
Yet More Modal Logics of Preference Change and Belief Revision

Jan van Eijck

Centre for Mathematics and Computer Science (CWI)
Kruislaan 413
1098 SJ Amsterdam, The Netherlands
jve@cwi.nl

Abstract

We contrast Bonanno’s ‘Belief Revision in a Temporal Framework’ [15] with preference change and belief revision from the perspective of dynamic epistemic logic (DEL). For that, we extend the logic of communication and change of [11] with relational substitutions [8] for preference change, and show that this does not alter its properties. Next we move to a more constrained context where belief and knowledge can be defined from preferences [29; 14; 5; 7], prove completeness of a very expressive logic of belief revision, and define a mechanism for updating belief revision models using a combination of action priority update [7] and preference substitution [8].

1 Reconstructing AGM Style Belief Revision

Bonanno’s paper offers a rational reconstruction of Alchourrón Gärdenfors Makinson style belief revision (AGM belief revision) [1] (see also [22] and [23]), in a framework where modalities B for single agent belief and I for being informed are mixed with a next time operator \bigcirc and its inverse \bigcirc^{-1} .

Both the AGM framework and Bonanno’s reconstruction of it do not explicitly represent the triggers that cause belief change in the first place. $I\phi$ expresses that the agent is informed that ϕ , but the communicative action that causes this change in information state is not represented. Also, ϕ is restricted to purely propositional formulas. Another limitation that Bonanno’s reconstruction shares with AGM is that it restricts attention to a single agent: changes of the beliefs of agents about the beliefs of other agents are not analyzed. In these respects [15] is close to Dynamic Doxastic Logic (DDL), as developed in [38; 39].

AGM style belief revision was proposed more than twenty years ago, and has grown into a paradigm in its own right in artificial intelligence. In the meanwhile rich frameworks of dynamic epistemic logic have emerged that are quite a bit more ambitious in their goals than AGM was when it was first proposed. AGM analyzes operations $+\phi$ for expanding with ϕ , $-\phi$ for retracting ϕ and $*\phi$ for revising with ϕ . It is formulated in a purely syntactic way, it hardly addresses issues of semantics, it does not propose sound and complete axiomatisations. It did shine in 1985, and it still shines now, but perhaps in a more modest way.

Bonanno's paper creates a nice link between this style of belief revision and epistemic/doxastic logic. While similar in spirit to Segerberg's work, it addresses the question of the rational reconstruction of AGM style belief revision more explicitly. This does add quite a lot to that framework: clear semantics, and a sound and complete axiomatisation. Still, it is fair to say that this rational reconstruction, nice as it is, also inherits the limitations of the original design.

2 A Broader Perspective

Meanwhile, epistemic logic has entered a different phase, with a new focus on the epistemic and doxastic effects of information updates such as public announcements [34; 24]. Public announcements are interesting because they create common knowledge, so the new focus on information updating fostered an interest in the evolution of multi-agent knowledge and belief under acts of communication.

Public announcement was generalized to updates with 'action models' that can express a wide range of communications (private announcements, group announcements, secret sharing, lies, and so on) in [4] and [3]. A further generalization to a complete logic of communication and change, with enriched actions that allow changing the facts of the world, was provided in [11]. The textbook treatment of dynamic epistemic logic in [20] bears witness to the fact that this approach is by now well established.

The above systems of dynamic epistemic logic do provide an account of knowledge or belief update, but they do not analyse belief revision in the sense of AGM. Information updating in dynamic epistemic logic is monotonic: facts that are announced to an audience of agents cannot be unlearned. Van Benthem [8] calls this 'belief change under hard information' or 'eliminative belief revision'. See also [19] for reflection on the distinction between this and belief change under soft information.

Assume a state of the world where p actually is the case, and where you know it, but I do not. Then public announcement of p will have the effect that I get to know it, but also that you know that I know it, that I know that you know that I know it, in short, that p becomes common knowledge. But if this announcement is followed by an announcement of $\neg p$, the effect will be inconsistent knowledge states for both of us.

It is clear that AGM deals with belief revision of a different kind: ‘belief change under soft information’ or ‘non-eliminative belief revision’. In [8] it is sketched how this can be incorporated into dynamic epistemic logic, and in the closely related [7] a theory of ‘doxastic actions’ is developed that can be seen as a further step in this direction.

Belief revision under soft information can, as Van Benthem observes, be modelled as change in the belief accessibilities of a model. This is different from public announcement, which can be viewed as elimination of worlds while leaving the accessibilities untouched.

Agent i believes that ϕ in a given world w if it is the case that ϕ is true in all worlds t that are reachable from w and that are minimal for a suitable plausibility ordering relation \leq_i . In the dynamic logic of belief revision these accessibilities can get updated in various ways. An example from [36] that is discussed in [8] is $\uparrow A$ for so-called ‘lexicographic upgrade’: all A worlds get promoted past all non- A worlds, while within the A worlds and within the non- A worlds the preference relation stays as before. Clearly this relation upgrade has as effect that it creates belief in A . And the belief upgrade can be undone: a next update with $\uparrow \neg A$ does not result in inconsistency.

Van Benthem [8] gives a complete dynamic logic of belief upgrade for the belief upgrade operation $\uparrow A$, and another one for a variation on it, $\uparrow A$, or ‘elite change’, that updates a plausibility ordering to a new one where the best A worlds get promoted past all other worlds, and for the rest the old ordering remains unchanged.

This is taken one step further in a general logic for changing preferences, in Van Benthem and Liu [9], where upgrade as relation change is handled for (reflexive and transitive) preference relations \leq_i , by means of a variation on product update called product upgrade. The idea is to keep the domain, the valuation and the epistemic relations the same, but to reset the preferences by means of a substitution of new pre-orders for the preference relations.

Treating knowledge as an equivalence and preference as a pre-order, without constraining the way in which they relate, as is done in [9], has the advantage of generality (one does not have to specify what ‘having a preference’ means), but it makes it harder to use the preference relation for modelling be-

belief change. If one models ‘regret’ as preferring a situation that one knows to be false to the current situation, then it follows that one can regret things one cannot even conceive. And using the same preference relation for belief looks strange, for this would allow beliefs that are known to be false. Van Benthem (private communication) advised me not to lose sleep over such philosophical issues. If we follow that advice, and call ‘belief’ what is true in all most preferred worlds, we can still take comfort from the fact that this view entails that one can never believe one is in a bad situation: the belief-accessible situations are by definition the best conceivable worlds. Anyhow, proceeding from the assumption that knowledge and preference are independent basic relations and then studying possible relations between them has turned out very fruitful: the recent theses by Girard [26] and Liu [33] are rich sources of insight in what a logical study of the interaction of knowledge and preference may reveal.

Here we will explore two avenues, different from the above but related to it. First, we assume nothing at all about the relation between knowledge on one hand and preference on the other. We show that the dynamic logic of this (including updating with suitable finite update models) is complete and decidable: Theorem 3.1 gives an extension of the reducibility result for LCC, the general logic of communication and change proposed and investigated in [11].

Next, we move closer to the AGM perspective, by postulating a close connection between knowledge, belief and preference. One takes preferences as primary, and imposes minimal conditions to allow a definition of knowledge from preferences. The key to this is the simple observation in Theorem 4.1 that a preorder can be turned into an equivalence by taking its symmetric closure if and only if it is weakly connected and conversely weakly connected. This means that by starting from weakly and converse weakly connected preorders one can interpret their symmetric closures as knowledge relations, and use the preferences themselves to define conditional beliefs, in the well known way that was first proposed in Boutillier [16] and Halpern [30]. The multi-agent version of this kind of conditional belief was further explored in [12] and in [5; 7]. We extend this to a complete logic of regular doxastic programs for belief revision models (Theorem 4.3), useful for reasoning about common knowledge, common conditional belief and their interaction. Finally, we make a formal proposal for a belief change mechanism by means of a combination of action model update in the style of [7] and plausibility substitution in the style of [9].

3 Preference Change in LCC

An epistemic preference model \mathbf{M} for set of agents I is a tuple (W, V, R, P) where W is a non-empty set of worlds, V is a propositional valuation, R is a function that maps each agent i to a relation R_i (the epistemic relation for i), and P is a function that maps each agent i to a preference relation P_i . There are no conditions at all on the R_i and the P_i (just as there are no constraints on the R_i relations in LCC [11]).

We fix a PDL style language for talking about epistemic preference models (assume p ranges over a set of basic propositions $Prop$ and i over a set of agents I):

$$\begin{aligned}\phi &::= \top \mid p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid [\pi]\phi \\ \pi &::= \sim_i \mid \geq_i \mid ?\phi \mid \pi_1; \pi_2 \mid \pi_1 \cup \pi_2 \mid \pi^*\end{aligned}$$

This is to be interpreted in the usual PDL manner, with $\llbracket \pi \rrbracket^{\mathbf{M}}$ giving the relation that interprets relational expression π in $\mathbf{M} = (W, V, R, P)$, where \sim_i is interpreted as the relation R_i and \geq_i as the relation P_i , and where the complex modalities are handled by the regular operations on relations. We employ the usual abbreviations: \perp is shorthand for $\neg\top$, $\phi_1 \vee \phi_2$ is shorthand for $\neg(\neg\phi_1 \wedge \neg\phi_2)$, $\phi_1 \rightarrow \phi_2$ is shorthand for $\neg(\phi_1 \wedge \neg\phi_2)$, $\phi_1 \leftrightarrow \phi_2$ is shorthand for $(\phi_1 \rightarrow \phi_2) \wedge (\phi_2 \rightarrow \phi_1)$, and $\langle \pi \rangle \phi$ is shorthand for $\neg[\pi]\neg\phi$.

$[\pi]\phi$ is true in world w of \mathbf{M} if for all v with $(w, v) \in \llbracket \pi \rrbracket^{\mathbf{M}}$ it holds that ϕ is true in v . This is completely axiomatised by the usual PDL rules and axioms ([37; 31]):

Modus ponens	and axioms for propositional logic
Modal generalisation	From $\vdash \phi$ infer $\vdash [\pi]\phi$
Normality	$\vdash [\pi](\phi \rightarrow \psi) \rightarrow ([\pi]\phi \rightarrow [\pi]\psi)$
Test	$\vdash [?\phi]\psi \leftrightarrow (\phi \rightarrow \psi)$
Sequence	$\vdash [\pi_1; \pi_2]\phi \leftrightarrow [\pi_1][\pi_2]\phi$
Choice	$\vdash [\pi_1 \cup \pi_2]\phi \leftrightarrow ([\pi_1]\phi \wedge [\pi_2]\phi)$
Mix	$\vdash [\pi^*]\phi \leftrightarrow (\phi \wedge [\pi][\pi^*]\phi)$
Induction	$\vdash (\phi \wedge [\pi^*](\phi \rightarrow [\pi]\phi)) \rightarrow [\pi^*]\phi$

In [11] it is proved that extending the PDL language with a extra modality $[A, e]\phi$ does not change the expressive power of the language. Interpretation of the new modality: $[A, e]\phi$ is true in w in \mathbf{M} if success of the update of \mathbf{M} with action model A to $\mathbf{M} \otimes A$ implies that ϕ is true in (w, e) in $\mathbf{M} \otimes A$. To see what *that* means one has to grasp the definition of update models A and

the update product operation \otimes , which we will now give for the epistemic preference case.

An action model (for agent set I) is like an epistemic preference model for I , with the difference that the worlds are now called events, and that the valuation has been replaced by a precondition map \mathbf{pre} that assigns to each event e a formula of the language called the precondition of e . From now on we call the epistemic preference models static models.

Updating a static model $\mathbf{M} = (W, V, R, P)$ with an action model $A = (E, \mathbf{pre}, \mathbf{R}, \mathbf{P})$ succeeds if the set

$$\{(w, e) \mid w \in W, e \in E, \mathbf{M}, w \models \mathbf{pre}(e)\}$$

is non-empty. The update result is a new static model $\mathbf{M} \otimes \mathbf{A} = (W', V', R', P')$ with

- $W' = \{(w, e) \mid w \in W, e \in E, \mathbf{M}, w \models \mathbf{pre}(e)\}$,
- $V'(w, e) = V(w)$,
- R'_i is given by $\{(w, e), (v, f) \mid (w, v) \in R_i, (e, f) \in \mathbf{R}_i\}$,
- P'_i is given by $\{(w, e), (v, f) \mid (w, v) \in P_i, (e, f) \in \mathbf{P}_i\}$.

If the static model has a set of distinguished states W_0 and the action model a set of distinguished events E_0 , then the distinguished worlds of $\mathbf{M} \otimes \mathbf{A}$ are the (w, e) with $w \in W_0$ and $e \in E_0$.




Figure 1. Static model and update

Figure 1 gives an example pair of a static model with an update action. The distinguished worlds of the model and the distinguished event of the

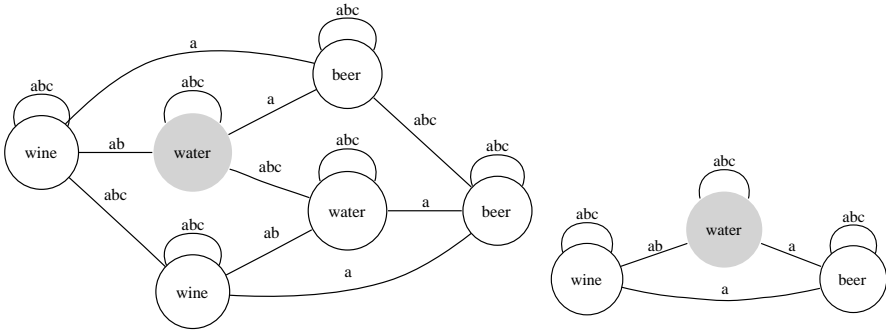


Figure 2. Update result, before and after reduction under bisimulation

action model are shaded grey. Only the R relations are drawn, for three agents a, b, c . The result of the update is shown in Figure 2, on the left. This result can be reduced to the bisimilar model on the right in the same figure, with the bisimulation linking the distinguished worlds. The result of the update is that the distinguished “wine” world has disappeared, without any of a, b, c being aware of the change.

In LCC, action update is extended with factual change, which is handled by propositional substitution. Here we will consider another extension, with preference change, handled by preference substitution (first proposed in [9]). A preference substitution is a map from agents to PDL program expressions π represented by a finite set of bindings

$$\{i_1 \mapsto \pi_1, \dots, i_n \mapsto \pi_n\}$$

where the i_j are agents, all different, and where the π_i are program expressions from our PDL language. It is assumed that each i that does not occur in the lefthand side of a binding is mapped to \geq_i . Call the set $\{i \in I \mid \rho(i) \neq \geq_i\}$ the *domain* of ρ . If $\mathbf{M} = (W, V, R, P)$ is a preference model and ρ is a preference substitution, then \mathbf{M}^ρ is the result of changing the preference map P of \mathbf{M} to P^ρ given by:

$$\begin{aligned} P^\rho(i) &:= P_i \text{ for } i \text{ not in the domain of } \rho, \\ P^\rho(i) &:= \llbracket \rho(i) \rrbracket^{\mathbf{M}} \text{ for } i = i_j \text{ in the domain of } \rho. \end{aligned}$$

Now extend the PDL language with a modality $[\rho]\phi$ for preference change,

with the following interpretation:

$$\mathbf{M}, w \models [\rho]\phi \quad : \iff \quad \mathbf{M}^\rho, w \models \phi.$$

Then we get a complete logic for preference change:

Theorem 3.1. The logic of epistemic preference PDL with preference change modalities is complete.

Proof. The preference change effects of $[\rho]$ can be captured by a set of reduction axioms for $[\rho]$ that commute with all sentential language constructs, and that handle formulas of the form $[\rho][\pi]\phi$ by means of reduction axioms of the form

$$[\rho][\pi]\phi \quad \leftrightarrow \quad [F_\rho(\pi)][\rho]\phi,$$

with F_ρ given by:

$$\begin{aligned} F_\rho(\sim_i) &:= \sim_i \\ F_\rho(\geq_i) &:= \begin{cases} \rho(i) & \text{if } i \text{ in the domain of } \rho, \\ \geq_i & \text{otherwise,} \end{cases} \\ F_\rho(? \phi) &:= ?[\rho]\phi, \\ F_\rho(\pi_1; \pi_2) &:= F_\rho(\pi_1); F_\rho(\pi_2), \\ F_\rho(\pi_1 \cup \pi_2) &:= F_\rho(\pi_1) \cup F_\rho(\pi_2), \\ F_\rho(\pi^*) &:= (F_\rho(\pi))^*. \end{aligned}$$

It is easily checked that these reduction axioms are sound, and that for each formula of the extended language the axioms yield an equivalent formula in which $[\rho]$ occurs with lower complexity, which means that the reduction axioms can be used to translate formulas of the extended language to PDL formulas. Completeness then follows from the completeness of PDL. Q.E.D.

4 Yet Another Logic ...

In this section we look at a more constrained case, by replacing the epistemic preference models by ‘belief revision models’ in the style of Grove [29], Board [14], and Baltag and Smets [5; 7] (who call them ‘multi-agent plausibility frames’). There is almost complete agreement that preference relations should be transitive and reflexive (pre-orders). But transitivity plus reflexivity of a binary relation R do not together imply that $R \cup R^{-1}$ is an equivalence. Figure 3 gives a counterexample. The two extra conditions of weak connectedness for R




Figure 3. Preorder with a non-transitive symmetric closure.

and for R^\sim remedy this. A binary relation R is weakly connected (terminology of [27]) if the following holds:

$$\forall x, y, z((xRy \wedge xRz) \rightarrow (yRz \vee y = z \vee zRy)).$$

Theorem 4.1. Assume R is a preorder. Then $R \cup R^\sim$ is an equivalence iff both R and R^\sim are weakly connected.

Proof. \Rightarrow : immediate.

\Leftarrow : Let R be a preorder such that both R and R^\sim are weakly connected. We have to show that $R \cup R^\sim$ is an equivalence. Symmetry and reflexivity are immediate. For the check of transitivity, assume $xR \cup R^\sim y$ and $yR \cup R^\sim z$. There are four cases. (i) $xRyRz$. Then xRz by transitivity of R , hence $xR \cup R^\sim z$. (ii) $xRyR^\sim z$. Then $yR^\sim x$ and $yR^\sim z$, and by weak connectedness of R^\sim , either $xR^\sim z$ or $zR^\sim x$, hence $xR \cup R^\sim z$, or $x = z$, hence xRz by reflexivity of R . Therefore $xR \cup R^\sim z$ in all cases. (iii) $xR^\sim yRz$. Similar. (iv) $xR^\sim yR^\sim z$. Then $zRyRx$, and zRx by transitivity of R . Therefore $xR \cup R^\sim z$. Q.E.D.

Call a preorder that is weakly connected and conversely weakly connected locally connected. The example in Figure 4 shows that locally connected preorders need not be connected. Taking the symmetric closure of this example generates an equivalence with two equivalence classes. More generally, taking the symmetric closure of a locally connected preorder creates an equivalence that can play the role of a knowledge relation defined from the preference order. To interpret the preference order as conditional belief, it is convenient to assume that it is also well-founded: this makes for a smooth definition of the notion of a ‘best possible world’.

A belief revision model \mathbf{M} (again, for a set of agents I) is a tuple (W, V, P) where W is a non-empty set of worlds, V is a propositional valuation and P is




Figure 4. Locally connected preorder that is not connected.

a function that maps each agent i to a preference relation \leq_i that is a locally connected well-preorder. That is, \leq_i is a pre-order (reflexive and transitive) that is well-founded (in terms of $<_i$ for the strict part of \leq_i , this is the requirement that there is no infinite sequence of w_1, w_2, \dots with $\dots <_i w_2 <_i w_1$), and such that both \leq_i and its converse are weakly connected.

In what follows we will use $<_i$ with the meaning explained above, \geq_i for the converse of \leq_i , $>_i$ for the converse of $<_i$, and \sim_i for $\leq_i \cup \geq_i$.

The locally connected well-preorders \leq_i can be used to induce accessibility relations \rightarrow_i^P for each subset P of the domain, by means of the following standard definition:

$$\rightarrow_i^P := \{(x, y) \mid x \sim_i y \wedge y \in \text{MIN}_{\leq_i} P\},$$

where $\text{MIN}_{\leq_i} P$, the set of minimal elements of P under \leq_i , is defined as

$$\{s \in P : \forall s' \in P (s' \leq s \Rightarrow s \leq s')\}.$$

This picks out the minimal worlds linked to the current world, according to \leq_i , within the set of worlds satisfying $\llbracket \phi \rrbracket^{\mathbf{M}}$. The requirement of wellfoundedness ensures that $\text{MIN}_{\leq} P$ will be non-empty for non-empty P . Investigating these \rightarrow^P relations, we see that they have plausible properties for belief:

Proposition 4.2. Let \leq be a locally connected well-preorder on S and let P be a non-empty subset of S . Then \rightarrow^P is transitive, euclidean and serial.

Proof. Transitivity: if $x \rightarrow^P y$ then $y \sim x$ and $y \in \text{MIN}_{\leq} P$. If $y \rightarrow^P z$ then $z \sim y$ and $z \in \text{MIN}_{\leq} P$. It follows by local connectedness of \leq that $z \sim x$ and by the definition of \rightarrow^P that $x \rightarrow^P z$.

Euclideanness: let $x \rightarrow^P y$ and $x \rightarrow^P z$. We have to show $y \rightarrow^P z$. From $x \rightarrow^P y$, $y \sim x$ and $y \in \text{MIN}_{\leq} P$. From $x \rightarrow^P z$, $z \sim x$ and $z \in \text{MIN}_{\leq} P$. From local connectedness, $y \sim z$. Hence $y \rightarrow^P z$.

Seriality: Let $x \in P$. Since \leq is a preorder there are $y \in P$ with $y \leq x$. The wellfoundedness of \leq guarantees that there are \leq minimal such y . Q.E.D.

Transitivity, euclideanness and seriality are the frame properties corresponding to positively and negatively introspective consistent belief (KD45 belief, [18]).

Figure 5 gives an example with both the \leq relation (shown as solid arrows in the direction of more preferred worlds, i.e., with an arrow from x to y for $x \geq y$) and the induced \rightarrow relation on the whole domain (shown as dotted arrows). The above gives us in fact knowledge relations \sim_i together with




Figure 5. Preference (solid arrows) and belief (dotted arrows).

for each knowledge cell a Lewis-style [32] counterfactual relation: a connected well-preorder, which can be viewed as a set of nested spheres, with the minimal elements as the innermost sphere. Compare also the conditional models of Burgess and Veltman [17; 40] (linked to Dynamic Doxastic Logic in [25]).

Baltag and Smets [7; 6] present logics of individual multi-agent belief and knowledge for belief revision models, and define belief update for this as a particular kind of action update in the style of [4], called action priority update. Here we sketch the extension to a system that also handles common knowledge and common conditional belief, and where the action update has belief change incorporated in it by means of relational substitution.

The set-up will be less general than in the logic LCC: in LCC no assumptions are made about the update actions, and so the accessibility relations

could easily deteriorate, e.g., as a result of updating with a lie. Since in the present set-up we make assumptions about the accessibilities (to wit, that they are locally connected well-preorders), we have to ensure that our update actions preserve these relational properties.

Consider the following slight modification of the PDL language (again assume p ranges over a set of basic propositions $Prop$ and i over a set of agents I):

$$\begin{aligned}\phi & ::= \top \mid p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid [\pi]\phi \\ \pi & ::= \sim_i \mid \leq_i \mid \geq_i \mid \rightarrow_i^\phi \mid \leftarrow_i^\phi \mid G \mid ?\phi \mid \pi_1; \pi_2 \mid \pi_1 \cup \pi_2 \mid \pi^*\end{aligned}$$

Call this language $\mathcal{L}_{\text{Pref}}$. This time, we treat \sim_i as a derived notion, by putting in an axiom that defines \sim_i as $\leq_i \cup \geq_i$. The intention is to let \sim_i be interpreted as the knowledge relation for agent i , \leq_i as the preference relation for i , \geq_i as the converse preference relation for i , \rightarrow_i^ϕ as the conditional belief relation defined from \leq_i as explained above, \leftarrow_i^ϕ as its converse, and G as global accessibility. We use \rightarrow_i as shorthand for \rightarrow_i^\top .

We have added a global modality G , and we will set up things in such way that $[G]\phi$ expresses that everywhere in the model ϕ holds, and that $\langle G \rangle \phi$ expresses that ϕ holds somewhere. It is well-known that adding a global modality and converses to PDL does not change its properties: the logic remains decidable, and satisfiability remains EXPTIME-complete [13].


The semantics of $\mathcal{L}_{\text{Pref}}$ is given relative to belief revision models as indicated above. Formula meaning $\llbracket \phi \rrbracket^{\mathbf{M}}$ and relational meaning $\llbracket \pi \rrbracket^{\mathbf{M}}$ are handled in the usual way. The interpretation of the knowledge relation of agent i is given by $\llbracket \sim_i \rrbracket^{\mathbf{M}} := \leq_i^{\mathbf{M}} \cup \geq_i^{\mathbf{M}}$, that for the preference relation of agent i by $\llbracket \leq_i \rrbracket^{\mathbf{M}} := \leq_i^{\mathbf{M}}$, that for the converse preference relation of agent i by its converse, that for the conditional belief of agent i by $\llbracket \rightarrow_i^\phi \rrbracket^{\mathbf{M}} := \rightarrow_i^{\llbracket \phi \rrbracket^{\mathbf{M}}}$, that for \leftarrow_i^ϕ by its converse. The global modality is interpreted as the universal relation, and test, sequential composition, choice and Kleene star are interpreted as usual.

The interplay between the modalities $[\sim_i]$ (knowledge) and $[\geq_i]$ (safe belief) is analysed by Baltag and Smets in [7], where they remark that the converse preference modality $[\geq_i]$ in belief revision models behaves like an S4.3 modality (reflexive, transitive and not forward branching), and lives happily together with the S5 modality for $[\sim_i]$.

To see how this all works out, let us have a look at the truth conditions for $[\rightarrow_i^\phi]\psi$. This is true in a world w in model \mathbf{M} if in all worlds v with $v \sim_i w$ and v minimal in $\llbracket \phi \rrbracket^{\mathbf{M}}$ under \leq_i it holds that ψ is true. This is indeed conditional

belief, relative to ϕ . Compare this with $[\geq_i]\psi$. This is true in a world w if in all worlds that are at least as preferred, ψ is true. Finally, $[\sim_i]\psi$ is true in w if ψ is true in all worlds, preferred or not, that i can access from w .

As a further example, consider a situation where Alexandru is drinking wine, while Jan does not know whether he is drinking wine or beer, and Sonja thinks that he is drinking tea. The actual situation is shaded grey, Jan's preference relation has solid lines, that of Sonja dotted lines. Reflexive arrows are not drawn, so Alexandru's preferences are not visible in the picture.



In the actual world it is true that Jan knows that Alexandru knows what Alexandru is drinking: $[\sim_j](\sim_a)w \vee (\sim_a)b$, and that Sonja believes Alexandru is drinking tea and that Alexandru knows it: $[s][\sim_a]t$. Under condition $\neg t$, however, Sonja has the belief in the actual world that Alexandru is drinking beer: $[\rightarrow_s^{\neg t}]b$. Moreover, Jan and Sonja have a common belief under condition $\neg t$ that Alexandru is drinking wine or beer: $[\rightarrow_j^{\neg t} \cup \rightarrow_s^{\neg t}; (\rightarrow_j^{\neg t} \cup \rightarrow_s^{\neg t})^*](w \vee b)$. As a final illustration, note that $[\leftarrow_s]\perp$ is true in a world if this is *not* among Sonja's most preferred worlds. Notice that if Sonja conditionalizes her belief to these worlds, she would believe that Alexandru is drinking beer: $[\rightarrow_s^{[s]\perp}]b$ is true in the actual world.

It should be clear from the example that this language is very expressive. To get at a complete logic for it, we need axioms and rules for propositional logic, S5 axioms for the global modality (Goranko and Passy [28]), axioms for forward connectedness of \geq and of \leq (see Goldblatt [27]), axioms for converse, relating \leq to \geq and \rightarrow to \leftarrow , as in temporal logic (Prior [35]), and the general axioms and rules for PDL (Seegerberg [37]). Finally, add the following definition of conditional belief in terms of knowledge and safe belief that can already be found in Boutillier [16] as an axiom:

$$[\rightarrow_i^\phi]\psi \quad \equiv \quad \langle \sim_i \rangle \phi \rightarrow \langle \sim_i \rangle (\phi \wedge [\geq_i](\phi \rightarrow \psi)).$$

This definition (also used in [7]) states that conditional to ϕ , i believes in ψ if either there are no accessible ϕ worlds, or there is an accessible ϕ world in

which the belief in $\phi \rightarrow \psi$ is safe. The full calculus for $\mathcal{L}_{\text{Pref}}$ is given in Figure 6.

Modus ponens	and axioms for propositional logic
Modal generalisation	From $\vdash \phi$ infer $\vdash [\pi]\phi$
Normality	$\vdash [\pi](\phi \rightarrow \psi) \rightarrow ([\pi]\phi \rightarrow [\pi]\psi)$
Inclusion of everything in G	$\vdash [G]\phi \rightarrow [\pi]\phi$
Reflexivity of G	$\vdash [G]\phi \rightarrow \phi$
Transitivity of G	$\vdash [G]\phi \rightarrow [G][G]\phi$
Symmetry of G	$\vdash \phi \rightarrow [G]\langle G \rangle \phi$
Knowledge definition	$\vdash [\sim_i]\phi \leftrightarrow [\leq_i \cup \geq_i]\phi$
Truthfulness of safe belief	$\vdash [\geq_i]\phi \rightarrow \phi$
Transitivity of safe belief	$\vdash [\geq_i]\phi \rightarrow [\geq_i][\geq_i]\phi$
\geq included in \leq^{\sim}	$\vdash \phi \rightarrow [\geq_i]\langle \leq_i \rangle \phi$
\leq included in \geq^{\sim}	$\vdash \phi \rightarrow [\leq_i]\langle \geq_i \rangle \phi$
Weak connectedness of \leq	$\vdash [\leq_i](\phi \wedge [\leq_i]\phi) \rightarrow \psi \vee [\leq_i](\psi \wedge [\leq_i]\psi) \rightarrow \phi$
Weak connectedness of \geq	$\vdash [\geq_i](\phi \wedge [\geq_i]\phi) \rightarrow \psi \vee [\geq_i](\psi \wedge [\geq_i]\psi) \rightarrow \phi$
Conditional belief definition	$\vdash [\rightarrow_i^{\phi}]\psi \leftrightarrow (\langle \sim_i \rangle \phi \rightarrow \langle \sim_i \rangle (\phi \wedge [\geq_i](\phi \rightarrow \psi)))$
\rightarrow included in \leftarrow^{\sim}	$\vdash \phi \rightarrow [\rightarrow_i^{\psi}]\langle \leftarrow_i^{\psi} \rangle \phi$
\leftarrow included in \rightarrow^{\sim}	$\vdash \phi \rightarrow [\leftarrow_i^{\psi}]\langle \rightarrow_i^{\psi} \rangle \phi$
Test	$\vdash [?\phi]\psi \leftrightarrow (\phi \rightarrow \psi)$
Sequence	$\vdash [\pi_1; \pi_2]\phi \leftrightarrow [\pi_1][\pi_2]\phi$
Choice	$\vdash [\pi_1 \cup \pi_2]\phi \leftrightarrow ([\pi_1]\phi \wedge [\pi_2]\phi)$
Mix	$\vdash [\pi^*]\phi \leftrightarrow (\phi \wedge [\pi][\pi^*]\phi)$
Induction	$\vdash (\phi \wedge [\pi^*](\phi \rightarrow [\pi]\phi)) \rightarrow [\pi^*]\phi$

Figure 6. Axiom system for $\mathcal{L}_{\text{Pref}}$.

Theorem 4.3. The axiom system for $\mathcal{L}_{\text{Pref}}$ is complete for belief revision models; $\mathcal{L}_{\text{Pref}}$ has the finite model property and is decidable.

Proof. Modify the canonical model construction for modal logic for the case of PDL, by means of Fischer-Ladner closures [21] (also see [13]). This gives a finite canonical model with the properties for \leq_i and \geq_i corresponding to the axioms (since the axioms for \leq_i and \geq_i are canonical). In particular, each \geq_i relation will be reflexive, transitive and weakly connected, each relation \leq_i will be weakly connected, and the \leq_i and \geq_i relations will be converses of each

other. Together this gives (Theorem 4.1) that the $\leq_i \cup \geq_i$ are equivalences. Since the canonical model has a finite set of nodes, each \leq_i relation is also well-founded. Thus, the canonical model is in fact a belief revision model. Also, the \rightarrow_i and \leftarrow_i relations are converses of each other, and related to the \geq_i relations in the correct way. The canonical model construction gives us for each consistent formula ϕ a belief revision model satisfying ϕ with a finite set of nodes. Only finitely many of the relations in that model are relevant to the satisfiability of ϕ , so this gives a finite model (see [13] for further details). Since the logic has the finite model property it is decidable. Q.E.D.

Since the axiomatisation is complete, the S5 properties of \sim_i are derivable, as well as the principle that knowledge implies safe belief:

$$[\sim_i]\phi \rightarrow [\geq_i]\phi.$$

The same holds for the following principles for conditional belief given in Board [14]:

Safe belief implies belief	$\vdash [\geq_i]\phi \rightarrow [\rightarrow_i^\psi]\phi$
Positive introspection	$\vdash [\rightarrow_i^\psi]\phi \rightarrow [\sim_i][\rightarrow\psi_i]\phi$
Negative introspection	$\vdash \neg[\rightarrow_i^\psi]\phi \rightarrow [\sim_i]\neg[\rightarrow_i^\psi]\phi$
Successful revision	$\vdash [\rightarrow_i^\phi]\phi$
Minimality of revision	$\vdash \langle \rightarrow_i^\phi \rangle \psi \rightarrow ([i^{\phi \wedge \psi}]\chi \leftrightarrow [\rightarrow_i^\phi](\psi \rightarrow \chi))$

We end with an open question: is \rightarrow_i^ϕ definable from \geq_i and \leq_i using only test, sequence, choice and Kleene star?

5 Combining Update and Upgrade

The way we composed knowledge and belief by means of regular operations may have a dynamic flavour, but appearance is deceptive. The resulting doxastic and epistemic ‘programs’ still describe what goes on in a static model. Real communicative action is changing old belief revision models into new ones. These actions should represent new hard information that cannot be undone, but also soft information like belief changes that can be reversed again later on. For this we can use update action by means of action models, with soft information update handled by means of action priority update [7; 6], or preference substitution as in [9]. Here we will propose a combination of these two.

Action models for belief revision are like belief revision models, but with the valuation replaced by a precondition map. We add two extra ingredients.

First, we add to each event a propositional substitution, to be used, as in LCC, for making factual changes to static models. Propositional substitutions are maps represented as sets of bindings

$$\{p_1 \mapsto \phi_1, \dots, p_n \mapsto \phi_n\}$$

where all the p_i are different. It is assumed that each p that does not occur in the lefthand side of a binding is mapped to p . The domain of a propositional substitution σ is the set $\{p \in Prop \mid \sigma(p) \neq p\}$. If $\mathbf{M} = (W, V, P)$ is a belief revision model and σ is an $\mathcal{L}_{\text{Pref}}$ propositional substitution, then $V_{\mathbf{M}}^{\sigma}$ is the valuation given by $\lambda w \lambda p. w \in \llbracket p^{\sigma} \rrbracket^{\mathbf{M}}$. In other words, $V_{\mathbf{M}}^{\sigma}$ assigns to w the set of basic propositions p such that p^{σ} is true in world w in model \mathbf{M} . \mathbf{M}^{σ} is the model with its valuation changed by σ as indicated. Next, we add relational substitutions, as defined in Section 3, one to each event. Thus, an action model for belief revision is a tuple $A = (E, \mathbf{pre}, \mathbf{P}, \mathbf{psub}, \mathbf{rsub})$ with E a non-empty finite set of events, \mathbf{psub} and \mathbf{rsub} maps from E to propositional substitutions and relational substitutions, respectively, and with \mathbf{rsub} subject to the following constraint:

If $e \sim_i f$ in the action model, then $\mathbf{rsub}(e)$ and $\mathbf{rsub}(f)$ have the same binding for i .

This ensures a coherent definition of the effect of relational substitution on a belief structure. The example in Figure 7 illustrates this. But note that the substitutions are subject to further constraints. In the action model in the




Figure 7. Unconstrained relational substitution creates havoc.

example, a single agent has a preference for $\neg p$ over p . In the update model, a substitution reverses the agent's preferences, but the agent cannot distinguish

this from an action where nothing happens. What should the result of the update look like? E.g., is there a preference arrow from (2, 3) to (1, 4)? This is impossible to answer, as action 3 asks us to reverse the preference and action 4 demands that we keep the initial preference. The constraint on substitutions rules out such dilemmas.

The relational substitution $\rho = \mathbf{rsub}(e)$ at event e in action model A is meant to be interpreted ‘locally’ at each world w in input model \mathbf{M} . If P is the preference map of \mathbf{M} , then let P_w^ρ be given by:

$$\begin{aligned} P_w^\rho(i) &:= P_i \cap |w|_{\llbracket \sim_i \rrbracket \mathbf{M}}^2 \text{ for } i \text{ not in the domain of } \rho, \\ P_w^\rho(i) &:= \llbracket \rho(i) \rrbracket^{\mathbf{M}} \cap |w|_{\llbracket \sim_i \rrbracket \mathbf{M}}^2 \\ &\text{for } i = i_j \text{ in the domain of } \rho. \end{aligned}$$

Thus, P_w^ρ is the result of making a change only to the local knowledge cell at world w of agent i (which is given by the equivalence class $|w|_{\llbracket \sim_i \rrbracket \mathbf{M}}$). Let

$$P^\rho(i) := \bigcup_{w \in W} P_w^\rho(i)$$

Then $P^\rho(i)$ gives the result of the substitution ρ on $P(i)$, for each knowledge cell $|w|_{\llbracket \sim_i \rrbracket \mathbf{M}}$ for i , and P^ρ gives the result of the substitution ρ on P , for each agent i .


Now the result of updating belief revision model $\mathbf{M} = (W, V, P)$ with action model $\mathbf{A} = (E, \mathbf{pre}, \mathbf{P}, \mathbf{psub}, \mathbf{rsub})$ is given by $\mathbf{M} \otimes \mathbf{A} = (W', V', P')$, where

- $W' = \{(w, e) \mid w \in W, e \in E, \mathbf{M}, w \models \mathbf{pre}(e)\}$,
- $V'(w, e) = V^\sigma(w)$,
- $P'(i)$ is given by the anti-lexicographical order defined from $P^\rho(i)$ and $\mathbf{P}(i)$ (see [7; 6]).

With these definitions in place, what are reasonable substitutions? A possible general form for a preference change could be a binding like this:

$$\geq_i \mapsto [\phi_1, \phi_2, \dots, \phi_n].$$

This is to be interpreted as an instruction to replace the belief preferences of i in the local knowledge cells by the new preference relation that prefers the ϕ_1 states above everything else, the $\neg\phi_1 \wedge \phi_2$ above the $\neg\phi_1 \wedge \neg\phi_2$ states, and so on, and the $\neg\phi_1 \wedge \neg\phi_2 \wedge \dots \wedge \neg\phi_{n-1} \wedge \phi_n$ states above the $\neg\phi_1 \wedge \neg\phi_2 \wedge \dots \wedge \neg\phi_n$ states. Such relations are indeed connected well-preorders.




Note that we can take $[\phi_1, \phi_2, \dots, \phi_n]$ as an abbreviation for the following doxastic program:

$$\begin{aligned}
 (\sim_i; ?\phi_1) &\cup (?\neg\phi_1; \sim_i; ?\neg\phi_1; ?\phi_2) \\
 &\cup (?\neg\phi_1; ?\neg\phi_2; \sim_i; ?\neg\phi_1; ?\neg\phi_2; ?\phi_3) \\
 &\cup \dots \\
 &\cup (?\neg\phi_1; \dots; ?\neg\phi_n; \sim_i; ?\neg\phi_1; \dots; ?\neg\phi_{n-1}; ?\phi_n)
 \end{aligned}$$


In general we have to be careful (as is also observed in [9]). If we have a connected well-preorder then adding arrows to it in the same knowledge cell may spoil its properties. Also, the union of two connected well-preorders need not be connected. So here is a question: what is the maximal sublanguage of doxastic programs that still guarantees that the defined relations are suitable preference relations? Or should belief revision models be further constrained to guarantee felicitous preference change? And if so: how?

6 Examples

Global amnesia: the event of agent a (Jason Bourne) forgetting all his beliefs, with everyone (including himself) being aware of this, is represented by the following action model (for the case of three agents a, b, c):



Alzheimer: the event of agent a forgetting everything, with the others being aware of this, while a wrongly believes that nothing has happened. It is tempting to model this with the following update model:



Note however that this does not satisfy the constraint on relation update (the two actions are connected, but the substitution for a is not the same), so it may result in incoherent models.

Lacunar amnesia (specific forgetting): forgetting everything about p . One way to model this is by means of an action model with a single action, accessible to all, with the relational substitution

$$\geq_i \mapsto (\geq_i \cup \sim_i; ?\neg p)^*$$

This will effectively add best-world arrows from everywhere in the knowledge cell to all $\neg p$ worlds.

Confession of faith in p , or publicly accepting p : an action model with a single action, accessible to all, with the relational substitution

$$\geq_i \mapsto (\geq_i \cup (? \neg p; \sim_i; ?p))^*.$$

This will make the p worlds better than the $\neg p$ worlds everywhere.

Submission to a guru: the act of adopting the belief of someone else, visible to all. A problem here is that the guru may know more than I do, so that the guru's preferences within my knowledge cell may not be connected. This means that the substitution $\leq_i \mapsto \leq_j$ — the binding that expresses that i takes over j 's beliefs — may involve growth or loss of knowledge for i . Consider the example of the wine-drinking Alexandru again: if Jan were to take over Sonja's beliefs, he would lose the information that Alexandru is drinking an alcoholic beverage.

Conformism: adopting the common beliefs of a certain group, visible to all: an action model with a single action accessible to all, with the following substitution for conformist agent i :

$$\geq_i \mapsto (\geq_i \cup \geq_j); (\geq_i \cup \geq_j)^*.$$

Belief coarsening: the most preferred worlds remain the most preferred, the next preferred remain the next preferred, and all further distinctions are erased. An action model with a single action accessible to all, and the following substitution for agent i :

$$\geq_i \mapsto \rightarrow_i \cup ?\top.$$

The union with the relation $?\top$ has the effect of adding all reflexive arrows, to ensure that the result is reflexive again.

Van Benthem's $\uparrow \phi$ is handled by a substitution consisting of bindings like this:

$$\geq_i \mapsto (? \phi; \sim_i ? \neg \phi) \cup (? \phi; \geq_i ? \phi) \cup (? \neg \phi; \geq_i ? \neg \phi).$$

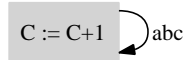
This is an alternative for an update with an action model that has $\neg \phi <_B \phi$. The example shows that conservative upgrade is handled equally well by action priority updating and by belief change via substitution. But belief change by substitution seems more appropriate for 'elite change'. For this we need a test for being in the best ϕ world that i can conceive, by means of the Panglossian formula $\langle \leftarrow_i^\phi \rangle \top$. The negation of this allows us to define elite change like this:

$$\geq_i \mapsto \rightarrow_i^\phi \cup (\geq_i; ?[\leftarrow_i^\phi] \perp).$$

This promotes the best ϕ worlds past all other worlds, while leaving the rest of the ordering unchanged. Admittedly, such an operation could also be performed using action priority updating, but it would be much more cumbersome.

7 Further Connections

To connect up to the work of Bonanno again, what about time? Note that perceiving the ticking of a clock can be viewed as information update. A clock tick constitutes a change in the world, and agents can be aware or unaware of the change. This can be modelled within the framework introduced above. Let t_1, \dots, t_n be the clock bits for counting ticks in binary, and let $C := C + 1$ be shorthand for the propositional substitution that is needed to increment the binary number t_1, \dots, t_n by 1. Then public awareness of the clock tick is modelled by:



Thus, perception of the ticking of a clock can be modelled as 'being in tune with change in the world'. Still, this is not quite the same as the 'next instance' operator \bigcirc , for the DEL framework is specific about what happens during the clock tick, while \bigcirc existentially quantifies over the change that takes place, rather in the spirit of [2].

In belief revision there is the AGM tradition, and its rational reconstruction in dynamic doxastic logic à la Segerberg. Now there also is a modal version in Bonanno style using temporal logic. It is shown in [10] that temporal logic has greater expressive power than DEL, which could be put to

use in a temporal logic of belief revision (although Bonanno's present version does not seem to harness this power). As an independent development there is dynamic epistemic logic in the Amsterdam/Oxford tradition, which was inspired by the logic of public announcement, and by the epistemic turn in game theory, à la Aumann. Next to this, and not quite integrated with it, there is an abundance of dynamic logics for belief change based on preference relations (Spohn, Shoham, Lewis), and again the Amsterdam and Oxford traditions. I hope this contribution has made clear that an elegant fusion of dynamic epistemic logic and dynamic logics for belief change is possible, and that this fusion allows to analyze AGM style belief revision in a multi-agent setting, and integrated within a powerful logic of communication and change.

Acknowledgements Thanks to Alexandru Baltag, Johan van Benthem and Floor Sietsma for helpful comments on several earlier versions of the paper, and for constructive discussions. Hans van Ditmarsch sent me his last minute remarks. I am grateful to the editors of this volume for their display of the right mix of patience, tact, flexibility and firmness (with Robert van Rooij exhibiting the first three of these qualities and Krzysztof Apt the fourth). The support of the NWO Cognition Programme to the 'Games, Action and Social Software' Advanced Studies project is gratefully acknowledged.

References

- [1] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [2] P. Balbiani, A. Baltag, H. Van Ditmarsch, A. Herzig, T. Hoshi, and T. Santos De Lima. What can we achieve by arbitrary announcements? A dynamic take on Fitch's knowability. In *Theoretical Aspects of Rationality and Knowledge (TARK-07)*, 2007.
- [3] A. Baltag and L.S. Moss. Logics for epistemic programs. *Synthese*, 139(2):165–224, 2004.
- [4] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In I. Bilboa, editor, *Proceedings of TARK'98*, pages 43–56, 1998.

- [5] A. Baltag and S. Smets. Conditional doxastic models: A qualitative approach to dynamic belief revision. *Electronic Notes in Theoretical Computer Science (ENTCS)*, 165:5–21, 2006.
- [6] A. Baltag and S. Smets. The logic of conditional doxastic actions, 2008. This volume.
- [7] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Texts in Logic and Games*. Amsterdam University Press, 2008. To appear.
- [8] J. van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 2:129–155, 2007.
- [9] J. van Benthem and F. Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics*, 14(2), 2004.
- [10] J. van Benthem and E. Pacuit. The tree of knowledge in action: Towards a common perspective. In Guido Governatori, Ian Hodkinson, and Yde Venema, editors, *Advances in Modal Logic, Volume 6*, pages 87–106. College Publications, 2006. available online at <http://www.aiml.net/volumes/volume6/>.
- [11] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
- [12] J. van Benthem, S. van Otterloo, and O. Roy. Preference logic, conditionals, and solution concepts in games. In H. Lagerlund, S. Lindström, and R. Sliwinski, editors, *Modality Matters*, pages 61–76, 2006.
- [13] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.
- [14] O. Board. Dynamic interactive epistemology. *Games and Economic Behaviour*, 49:49–80, 2002.
- [15] G. Bonanno. Belief revision in a temporal framework, 2008. This volume.
- [16] C. Boutilier. Toward a logic of qualitative decision theory. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Proceedings of the 4th International Conference on Principle of Knowledge Representation and Reasoning (KR-94)*, pages 75–86. Morgan Kaufmann, 1994.

- [17] J.P. Burgess. Quick completeness proofs for some logics of conditionals. *Notre Dame Journal of Formal Logic*, 22(1):76–84, 1981.
- [18] B.F. Chellas. *Modal Logic: An Introduction*. Cambridge University Press, 1980.
- [19] H.P. van Ditmarsch. Belief change and dynamic logic. In J. Delgrande, J. Lang, H. Rott, and J.-M. Tallon, editors, *Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics*, number 05321 in Dagstuhl Seminar Proceedings, Dagstuhl, 2005. IBFI, Schloss Dagstuhl.
- [20] H.P. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2006.
- [21] M.J. Fischer and R.E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, 1979.
- [22] P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. MIT Press, Cambridge Mass, 1988.
- [23] P. Gärdenfors and H. Rott. Belief revision. In D.M. Gabbay, C.J. Hogger, and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 4. Oxford University Press, Oxford, 1995.
- [24] J. Gerbrandy. Dynamic epistemic logic. In L.S. Moss et al., editors, *Logic, Language and Information, Vol. 2*. CSLI Publications, Stanford, 1999.
- [25] P. Girard. From onions to broccoli: Generalizing Lewis’s counterfactual logic. *Journal of Applied Non-Classical Logics*, 17(3):213–229, 2007. Special Issue on Belief Revision and Dynamic Logic.
- [26] P. Girard. *Modal Logic for Belief and Preference Change*. PhD thesis, Stanford University, January 2008.
- [27] R. Goldblatt. *Logics of Time and Computation, Second Edition, Revised and Expanded*, volume 7 of *CSLI Lecture Notes*. CSLI, Stanford, 1992 (first edition 1987). Distributed by University of Chicago Press.

- [28] V. Goranko and S. Passy. Using the universal modality: gains and questions. *Journal of Logic and Computation*, 2:5–30, 1992.
- [29] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
- [30] Joseph Y. Halpern. Defining relative likelihood in partially-ordered preferential structures. *Journal of Artificial Intelligence Research*, 7:1–24, 1997.
- [31] D. Kozen and R. Parikh. An elementary proof of the completeness of PDL. *Theoretical Computer Science*, 14:113–118, 1981.
- [32] D. Lewis. *Counterfactuals*. Basil Blackwell, Oxford, 1973.
- [33] F. Liu. *Changing for the Better: Preference Dynamics and Agent Diversity*. PhD thesis, ILLC, Amsterdam, February 2008.
- [34] J. A. Plaza. Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216, 1989.
- [35] A. Prior. *Past, Present and Future*. Clarendon Press, Oxford, 1967.
- [36] H. Rott. Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In Henrik Lagerlund and Sten Lindström und Rysiek Sliwinski, editors, *Modality Matters. Twenty-Five Essays in Honour of Krister Segerberg*, volume 53 of *Uppsala Philosophical Studies*, pages 359–384. Uppsala Universitet, 2006.
- [37] K. Segerberg. A completeness theorem in the modal logic of programs. In T. Traczyck, editor, *Universal Algebra and Applications*, pages 36–46. Polish Science Publications, 1982.
- [38] K. Segerberg. Belief revision from the point of view of doxastic logic. *Bulletin of the IGPL*, 3:534–554, 1995.
- [39] K. Segerberg. Two traditions in the logic of belief: bringing them together. In H.J. Ohlbach and U. Reyle, editors, *Logic, language and reasoning*, pages 135–147. Kluwer Academic Publishers, 1999.
- [40] F. Veltman. *Logics for Conditionals*. PhD thesis, University of Amsterdam, Amsterdam, 1985.