds from the old situation. If there are no $A \wedge B$ -worlds, the new topmost worlds are the best B -worlds from the old situation, which are also A . Thus, the following two principles hold for successive factual belief change:

$$\begin{aligned} & [\uparrow A][\uparrow B]B_i\varphi \rightarrow [\uparrow(A \wedge B)]B_i\varphi, \\ & E(A \wedge B) \wedge [\uparrow(A \wedge B)]B_i\varphi \rightarrow [\uparrow A][\uparrow B]B_i\varphi. \end{aligned}$$

As long as the two incoming beliefs are consistent with each other, getting the triggers successively or as a conjunction has no effect on current beliefs. But even so, two successive steps $\uparrow A$, $\uparrow B$ rearrange the total ordering in a different way from one step $\uparrow(A \wedge B)$, once we look at the level next to the top. In general, there may not be any iteration law compressing the total effect of two revision steps to just one with the same consequences for conditional belief. And, why *should* there be?

Many agents and common beliefs Standard belief revision policies have no ‘social’ scenario: they describe what a single agent does when confronted with surprising facts. But there is a natural generalization to a more interactive setting, where an agent is confronted with information from other sources, which need to be integrated into one new plausibility ordering. In that case, we must analyze *belief merge* (cf. [MAY 98]), and perhaps more general forms of ‘judgment aggregation’ ([LIS 04]). And there is an issue what would be plausible ‘AGM postulates’ for this interactive setting. Construed either way, it is important how agents achieve *common beliefs* after upgrade. This call for a generalization of our Theorems 5 and 6 to include common belief in groups after plausibility upgrade (cf. [BEN 06d] for the case of common knowledge after update), something which we leave as an open problem here.

5. Belief revision postulates as modal frame correspondences

Now, what about the more standard *postulational* approach to belief revision? The latter *modus operandi* advocates no specific mechanism for relation change, but the postulates of [GAR 87] rather constrain the whole family of options. A corresponding modal style way of thinking exists, viz. that of Segerberg’s *dynamic doxastic logic DDL* (cf., e.g., [SEG 98]). This system provides an abstract modal framework where one merely assumes that *some* relation change takes place on the current model: either functional, or even non-deterministic relational. The main operator will look like this:

$$\mathcal{M}, s \models [*A]\varphi \text{ iff } \mathcal{M} * \llbracket A \rrbracket, s \models \varphi \text{ where } \llbracket A \rrbracket \text{ is the set of worlds in } \mathcal{M} \text{ satisfying } A, \text{ and } \mathcal{M} * \llbracket A \rrbracket \text{ is some changed version of the model } \mathcal{M}.$$

Current *DDL* uses models that resemble Lewis sphere systems for conditional logic, or generalized versions,¹² and $\mathcal{M} * \llbracket A \rrbracket$ is then specified by the new comparative relation, leaving \mathcal{M} ’s set of worlds the same. We refer to the cited literature for details. Clearly, the axioms of the minimal modal logic K will be valid on any such models

12. Sometimes, for greater generality, *DDL* uses modal neighborhood semantics generalizing world-to-world accessibilities to world-to-set relations R_sX .

for belief change, just by the nature of the above format. On top of that, additional postulates will constrain the relation changes that correspond to ‘bona fide’ belief revision policies. And in the limit, a particular set of axioms might even determine one particular revision policy.

The final main contribution of this paper places belief revision theory once more on standard modal ground. The analysis of general postulates can be done systematically by a standard modal technique, viz. *frame correspondences*. How this works can be demonstrated quite simply on the static models of this paper.¹³

Take a functional framework of arbitrary relation changing operations $\spadesuit A$ over simple models consisting of worlds and a ternary comparison relation $\leq_s xy$:

$\spadesuit A$ takes any model \mathcal{M} and a set of worlds A in it, and yields a new model $\mathcal{M}\spadesuit A$ with the same set of worlds but some possibly changed relation \leq_s .

We now analyze some *AGM*-type general assertions in a correspondence format.¹⁴ It then quickly becomes clear what constraints they express on relation change. For a start, ‘Success’ says something weak, true for both the earlier $\uparrow P$ and $\uparrow P$:

FACT 16. — The formula $[\spadesuit p]Bp$ says that the best worlds in $\mathcal{M}\spadesuit p$ are all in p .

This trivial observation needs no proof. But actually, it seems as if we could safely demand something much stronger on relation change for belief revision, viz. that

The best worlds in $\mathcal{M}\spadesuit p$ are precisely the best p -worlds in \mathcal{M} UC

This, too, can be expressed. But then, we need the following stronger Ramsey-style dynamic formula, involving two different proposition letters p and q :

FACT 17. — The formula $B(q|p) \leftrightarrow [\spadesuit p]Bq$ expresses UC .

But actually, this preoccupation with the Upper Classes still fails to constrain the *total* relation change. For that, we really need to look at the social order in all classes after the Revolution, i.e., at *conditional beliefs* following relation upgrade.

As a deeper illustration, we consider the crucial reduction axiom for $\uparrow P$, now stated using proposition letters instead of schematic variables for arbitrary formulas. As these refer to bare sets, we suppress the earlier dynamic modalities $[\uparrow P]\psi$ which

13. *DDL* also has more baroque ‘vegetablian’ model structures with onions and broccoli, but [GIR 07] shows that ‘onion models’ are in one-to-one correspondence with Lewis-style models for conditional logic, and ‘broccoli models’ with Burgess-Veltman partial-order models.

14. A technical point: Standard modal frame correspondences come in the following format:

The modal K4-axiom $\Box p \rightarrow \Box\Box p$ is true at world s in a frame $\mathcal{F} = (W, R)$ iff the relation R is transitive at s : i.e., $\mathcal{F}, s \models \forall y(Rxy \rightarrow \forall z(Ryz \rightarrow Rxz))$.

‘Frame truth’ means the formula is true under all valuations on frame F for its proposition letters. Thus, it does not matter whether we use the formula $\Box p \leftrightarrow \Box\Box p$ or the schema $\Box\varphi \leftrightarrow \Box\Box\varphi$. Not so for *PAL* and related *DEL* systems, given the earlier difference between plain validity and schematic validity. In what follows we use proposition letters, and we mean it...

kept track of possible ‘transfer effects’. The following argument shows that it determines lexicographic reordering of models completely: a show-case for our correspondence take on the postulational approach to belief revision:

THEOREM 18. — *The formula $[\spadesuit p]B(q|r) \leftrightarrow (E(p \wedge r) \wedge B(q | p \wedge r)) \vee (\neg E(p \wedge r) \wedge B(q|r))$ holds in a frame iff the operation interpreting $\spadesuit p$ is lexicographic upgrade.*

PROOF. — Suppose that $\leq_s xy$ in $\mathcal{M}_{\spadesuit p}$. We first show that \leq_s is the relation produced by lexicographic upgrade. Let r be the set $\{x, y\}$ and $q = \{x\}$. Then the left-hand side of our formula is true. Now we have two options on the right-hand side. *Case 1:* one of x, y is in p , and hence $p \wedge r = \{x, y\}$ (1.1) or $\{y\}$ (1.2) or $\{x\}$ (1.3). Moreover, $B(q|p \wedge r)$ holds in \mathcal{M} at s . In case (1.1), we have $\leq_s xy$ in \mathcal{M} . In case (1.2), we must have $y = x$, and again $\leq_s xy$ in \mathcal{M} . Case (1.3) can only occur when $x \in p$ and not $y \in p$. Thus, all new relational pairs in $\mathcal{M}_{\spadesuit p}$ satisfy the description of the lexicographic reordering. *Case 2* is when none of x, y are in p , and it can be analyzed analogously, using the truth of the disjunct $\neg E(p \wedge r) \wedge B(q|r)$.

Conversely, we must show that all pairs which satisfy the description of lexicographic upgrade do make it into the new order. Here is one example; the other case is similar. Suppose that $x \in p$ while $y \notin p$. Next, set $r = \{x, y\}$ and $q = \{x\}$. Then $p \wedge r = \{x\}$, and we have $E(p \wedge r) \wedge B(q|p \wedge r)$ for obvious reasons. The left-hand side formula $[\spadesuit p]B(q|r)$ is then also true, since our axiom is supposed to hold for any interpretation of the proposition letters q, r - and it tells us that, in the model $\mathcal{M}_{\spadesuit p}$, the best worlds in $\{x, y\}$ are in $\{x\}$: i.e., $\leq_s xy$. ■

The setting of the preceding correspondence argument can be made more precise - something we will do in an extended version of this paper. One can work inside a universe of plausibility frames and transition relations between them, with second-order quantifiers ranging over sets of worlds inside and across these frames.

In this setting, the *AGM*-postulates are modal principles to be analyzed in the same style. Some of them have already been discussed briefly in Sections 4.2 and 4.3. In general, these postulates involve an interplay of *two* abstract operations that change models: *update* $!P$ and *upgrade* $\spadesuit P$, leading to mixed principles such as

$$(a) \quad [\spadesuit(p \wedge q)]Br \rightarrow [!q][\spadesuit p]Br$$

$$(b) \quad [\spadesuit p]Eq \wedge [!q][\spadesuit p]Br \rightarrow [\spadesuit(p \wedge q)]Br$$

These, too, can be analyzed in correspondence style. Instead of doing this now¹⁵, we merely show how an abstract postulational *AGM*-style analysis also works for *PAL*-style information updates $!P$. First, for the sake of uniformity, instead of eliminating worlds, we can also describe this operation as changing the epistemic accessibility relations, by cutting all links between P -worlds and $\neg P$ -worlds. Thus, we are in the same format as before, and again, the earlier crucial reduction axiom for public announcement turns out to capture this operation completely:

15. This analysis, too, will be included in the final extended version of this paper.

THEOREM 19. — *Eliminative update is determined completely by the formula*

$$[\spadesuit p]Kq \leftrightarrow (p \rightarrow K[\spadesuit p]q)$$

PROOF. — From left to right, the formula implies the following. Take q equal to the set of worlds which are \sim -accessible from the current s inside the set p . Assume also that s is in p . Then the right-hand side says that all worlds still \sim -accessible from s after the operation $\spadesuit p$ are in q : i.e., they were accessible before, and they were members of p . Thus, the relation change leaves only already existing links from p -worlds to p -worlds. By a similar argument in the converse direction, we see that indeed, *all* such links are preserved into the new model after the operation $\spadesuit p$. This is precisely the link-cutting version of epistemic update described before. ■

Again, this correspondence argument could be sharpened up, by defining the relevant universe of epistemic frames and transition relations more explicitly, and stipulating how individual worlds can be related across frames. In such a setting, we would need *three axioms* to zero in more accurately on eliminative update for public announcement. First, we use the equivalence (a) $\langle !p \rangle T \leftrightarrow p$ to make sure that inside a given model \mathcal{M} , the only worlds surviving into $\mathcal{M}\spadesuit p$ are those in the set denoted by p . Next, we use a reduction axiom (b) $\langle !p \rangle E q \leftrightarrow p \wedge E \langle !p \rangle q$ for the existential modality $E q$ (“ q is true in some world”) to make sure that the domain of $\mathcal{M}\spadesuit p$ does not contain any objects beyond the set p in \mathcal{M} . Finally, the above axiom (c) for knowledge ensures that the epistemic relations are the same in \mathcal{M} and $\mathcal{M}\spadesuit p$, so that our update operation really takes a *submodel*.

These observations point at a more general correspondence theory for languages with model changing modalities - but this is far beyond what we need here.

6. Extensions: richer triggers, further policies, temporal perspective

6.1. Other triggers for belief revision

The term ‘trigger’ has been used a lot, but not much has been said about them. In general *DEL*-style update logics, triggers for information change can be much more complex than just public announcements. In particular, *event models* $\mathcal{A} = (E, \{\sim_i \mid i \in G\}, PRE(e) \mid e \in E)$ model relevant events, and epistemic relations \sim_i encode what agents cannot distinguish. The preconditions $PRE(e)$ describe just when an event e can take place. We also put an ‘actual event’ e to get (\mathcal{A}, e) . This format describes scenarios where not all agents have the same observational access to what is happening, as in conversations, games, emails with private *bcc* actions, or indeed, any sophisticated human activity. This calls for a more complex notion of *product update*, turning the current epistemic model (\mathcal{M}, s) into a

product model $(\mathcal{M} \times \mathcal{A}, (s, e))$ with domain $\{(s, e) \mid s \text{ a world in } \mathcal{M}, e \text{ an event in } \mathcal{A}\}$, $(\mathcal{M}, s) \models PRE(e)$, and new accessibility relation: $(s, e) \sim_i (t, f)$ iff both $s \sim_i t$ and $e \sim_i f$.

The valuation for atoms p at (s, e) is that at s in \mathcal{M} - though this can be generalized to deal with genuine world change. A product model $\mathcal{M} \times \mathcal{A}$ can be larger than \mathcal{M} itself, recording information of different agents about the facts and what the others know in complex scenarios of conversation and observation such as card games. This is the real arena where *DEL* lives today. The full language has the following syntax:

$$p \mid \varphi \mid \varphi \vee \psi \mid K_i \varphi \mid C_G \varphi \mid [\mathcal{A}, e] \varphi: (\mathcal{A}, e) \text{ any event model with actual event } e.$$

Semantically, $\mathcal{M}, s \models [\mathcal{A}, e] \varphi$ iff $\mathcal{M} \times \mathcal{A}, (s, e) \models \varphi$. [BAL 98] show that the resulting logic *LEA* is effectively axiomatizable and decidable. Here is a typical valid reduction axiom for knowledge of agents:

$$[\mathcal{A}, e] K_i \varphi \leftrightarrow (PRE(e) \rightarrow \bigwedge_{f \sim_i e} K_i [\mathcal{A}, f] \varphi)$$

Cf. [BEN 06a], [DIT 07], [BEN 06d] for further dynamic-epistemic logics working in the same compositional style as above.

Richer scenarios This is also a natural continuation of our present analysis of belief revision. Event models provide much richer triggers for information update and belief revision. These include cases where agents have possibly mistaken beliefs about which event they are witnessing. (I think I see you draw a red card, but it was an orange one.) Indeed, our beliefs are seldom changed by simple actions like the above $!P$ or $\uparrow P$. They change in complex conversations, games, and other real-life phenomena studied in *DEL* with product update. What is good for the one, is good for the other. Also, product update is easy to generalize to scenarios where *real change* occurs to the world ([BEN 06d]). The compositional methodology of our update logics extends automatically to deal with ‘update’ in the non-*AGM* Katsuno-Mendelzon sense of information about real changes that have taken place.

Even though no specific theorem is formulated here, our claim is that

The *PAL* analysis in Section 4 extends straightforwardly to a complete *DEL*-style dynamic logic for product update performing both information update and belief revision on epistemic-doxastic event models.

For a recent system of this sort, developed independently from this paper, cf. [BAL 06b] and [BAL 06a]. These propose one general upgrade rule for plausibility relations giving priority to the plausibility ordering among the last-observed events.

A complication Still, there are also difficulties with product update for complex triggers. What if signals *disagree*? Say, I believe that p is the case, but now observe a signal suggesting that p holds, with equal strength in the opposite direction. How should I *merge* my earlier view and the new input into one coherent picture? [AUC 03] uses Spohn-style ‘graded models’ for this purpose, with numerical computations on plausibility strength of worlds. [LIU 06b] proposes mechanisms of relation change via ‘utility upgrade’. While these systems can be made to work in the *DEL*-spirit, they do raise issues about whether belief revision should not eventually be construed in terms of *strengths* of beliefs, rather than just beliefs themselves.

6.2. Temporal perspective

DEL and *AGM* are in the same boat with respect to many further issues. E.g., it has often been observed that many informational processes involve both the *temporal past*, i.e., the history of what has happened so far, and the *temporal future*. E.g., our beliefs about an agent may also depend on hypotheses about its long-term future behavior. This brings us to the realm of *epistemic temporal logics*, which occur both in the philosophical literature ([BEL 01]) and the computational one ([FAG 95], [PAR 03]). I skip this angle here, since it seems to pose similar issues in both settings. But what is true is that temporal logics offer a much richer canvas on which to study phenomena of information update, learning, and belief revision (cf. [KEL 96]). [BEN 06f] is a survey of the area including a comparison with *DEL*.

‘Backward’ versus ‘forward’ in update In comparing all these logics, the following contrast plays a role. Some logics in the temporal tradition are ‘forward-looking’. Unlike *DEL*, they derive future states not from informational events transforming about the current setting, but from commands of the *STIT*-type: ‘see to it that φ comes about’. In this style of analysis, one does not have to tell the agent *how* to do this. To some extent, *AGM* is more like this, as no concrete instruction for change is provided, and one just assumes that the command ‘join the Believers in φ ’ can be obeyed in some manner. *DEL* on the other hand, tends to think of such forward instructions as ‘wishful thinking’, and rather analyzes concrete given event scenarios for the changes which they turn out to produce.

The forward style of thinking sits better, perhaps, with the ‘Grand Stage’ idea of some temporal universe already containing all histories that are possible lines of investigation. An update $!P$ is then an instruction to make a *minimal move to some already available future state* where one knows that P . And the same holds for belief revision.¹⁶ In a setting like that, no definable explicit construction takes place for ‘the next model’ as happens in *DEL* - and it is the externally supplied temporal *model* which decides which next stage is reachable by some minimal move of coming to know or believe some proposition φ . By contrast, *DEL* may be viewed as a sort of ‘mini process algebra’ of successive model construction.

7. Conclusion

The main point of this paper is just to show that ‘it can be done’: (a) dealing with concrete mechanisms for *AGM* belief change within the *DEL* paradigm, and then describing their properties completely; and conversely, (b) analyzing abstract revision postulates in a standard modal correspondence style. The result is one merged theory

16. Interestingly, the first modal analysis of *AGM* that I know of, in [BEN 89] (a Logic Colloquium lecture from 1987 reacting to Gärdenfors’ earlier work) works with three modalities $[+A]$, $[-A]$ and $[*A]$, for update, contraction and revision, respectively, over a temporal universe of successive ‘information stages’ ordered by inclusion.

of information and belief revision, which uses standard modal techniques, the ‘lingua franca’ of our field. Moreover, we can now freely transfer issues and results between the two research areas - provided that we see through superficial differences in ‘life-styles’ and idiosyncratic discussion topics, such as Ramsey Tests in *AGM*, or modal bisimulation folklore in *DEL*.

Finally, this proposal has not dropped out of the blue. Related work includes, in particular, dynamic doxastic logic ([SEG 95], [SEG 99], [LIN 06]) and dynamic logics of preferences ([BEN 07]), but also relevant are conditional logic, non-monotonic logic, preference logic, database update ([FAG 86]), *STIT* logic, and epistemic temporal logic - witness the references throughout our text.

Acknowledgements

I thank Alexandru Baltag, Hans van Ditmarsch, Wiebe van der Hoek, Patrick Girard, Hannes Leitgeb, Fenrong Liu, Hans Rott, Tomasz Sadzik, Sonja Smets, and Frank Veltman, and as well as the referees and editors of this Journal for their comments and suggestions.

8. References

- [AUC 03] AUCHER G., “A Combination System for Update Logic and Belief Revision”, Master’s thesis, ILLC, University of Amsterdam, 2003.
- [BAL 98] BALTAG A., MOSS L., SOLECKI S., “The Logic of Common Knowledge, Public Announcements, and Private Suspicions”, GILBOA I., Ed., *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*, 1998, p. 43–56.
- [BAL 06a] BALTAG A., SMETS S., “Dynamic Belief Revision over Multi-Agent Plausibility Models”, in *Proceedings of the 7th Conference on Logic and the Foundations of Game and Decision Theory (LOFT 06)*, Liverpool, 2006.
- [BAL 06b] BALTAG A., SMETS S., “The Logic of Conditional Doxastic Actions: A theory of dynamic multi-agent belief revision”, in S. Artemov and R. Parikh, Ed., *Proceedings of the Workshop on Rationality and Knowledge*, ESSLLI, Malaga, 2006.
- [BAR 99] BARWISE J., VAN BENTHEM J., “Interpolation, Preservation, and Pebble Games”, *Journal of Symbolic Logic*, vol. 64:2, 1999, p. 881-903.
- [BEL 01] BELNAP N., PERLOFF M., XU M., *Facing the Future*, Oxford University Press, Oxford, 2001.
- [BEN 89] VAN BENTHEM J., “Semantic Parallels in Natural Language and Computation”, EBBINGHAUS H.-D. *et al.*, Ed., *Logic Colloquium, Granada 1987*, p. 331-375, North-Holland, Amsterdam, 1989.
- [BEN 93] VAN BENTHEM J., VAN EIJCK J., FROLOVA A., “Changing Preferences”, Tech Report, CS-93-10, Center for Mathematics and Computer Science, Amsterdam, 1993.

- [BEN 03] VAN BENTHEM J., “Structural Properties of Dynamic Reasoning”, PEREGRIN J., Ed., *Meaning: the Dynamic Turn*, Elsevier, Amsterdam, 2003, p. 15-31.
- [BEN 04] VAN BENTHEM J., LIU F., “Diversity of Logical Agents in Games”, *Philosophia Scientiae*, vol. 8(2), 2004, p. 163-178.
- [BEN 06a] VAN BENTHEM J., “‘One is a Lonely Number’: On the Logic of Communication”, Tech Report, PP-2002-27, ILLC, University of Amsterdam. In Z. Chatzidakis, P. Koepke, W. Pohlers, Eds., *Logic Colloquium '02*, ASL & A.K.Peters, Wellesley MA, 96-129, 2006.
- [BEN 06b] VAN BENTHEM J., “Open Problems in Logical Dynamics”, in T. Rozhkovskaya, ed., *Mathematical Problems from Applied Logics*, Springer, New York, 2006.
- [BEN 06c] VAN BENTHEM J., GIRARD P., ROY O., “Ceteris Paribus Preference Logic”, Manuscript. Stanford University and ILLC Amsterdam, 2006.
- [BEN 06d] VAN BENTHEM J., VAN EIJCK J., KOOI B., “Logics of Communication and Change”, *Information and Computation*, vol. 204(11), 2006, p. 1620-1662, previous version in *TARK 2005* as ‘Common Knowledge in Update Logics’.
- [BEN 06e] VAN BENTHEM J., VAN OTTERLOO S., ROY O., “Preference Logic, Conditionals and Solution Concepts in Games”, LAGERLUND H., LINDSTRÖM S., SLIWINSKI R., Eds., *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*, p. 61-77, Uppsala Philosophical Studies 53, 2006.
- [BEN 06f] VAN BENTHEM J., PACUIT E., “The Tree of Knowledge in Action: Towards a Common Perspective”, Proceedings *Advances in Modal Logic*, 2006.
- [BEN 07] VAN BENTHEM J., LIU F., “Dynamic Logic of Preference Upgrade”, Tech Report, PP-2005-29, ILLC, University of Amsterdam. To appear in the *Journal of Applied Non-Classical Logic*, 2007.
- [BOU 94] BOUTILIER C., “Conditional Logics of Normality: a Modal Approach”, *Artificial Intelligence*, vol. 68, 1994, p. 87-154.
- [BUR 84] BURGESS J., “Basic Tense Logic”, vol. 2 of *Handbook of Philosophical Logic*, chapter 2, p. 89-133, Reidel, 1984.
- [DIT 06] VAN DITMARSCH H., “Prolegomena to Dynamic Logic for Belief Revision”, to appear in: *Knowledge, Rationality & Action (Synthese)*, 2006.
- [DIT 07] VAN DITMARSCH H., VAN DER HOEK W., KOOI B., *Dynamic Epistemic Logic*, Springer, Berlin, 2007, to appear.
- [FAG 86] FAGIN R., KUPER G., ULLMAN J., VARDI M., “Updating Logical Databases”, KANNELAKIS P., PREPARATA F., Eds., *Advances in Computing Research*, vol. III, p. 1-18, JAI Press, 1986.
- [FAG 95] FAGIN R., HALPERN J. Y., MOSES Y., VARDI M. Y., *Reasoning about Knowledge*, The MIT Press, 1995.
- [GAR 87] GARDENFORS P., *Knowledge in Flux*, The MIT Press, Cambridge (Mass.), 1987.
- [GAR 95] GARDENFORS P., ROTT H., “Belief Revision”, GABBAY D. M., HOGGER C. J., ROBINSON J. A., Eds., *Handbook of Logic in Artificial Intelligence and Logic Programming*, vol. 4, Oxford University Press, 1995.

- [GER 99] GERBRANDY J., “Bisimulations on Planet Kripke”, PhD thesis, ILLC, Amsterdam, 1999.
- [GIR 07] GIRARD P., “From Onions to Broccoli: Generalizing Lewis’ Conditional Logic”, Tech Report, PP-2005-12, ILLC, University of Amsterdam. To appear in the *Journal of Applied Non-Classical Logics*, 2007.
- [KEL 96] KELLY K., *The Logic of Reliable Inquiry*, Oxford University Press, Oxford, 1996.
- [KOO 05] KOOI B., “As the World Turns: on the Logic of Public Updates”, Manuscript, University of Groningen, 2005.
- [LEW 73] LEWIS D., *Counterfactuals*, Harvard University Press, 1973.
- [LIN 06] LINDSTROM S., SEGERBERG K., “Modal Logic and Philosophy”, in P. Blackburn, J. van Benthem and F. Wolter, Eds., *Handbook of Modal Logic*, 2006.
- [LIS 04] LIST C., PETTIT P., “Aggregating Sets of Judgments. Two Impossibility Results Compared”, *Synthese*, vol. 140, 2004, p. 207-235.
- [LIU 06a] LIU F., “Diversity of Agents”, Tech Report, PP-2006-37, ILLC, University of Amsterdam. In *Proceedings of the Workshop on Logics for Resource Bounded Agents*, ESSLLI, Malaga, 2006.
- [LIU 06b] LIU F., “Preference Change and Information Processing”, Tech Report, PP-2006-41, ILLC, University of Amsterdam. In *Proceedings of the 7th Conference on Logic and the Foundations of Game and Decision Theory (LOFT 06)*, Liverpool, 2006.
- [LUT 06] LUTZ C., “Complexity and Succinctness of Public Announcement Logic”, 2006, in *Proceedings of the Fifth International Conference on Autonomous Agents and Multiagent Systems (AAMAS06)*.
- [MAY 98] MAYNARD-REID II P., SHOHAM Y., “From Belief Revision to Belief Fusion”, *Proceedings of LOFT-98*, 1998, Torino.
- [MIL 05] MILLER J., MOSS L., “The Undecidability of Iterated Modal Relativization”, *Studia Logica*, vol. 97:3, 2005, p. 373-407.
- [PAR 03] PARIKH R., RAMANUJAM R., “A Knowledge-Based Semantics of Messages”, *Journal of Logic, Language and Information*, vol. 12, 2003, p. 453-467.
- [PLA 89] PLAZA J. A., “Logics of Public Communications”, *Proceedings 4th International Symposium on Methodologies for Intelligent Systems*, 1989.
- [ROT 06] ROTT H., “Shifting Priorities: Simple Representations for 27 Iterated Theory Change Operators”, LAGERLUND H., LINDSTRÖM S., SLIWINSKI R., Eds., *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*, p. 359-384, Uppsala Philosophical Studies 53, 2006.
- [SEG 95] SEGERBERG K., “Belief Revision from the Point of View of Doxastic Logic”, *Bulletin of the IGPL*, vol. 3, 1995, p. 534-553.
- [SEG 98] SEGERBERG K., “Irrevocable Belief Revision in Dynamic Doxastic Logic”, *Notre Dame Journal of Formal Logic*, vol. 39, 1998, p. 287-306.
- [SEG 99] SEGERBERG K., “Default Logic as Dynamic Doxastic Logic”, *Erkenntnis*, vol. 50, 1999, p. 333-352.

- [SHO 88] SHOHAM Y., *Reasoning About Change: Time and Change from the Standpoint of Artificial Intelligence*, MIT Press, 1988.
- [SPO 88] SPOHN W., “Ordinal Conditional Functions: A Dynamic Theory of Epistemic States”, HARPER W. L. *et al.*, Ed., *Causation in Decision, Belief Change and Statistics II*, p. 105-134, Kluwer, Dordrecht, 1988.
- [VEL 85] VELTMAN F., “Logics for Conditionals”, PhD thesis, University of Amsterdam, 1985.
- [VEL 96] VELTMAN F., “Defaults in Update Semantics”, *Journal of Philosophical Logic*, vol. 25, 1996, p. 221-261.
- [YAM 06] YAMADA T., “Acts of Commands and Changing Obligations”, *Proceedings 7th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA VII)*, 2006.