Belief change and dynamic logic*

Hans van Ditmarsch, University of Otago, New Zealand

Abstract

In this paper we compare standard ways to perform belief change with attempts to model such change with dynamic modal operators. We address both belief expansion, belief contraction, belief revision, and update. Public announcement logic is an implementation of a belief expansion operator in a dynamic epistemic setting. The postulates of success and minimal change cannot be satisfied in their original AGM formulation. Dynamic doxastic logic provides various implementations of a belief revision operator in a dynamic epistemic setting. We provide an example of the application of such a dynamic doxastic semantics, which can also be seen as a realization of a proposal by Spohn.

The purpose of this paper is to provide an overview of different approaches to dynamic logics for belief change, including some of our own contributions to the area, in a way that is accessible to an interdisciplinary audience. It does not contain new technical results not published elsewhere. However, the comparison of different approaches provides, we hope, new insights not seen before. A final version of this paper may include new technical results.

1 Overview

Section 4 contains the meat of this paper. Section 2 gives logical preliminaries necessary to understand that section. It introduces the public announcement logic which in Section 4 is referred to as an implementation of a belief expansion operator. Section 3 gives an overview of the belief revision terminology that is assumed in Section 4. Readers familiar with either one, or the other, or both, may choose to skip one or both sections.

In Section 4 we first address the early, non-epistemic, approaches to belief revision by van Benthem and (at that time his student) de Rijke. Then, we

*This draft paper is based on the presentation given during the Dagstuhl Seminar 05321 'Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics'. I am indebted to the Dagstuhl organization for their supportive environment and to the seminar organizers Jim Delgrande, Jérôme Lang, Hans Rott, and Jean-Marc Tallon for inviting me to this workshop and for inspiring discussions and presentations. I thank Jonathan Ben-Naim for organizing the after-conference formalities and proceedings. Sections 3 and 4 of this draft paper are mainly based on material that may later be included in a chapter in the forthcoming Handbook on the Philosophy of Information.

Dagstuhl Seminar Proceedings 05321
Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics
http://drops.dagstuhl.de/opus/volltexte/2005/337

cover the approach by Segerberg and his collaborators Lindström and Rabinowicz. This is followed by some elementary examples on how to interpret belief change in a strictly dynamic epistemic setting (as in work by Baltag, Gerbrandy, van Ditmarsch, and others). Public announcement logic (Plaza, Baltag et al., Gerbrandy) is then illustrated as an implementation of an expansion operator, followed by an example of a dynamic doxastic logic that implements a revision operator (Aucher, van Ditmarsch). The concluding remarks include references to the literature on factual updates in a dynamic epistemic setting (Baltag, van Benthem, Kooi, ...). All proper references are found in the sections themselves.

2 Preliminaries: public announcement logic

Public announcement logic is a dynamic epistemic logic and is an extension of standard multi-agent epistemic logic. Intuitive explanations of the epistemic part of the semantics can be found in [FHMV95, vdHV02, vDvdHK06]. We give a concise overview of, in that order, the language, the structures on which the language is interpreted, and the semantics.

Given are a set of agents N and a set of atoms P. The language of public announcement logic is inductively defined as

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \psi) \mid K_n\varphi \mid C_G\varphi \mid [\varphi]\psi$$

where $p \in P$, $n \in N$, and $G \subseteq N$ are arbitrary. For $K_n\varphi$, read 'agent n knows formula φ'. For $C_G\varphi$, read 'group of agents G commonly know formula φ'. For $[\varphi]\psi$, read 'after public announcement of φ, formula ψ (is true)'. Note that in Section 4 we write $[\oplus\varphi]\psi$ instead of $[\varphi]\psi$, as we present public announcement logic there as an example of belief expansion.

Next, we introduce the structures. An *epistemic model* $M = \langle S, \sim, V\rangle$ consists of a *domain* S of (factual) *states* (or 'worlds'), *accessibility* $\sim : N \to \mathcal{P}(S \times S)$, and a *valuation* $V : P \to \mathcal{P}(S)$. For $s \in S$, (M, s) is an *epistemic state* (also known as a pointed Kripke model). For $\sim(n)$ we write $\sim_n$, and for $V(p)$ we write V_p. So, access $\sim$ can be seen as a set of equivalence relations $\sim_n$, and V as a set of valuations V_p. Given two states s, s' in the domain, $s \sim_n s'$ means that s is indistinguishable from s' for agent n on the basis of its information. The group accessibility relation $\sim_G$ is the transitive and reflexive closure of the union of all access for the individuals in G: $\sim_G \equiv (\bigcup_{n\in G} \sim_n)^*$. This is access to interpret common knowledge for group G.

Finally, we give the semantics. Assume an epistemic model $M = \langle S, \sim, V\rangle$.

$$\begin{array}{lll}
M, s \models p & \text{iff} & s \in V_p \\
M, s \models \neg\varphi & \text{iff} & M, s \not\models \varphi \\
M, s \models \varphi \wedge \psi & \text{iff} & M, s \models \varphi \text{ and } M, s \models \psi \\
M, s \models K_n\varphi & \text{iff} & \text{for all } t \in S : s \sim_n t \text{ implies } M, t \models \varphi \\
M, s \models C_G\varphi & \text{iff} & \text{for all } t \in S : s \sim_G t \text{ implies } M, t \models \varphi \\
M, s \models [\varphi]\psi & \text{iff} & M, s \models \varphi \text{ implies } M|\varphi, s \models \psi
\end{array}$$

Epistemic model $M|\varphi = \langle S', \sim', V' \rangle$ is defined as

$$\begin{array}{lcl} S' & = & \{s' \in S \mid M, s' \models \varphi\} \\ \sim'_n & = & \sim_n \cap (S' \times S') \\ V'_p & = & V_p \cap S' \end{array}$$

The dynamic modal operator $[\varphi]$ is interpreted as an epistemic state transformer. Announcements are assumed to be truthful, and this is commonly known by all agents. Therefore, the model $M|\varphi$ is the model M restricted to all the states where φ is true, including access between states. The dual of $[\varphi]$ is $\langle\varphi\rangle$: $M, s \models \langle\varphi\rangle\psi$ iff $M, s \models \varphi$ and $M|\varphi, s \models \psi$.

Formula φ is valid on model M, notation $M \models \varphi$, if and only if for all states s in the domain of M: $M, s \models \varphi$. Formula φ is valid, notation $\models \varphi$, if and only if for all models M: $M \models \varphi$. Logical consequence $\Psi \models \varphi$ is defined as "for all (M, s), if $M, s \models \psi$ for all $\psi \in \Psi$, then $M, s \models \varphi$." For $\{\psi\} \models \varphi$, write $\psi \models \varphi$.

A proof system for this logic is presented, and shown to be complete, in [BMS98], with precursors – namely for public announcement logic *without* common knowledge – in [Pla89, Ger99]. For a concise completeness proof, see [vDvdHK06].

The notion of belief, also referred to in this paper, is weaker than that of knowledge, in that belief of a proposition does not entail its truth, in other words, $B\varphi \rightarrow \varphi$ is *not* valid for belief. In that case, and also for even weaker notions of modal operators, the accessibility relation is not an equivalence relation. Instead of $s \sim_n t$, one then writes (even though we have avoided that level of technical detail in this paper) $R_n(s, t)$, or $s \rightarrow_n t$, or $(s, t) \in R_n$. Note that the example structures in Figure 1 carry such more arbitrary accessibility relations. Group access $\rightarrow_G$ is similarly to $\sim_G$ defined as $\rightarrow_G \equiv (\bigcup_{n \in G} \rightarrow_n)^*$, this access is then used to interpret a notion of common belief. There are different ways to formalize public announcements in this more general setting, to which we will pay no attention in this paper. For examples, see [BM04, Ger99]. More complex dynamics than announcements are also not addressed in this paper. For that, see those same references, or [vDvdHK06].

3 Preliminaries: belief revision

The traditional emphasis in what is known as the area of 'belief revision' is *theory revision*; how to change a deductively closed set of formulas $\mathcal{K}$ into another deductively closed set of formulas. Overview publications for this area are [AGM85] and [Gär88]. A theory typically consists of objective, i.e. non-epistemic, 'beliefs' that are changed relative to expansion, contraction, or revision, and also typically from the point of view of a single agent. Note that 'belief' means nothing but 'formula' here: there is no explicit representation of the status of such a formula as 'believed by the agent', as in epistemic logic. In the case of an *expansion*, new information described by a formula φ is incorporated into a theory $\mathcal{K}$ by somehow 'adding' φ to the theory. For the result we write $\mathcal{K} \oplus \varphi$. In case the negation was already in $\mathcal{K}$, the result will be the inconsistent theory $\mathcal{K}_\perp$ that consists of all formulas. The result of a *contraction* should be that the formula with which the theory $\mathcal{K}$ is contracted is longer be believed, for which we write $\varphi \notin \mathcal{K} \ominus \varphi$. Contraction with a validity cannot be successful, as all validities are in all theories. Contraction with a validity therefore leaves a theory unchanged. In the case of a revision, the negation $\neg\varphi$ of the revision formula φ is typically in the theory $\mathcal{K}$ but for the revision to 'succeed' the revised (consistent) theory should (also) be consistent. A process of mere expansion therefore does not work, and one first has to contract the theory in a way that removes $\neg\varphi$ from it and all its dependencies. This can then be followed by an expansion with φ. For the result of the revision of theory $\mathcal{K}$ with formula φ we write $\mathcal{K} \circledast \varphi$. The Levi-identity states that $\mathcal{K} \circledast \varphi = \mathcal{K} \ominus \neg\varphi \oplus \varphi$: a revision can be seen as a contraction followed by an expansion. The 'theories' $\mathcal{K}$ that we are changing can have any shape, such as first-order theories. Here, we restrict ourselves to propositional theories, and we investigate possibilities to extend this to theories of propositional modal formulas.

Yet another issue in traditional belief revision comes under the name of 'update'. An *update*—unfortunately a clash cannot be avoided with the more general meaning of that term in dynamic epistemic logic, where it incorporates belief revision as well—is a *factual* change, as opposed to a belief change in the three previously distinguished notions. The latter merely express a different agent stance towards a non-changing world, but in an 'update' the world itself changes. The standard reference for that is [KM91]. We will pay only minor attention to such updates in this paper.

4 Belief Change and Dynamic Logic

Belief change with dynamic non-epistemic logic The different 'theory change operators' $\oplus$, $\ominus$, and $\circledast$ can be reinterpreted as dynamic modal operators. A straightforward way to implement that, is some logic in which $[\circledast\varphi]\psi$ expresses that after revision with φ, ψ holds—where ψ actually means, as in [AGM85], 'ψ is believed by the agent'. This approach was suggested by van Benthem in [vB94][1] and further developed by de Rijke in [dR94]. They propose a semantical counterpart of a total order on theories, in the form of an 'updating' and 'downdating' relation between states or worlds, standing for theories, and interpret the modal operator as a transition in such a structure according to these relations. 'Updating' models expansion: it relates the current state to states that result from expansion. 'Downdating' models contraction. It relates states that result from contraction to the current state. Revision is indeed downdating followed by updating. In this overview we focus on approaches that extend epistemic logics, therefore we do not give more details on this non-epistemic approach.

[1] It is only one of many topics covered in that publication, namely Section 6, pages 714–715, 'Cognitive procedures over information patterns'. Note this work is similar to a 1991 technical report.

Belief change with dynamic epistemic logic In the approach by Segerberg and collaborators [LR99a, Seg99b, Seg99a, LR99b] beliefs are represented explicitly. We now identify a theory $\mathcal{K}$ with the believed formulas (or some subset of the believed formulas) in an epistemic state: $\mathcal{K} = \{\psi \mid M, s \models B\psi\}$. As in [dR94] they express belief change with dynamic modal operators $[\oplus\varphi]$, $[\ominus\varphi]$, and $[\circledast\varphi]$. In a typical revision where we have that $\neg\varphi \in \mathcal{K}$, $\varphi \in \mathcal{K} \circledast \varphi$, and $\neg\varphi \notin \mathcal{K} \circledast \varphi$, we now get

- $M, s \models B\neg\varphi$
- $M, s \models [\circledast\varphi]B\varphi$
- $M, s \models [\circledast\varphi]\neg B\neg\varphi$

For contraction, we want that in case $M, s \models B\varphi$, after contraction φ is no longer believed, i.e., $M, s \models [\ominus\varphi]\neg B\varphi$. Similarly, for expansion we aim to achieve $M, s \models [\oplus\varphi]B\varphi$.

This approach is known as *dynamic doxastic logic* or **DDL**. Similar to [dR94] it presumes a transition relation between states representing theories, but this is now differently realized, namely using what is known as a Segerberg-style semantics wherein factual and epistemic information—under the terms of world component and doxastic component—are strictly separated. A dynamic operator is interpreted as a transition along the 'lines' of minimal theory change set out by this given structure, with the additional restriction that the transitions describe epistemic (doxastic) change only, and not factual change. This restriction is enforced by not allowing the 'world component' to change in the transition relation but only the 'doxastic component' [LR99a, p.18].

There are now two options: either we restrict ourselves to beliefs in *objective* (boolean, non-epistemic) formulas, and we get what is known as basic **DDL**, as in [LR99a, Seg99b]. Or we allow higher-order beliefs, as in, for example, the public announcement logic introduced in Section 2. We thus get 'full' or 'unlimited' **DDL**, also discussed in [LR99a] but mainly in [LR99b]. Without the restriction to belief of objective formulas, a number of problems come to the fore related to higher order belief, knowledge growth, 'success' of revision, and multi-agent belief. We address these issues in the dynamic epistemic setting in this chapter, that provides a 'third way' given the two views on belief revision with dynamic logic presented so far. In dynamic epistemics, unlike the two approaches to dynamic belief revision already presented, the transition that interprets the dynamic operators is *induced* by the current information state and the revision formula, and does not *assume* such a transition relation. We demonstrate the various possible transitions by a simple example.

Examples of belief change with dynamic epistemic logic Consider expressing and changing uncertainty about the truth of a single fact p, and assume an information state where the agent (whose beliefs are interpreted by the unlabeled accessibility relation depicted) may be uncertain about p and where p is actually false (indicated by 'designating' the actual state by underlining it). Figure 1 lists all conceivable sorts of belief change.

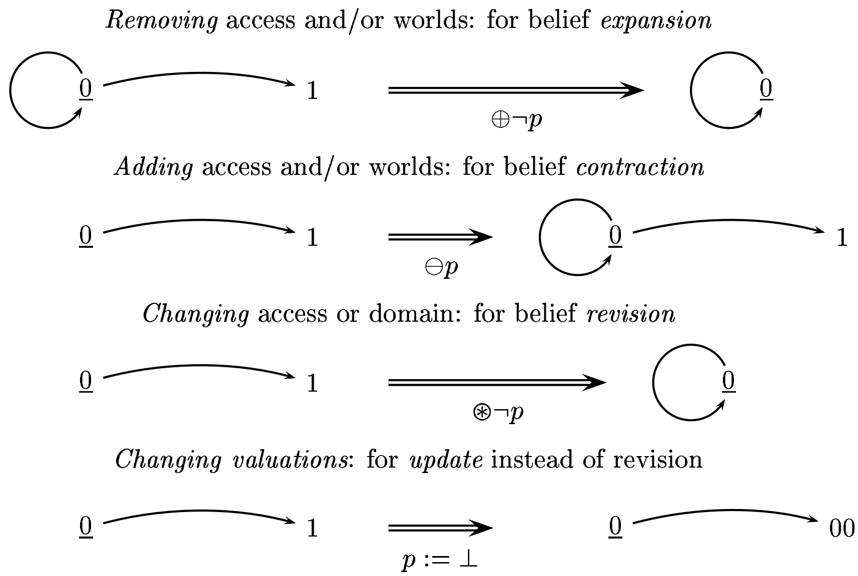

Figure 1: Possible changes of belief mirrored in Kripke structure transitions

In the top structure, uncertainty about the fact p (i.e., absence of belief in p and absence of belief in $\neg p$) is changed into belief in $\neg p$. On the left, $\neg Bp$ is true, and on the right $B\neg p$. In the second from above, belief in p is weakened to uncertainty about p, and in the third from above we change from Bp to $B\neg p$. Note that also in this semantic setting of Kripke-structure transformation, belief revision can again be seen as a contraction followed by an expansion, so we may in principle consider semantic alternatives for the Levi-identity. The last information state transition in Figure 1 depicts factual change. The state with changed valuation has suggestively been renamed from 1 to 00, although formally, of course, it is only the valuation of a named state that changes. The 'assignment' or substitution $p := \perp$ indicates that the valuation of atom p is revised into the valuation of the assigned formula. As this is $\perp$, the new valuation of p (seen as a subset of the domain) is now the empty set of states.

Public announcement as belief expansion The public announcement logic already discussed in detail can be seen as an implementation of a belief expansion operator for higher-order belief (i.e., both fully introspective beliefs of a single agent but also beliefs of agents in a group about the beliefs of other agents in that group), where the next information state is computed from 'merely' the current information state and the expansion formula. This computation is straightforward for expansion, as restricting the domain or accessibility relation can easily be seen as structurally related to the existing model. The semantics of the public announcements already presented operates just like that: it restricts a model M to the submodel $M|\varphi$ consisting of the worlds where the announcement is true. From here on, for $[\varphi]\psi$ we write $[\oplus\varphi]\psi$, and we focus on knowledge K.

Knowledge growth In such a higher-order setting we cannot maintain the expansion postulates. First, we have to revise our ideas about 'minimal change'. In particular, it can no longer be maintained that expanded theories contain their predecessors:

Identify a theory $\mathcal{K}$ as before with the set of known formulas in an information state: $\{\psi \mid M, s \models K\psi\}$. Let $M, s \models \varphi$. Suppose $\mathcal{K} \subset \mathcal{K} \oplus \varphi$, then there must be at least one formula ψ such that $\psi \in \mathcal{K} \oplus \varphi$ but $\psi \notin \mathcal{K}$. From $\psi \in \mathcal{K} \oplus \varphi$ follows by positive introspection that $K\psi \in \mathcal{K} \oplus \varphi$. From $\psi \notin \mathcal{K}$ follows by negative introspection that $\neg K\psi \in \mathcal{K}$. From $\neg K\psi \in \mathcal{K} \subset \mathcal{K} \oplus \varphi$ and $K\psi \in \mathcal{K} \oplus \varphi$ follows a contradiction.

Therefore, strict knowledge *growth* is contradictory for introspective agents (we did not use the truth axiom $K\varphi \rightarrow \varphi$, so the results hold for introspective belief as well), as observed by many authors: one does not care to preserve ignorance of the expansion formula, when expanding a theory. Therefore, one cannot adhere closely to the *expansion postulate* which states that $\mathcal{K} \subseteq \mathcal{K} \oplus \varphi$ (also known as postulate $\mathcal{K} \oplus 3$). Fortunately, knowledge *change* in a way that reflects the ideas behind expansion is still possible. And also, knowledge growth is possible for *fragments* of the public announcement language; for an example, see [vvK05].

Success The *success postulate*, which states that $\varphi \in \mathcal{K} \oplus \varphi$ (also known as postulate $\mathcal{K} \oplus 2$), cannot be maintained either. The most basic example illustrating that, is an announcement of the Moore-sentence $p \wedge \neg Kp$. This sentence cannot be believed, or known, after its announcement [Hin62, Moo42][2]. Now

[2] Hintikka's 'Knowledge and belief' [Hin62, p.64] provides a list of excellent references on this topic. This also reveals an interesting development of the notion—we explain this using the logical notation used in this chapter. In [Moo12, p.78] Moore writes that if I assert a proposition φ, I *express* or *imply* that I *think* or *know* φ, in other words I express $B\varphi$ or $K\varphi$. But φ cannot be said to *mean* $B\varphi$ [Moo12, p.77] as this would cause, by substitution, an infinite sequence $BB\varphi$, $BBB\varphi$, ad infinitum. "But thus to believe that somebody believes, that somebody believes, that somebody believes ... quite indefinitely, without *ever* coming to anything which is what is believed, is to believe nothing at all" [Moo12, p.77]. All this is in the context of a discussion on whether moral judgements are judgements about our feelings, or about our beliefs. Moore does not state in [Moo12] (to our knowledge) that $\varphi \wedge \neg B\varphi$ cannot be believed. In Moore's "A reply to my critics", a chapter in the 'Library of Living Philosophers' volume dedidated to him, he writes " 'I went to the pictures last Tuesday, but I don't believe that I did' is a perfectly absurd thing to say, although *what* is asserted is something which is perfectly possibly logically" [Moo42, p.543]. The absurdity follows from the implicature 'asserting φ implies $B\varphi$' pointed out in [Moo12]. In other words, $B(p \wedge \neg Bp)$ is 'absurd' for the example of factual information p. As far as we know, this is the first full-blown occurrence of a Moore-sentence. Then in [Moo44, p.204] Moore writes " 'I believe he has gone out, but he has not' is absurd. This, though absurd, is not self-contradictory; for it may quite well be true." This is an example of $\neg p \wedge Bp$. Together with [Moo42] it also sufficiently shows, we think, that Moore really had either of the general forms $\varphi \wedge \neg B\varphi$ or $\varphi \wedge B\neg\varphi$ in mind. Note that he does not claim that $B(\varphi \wedge \neg B\varphi)$ is inconsistent ('self-contradictory') as such, but 'only'

this is first of all for the obvious reason that it is a Moore-sentence, and by definition that is a sentence that cannot be believed, but it should be pointed out that as an announcement it can very well be true and therefore executed: after it, p is (publicly) known, which in fact entails the *negation* of $p \wedge \neg Kp$. But this means that expansion with $p \wedge \neg Kp$ *cannot* be successful: yet another barrier to satisfy the AGM postulates for higher-order belief expansion. Gerbrandy [Ger99] calls this phenomenon an *unsuccessful update*; [Ger05] is similar to [Ger99]. The matter is also taken up in [vDK05].

For truthful public announcements, the formulas φ that always become known after their announcement can be properly said to be the *successful formulas* and characterized by the validity of $[\oplus\varphi]\varphi$. This entails the validity of $\varphi \rightarrow [\oplus\varphi]C_N\varphi$ (in this setting where common knowledge of a formula entails its truth). The latter says that if φ is true, announcing φ makes it common knowledge, which more properly grasps what 'success' means in a higher-order setting. An intriguing question is: *Which* formulas are successful? An answer to that question would address knowledge expansion satisfactorily in this higher-order setting. But the answer is as yet unclear. Obvious inductive definitions fail. Even when both φ and ψ are successful, $\neg\varphi$ may be unsuccessful (for $\varphi = p \wedge \neg Kp$), $\varphi \wedge \psi$ may be unsuccessful (for $\varphi = p$ and $\psi = \neg Kp$), and as well $[\oplus\varphi]\psi$ and $\varphi \rightarrow \psi$ may be unsuccessful.

There are relevant successful fragments of the language. For example, *public knowledge formulas* are successful: $[\oplus C_N\varphi]C_N\varphi$ is valid. This follows from bisimulation invariance under point-generated submodel constructions. Another successful fragment form the *preserved formulas* (introduced for the language without announcements by van Benthem in [vB02]) that are inductively defined as $\varphi ::= p \mid \neg p \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid K_n\varphi \mid C_G\varphi \mid [\oplus\neg\varphi]\psi$ (where $G \subseteq N$). From $\varphi \rightarrow [\oplus\psi]\varphi$ for arbitrary ψ, follows $\varphi \rightarrow [\oplus\varphi]\varphi$ which is equivalent to $[\oplus\varphi]\varphi$; therefore preserved formulas are successful formulas. The inductive case $[\oplus\neg\varphi]\psi$ in the 'preserved formulas' may possibly puzzle the reader. Its proof [vDK05] is quite elementary (and proceeds by induction on formula structure) and shows that the puzzling *negation* in the announcement clause is directly related to the truth of the announcement as a *condition*:

Let $M, s \models [\oplus\neg\varphi]\psi$, and $M' \subseteq M$ such that $s \in M'$. Assume $M', s \models \neg\varphi$. Then $M, s \models \neg\varphi$ by *contraposition* of the inductive hypothesis for φ. From that and $M, s \models [\oplus\neg\varphi]\psi$ follows $M|\neg\varphi, s \models \psi$. From the inductive hypothesis for ψ follows $M'|\neg\varphi, s \models \psi$. Therefore $M', s \models [\oplus\neg\varphi]\psi$ by definition.

Dynamic doxastic logic as belief revision We now present a (different from DDL) dynamic doxastic semantics that can be seen as the implementation of a belief revision operator. Assuming that the new information state (pointed Kripke model) is constructed from the current information state and the revision

that asserting $\varphi \wedge \neg B\varphi$ implies $B\varphi$, which contradicts $\neg B\varphi$ [Moo42, pp.204–205]. The further development of this notion, addressed in our contribution, firstly puts Moore-sentences in a *multi-agent* perspective of announcements of the form 'p is true and *you* don't believe that', and, secondly, puts Moore-sentences in a *dynamic* perspective of announcements that cannot be believed after being announced. Both perspectives appear to go beyond Moore.

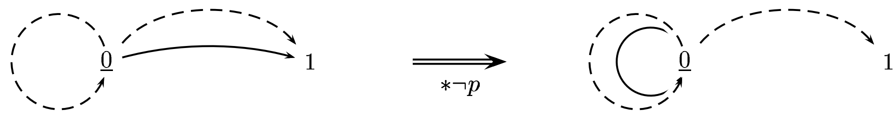

Figure 2: Belief revision related to degrees of belief

formula (and does not assume an underlying transition structure), it seems harder to provide a mechanism that explains *adding* or *changing* access than one for merely deleting access—for the same reason that contraction or revision needs an entrenchment relation or something similar. A transition as in Figure 1 is hard to justify by the structure of the belief state before the revision. And another problem is that the dynamic epistemic logics presented so far do not provide a way to model 'forgetting' knowledge or beliefs, as, from an agent's point of view, belief and knowledge are indistinguishable in these logics (and 'belief' always means 'conviction').

In a setting where *degrees of belief* and possibly knowledge too are represented, one can provide such a structural justification. For a simple example, we add another degree of belief to the 'revision' example in Figure 1: in Figure 2 the dotted line interprets a stronger degree of belief. It contains (entails) the weaker, or most normal, belief that is represented as before with the solid line. There is a one-to-one correspondence between such accessibility relations satisfying inclusion, and preferences between worlds or a 'system of spheres' as in [Lew73] and also propagated in [Gro88, Spo88].

The basic idea in [Lew73] is that, given a domain of worlds, from the perspective of a given world in that domain, some worlds may be preferable over others. The worlds for which a preference exists are the plausible worlds, and the preference is typically a partial order (plus additional constraints). In the left model in Figure 2, we have that, given world 0, the set of plausible worlds is $\{0, 1\}$, and that world 1 is preferred over world 0, for which we write $1 <^0 0$; relation $<^0$ is the preference relation associated with world 0. The set of plausible worlds given world 1 is empty; $<^1 = \emptyset$. The belief revision resulting in the model on the right now consists of changing preferences between the plausible worlds: in the resulting model world 0 is preferred over world 1: $0 <^0 1$. The relation between preferences and accessibility is fairly simple. In general, given a partial order with degrees of belief x we can define $s \rightarrow^x s'$ if and only if the degree of world s' in the preference order $<^s$ is at most x. In Figure 2 we have two degrees of belief, and therefore two accessibility relations; the 'at most' is to ensure an inclusion relation. To accessibility relations $\rightarrow^x$ are associated 'degree of belief' modalities B^x in the obvious way.

Such a modal setting for reasoning about preferences also applies to a multi-agent situation, one can also restrict oneself to introspective belief or knowledge, and further demand additional (frame characterizable) restrictions expressing that agents are knowledgable about their own preferences. Concerning the static picture, such ideas have emerged under the name of 'dynamic in-

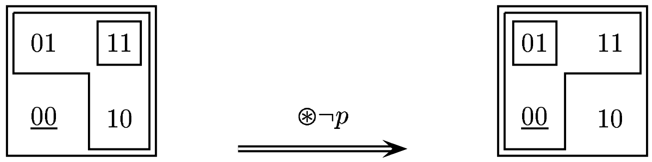

Figure 3: This depicts Spohn-like belief revision in a multimodal setting.

teractive epistemology' in the game theoretical and philosophical community [Sta96, AS05, Boa04]—the word 'dynamic' refers to the conditional modal operators in those approaches that are used to model belief revision, not to the dynamic modal approach intended here.[3] The ideas have also surfaced as dynamic doxastic logic(s) in [vDL03, Auc03, vD05]. We close this paper with an example of the latter.

Example of introspective belief revision A proposal by Aucher [Auc03] (or its possibly easier to procure version [Auc05]) can be seen as an implementation of one of Spohn's proposals in [Spo88][4]. An illustration is depicted in Figure 3. On the left in that figure, the agent believes atomic propositions p and q – the name 11 stands for the world where p and q are both true. In particular, Bp is true. Note that in this case there are three degrees of belief, let us say 0, 1, and 2. Degree of belief 0 stands for most normal belief and $\rightarrow^0$ therefore corresponds to B. Apart from that, we have B^1 and B^2—in this case B^2 is equal to knowledge K. For example, $B^1(p \vee q)$ is valid on the model: the agent has a somewhat stronger belief in the (weaker proposition than $p \wedge q$ namely that) $p \vee q$.

On the right the agent believes $\neg p$ (and q). In other words, $B\neg p$ is true. On the left, it is therefore true that after revision with $\neg p$, the agent believes $\neg p$: $[\circledast\neg p]B\neg p$ is true. The belief revision is therefore successful. This revision is computationally achieved by the following recipe: determine the minimal world where the revision formula $\neg p$ is true. This is 01. Now deduct the degree of *that* world from the degree of *any* $\neg p$-world. We thus get a degree 0 for world 01, and a degree 1 for world 00. For worlds where the revision formula is false,

[3]The relation between conditional modal operators and dynamic modal operators, and how appropriate it is to model belief revision in the former setting compared to the latter, seems to us only incompletely understood and merits further investigation. See [vD05], appendix A, and see [vBvEK05] for the encompassing notion of 'relativized common knowledge'—a proposal to generalize conditional (individual) knowledge.

[4]Namely the revision also known as 'minimal Spohn': when revising with φ, make the minimal φ-worlds the 'most normal' worlds, such that they are believed after the revision; for details, see Definition 6 on page 117 in [Spo88]. In Spohn's terms the revision in the example below would be called $\{00, 01\}$-1 conditionalization of the current ordinal conditional function, where $\{00, 01\}$ is the denotation of the revision formula $\neg p$ in the current epistemic state, and '1' is the decreased 'firmness α' with which the p worlds are updated.

i.e., where p is true, do the same, but ensure that the most normal p-world is slightly less preferable than the most normal $\neg p$. In this particular example we ensure this by 'adding 1 to the current degree'. This results in a degree 1 for world 11 and a degree 2 for world 10. This completes the computation.

The above presents only *one* example of *one* dynamic belief revision operator that can be seen as an implementation of the AGM postulates. This particular operator is successful on propositional formulas, and for those can also be considered as effecting minimal change. Other examples are given in [vD05], the list of over twenty different theory change operators in [Rot04] seems also particularly suitable for implementation in this setting.

A final word on 'update'. As mentioned, belief update as opposed to belief revision is also investigated in dynamic epistemics under the name of 'factual change'. This is investigated in, for example, [BMS99, vE04, vBvEK05, vDvdHK05]. These ideas also deserve to be properly applied to the belief revision arena.

References

[AGM85] 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.

[AS05] G.B. Asheim and Y. Søvik. Preference-based belief operators. *Mathematical Social Sciences*, 50(1):61–82, 2005.

[Auc03] G. Aucher. A combined system for update logic and belief revision. Master's thesis, ILLC, University of Amsterdam, Amsterdam, the Netherlands, 2003.

[Auc05] G. Aucher. A combined system for update logic and belief revision. In M.W. Barley and N. Kasabov, editors, *Intelligent Agents and Multi-Agent Systems – 7th Pacific Rim International Workshop on Multi-Agents (PRIMA 2004)*, pages 1–17. Springer, 2005. LNAI 3371.

[BM04] A. Baltag and L.S. Moss. Logics for epistemic programs. *Synthese*, 139:165–224, 2004. Knowledge, Rationality & Action 1–60.

[BMS98] A. Baltag, L.S. Moss, and S. Solecki. The logic of common knowledge, public announcements, and private suspicions. In I. Gilboa, editor, *Proceedings of the 7th conference on theoretical aspects of rationality and knowledge (TARK 98)*, pages 43–56, 1998.

[BMS99] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. Technical report, Centrum voor Wiskunde en Informatica, Amsterdam, 1999. CWI Report SEN-R9922.

[Boa04] O. Board. Dynamic interactive epistemology. *Games and Economic Behaviour*, 49:49–80, 2004.

[dR94] M. de Rijke. Meeting some neighbours. In J. van Eijck and A. Visser, editors, *Logic and information flow*, pages 170–195, Cambridge MA, 1994. MIT Press.

[FHMV95] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge MA, 1995.

[Gär88] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge, MA, 1988.

[Ger99] J.D. Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, University of Amsterdam, 1999. ILLC Dissertation Series DS-1999-01.

[Ger05] J.D. Gerbrandy. The surprise examination in dynamic epistemic logic. *Synthese*, 2005. To appear.

[Gro88] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

[Hin62] J. Hintikka. *Knowledge and Belief*. Cornell University Press, Ithaca, NY, 1962.

[KM91] H. Katsuno and A. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, pages 387–394, 1991.

[Lew73] D. Lewis. *Counterfactuals*. Harvard University Press, Cambridge MA, 1973.

[LR99a] S. Lindström and W. Rabinowicz. Belief change for introspective agents, 1999. `http://www.lucs.lu.se/spinning/`.

[LR99b] S. Lindström and W. Rabinowicz. DDL unlimited: dynamic doxastic logic for introspective agents. *Erkenntnis*, 50:353–385, 1999.

[Moo12] G.E. Moore. *Ethics*. Oxford University Press, 1912. Consulted edition: The Home University Library of Modern Knowledge, volume 54, OUP, 1947.

[Moo42] G.E. Moore. A reply to my critics. In P.A. Schilpp, editor, *The Philosophy of G.E. Moore*, pages 535–677. Northwestern University, Evanston IL, 1942. The Library of Living Philosophers (volume 4).

[Moo44] G.E. Moore. Russell's "theory of descriptions". In P.A. Schilpp, editor, *The Philosophy of Bertrand Russell*, pages 175–225. Northwestern University, Evanston IL, 1944. The Library of Living Philosophers (volume 5).

[Pla89] 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.

[Rot04] H. Rott. Adjusting priorities: Simple representations for 27 iterated theory change operators. Manuscript, 2004.

[Seg99a] K. Segerberg. Default logic as dynamic doxastic logic. *Erkenntnis*, 50:333–352, 1999.

[Seg99b] 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, Dordrecht, 1999. Kluwer Academic Publishers.

[Spo88] W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, volume II, pages 105–134, 1988.

[Sta96] R. Stalnaker. Knowledge, belief and counterfactual reasoning in games. *Economics and Philosophy*, 12:133–163, 1996.

[vB94] J.F.A.K. van Benthem. Logic and the flow of information. In *Proceedings of the 9th International Congress of Logic, Methodology and Philosophy of Science (1991)*. Elsevier Science B.V., 1994.

[vB02] J.F.A.K. van Benthem. One is a lonely number: on the logic of communication. Technical report, ILLC, University of Amsterdam, 2002. Report PP-2002-27 (material presented at the Logic Colloquium 2002).

[vBvEK05] J.F.A.K. van Benthem, J. van Eijck, and B.P. Kooi. Logics of communication and change. manuscript, 2005.

[vD05] H.P. van Ditmarsch. Prolegomena to dynamic logic for belief revision. *Knowledge, Rationality & Action*, 2005. To appear.

[vdHV02] W. van der Hoek and L.C. Verbrugge. Epistemic logic: a survey. In L.A. Petrosjan and V.V. Mazalov, editors, *Game theory and applications*, volume 8, pages 53–94, 2002.

[vDK05] H.P. van Ditmarsch and B.P. Kooi. The secret of my success. *Synthese*, 2005. To appear.

[vDL03] H.P. van Ditmarsch and W.A. Labuschagne. A multimodal language for revising defeasible beliefs. In E. Álvarez, R. Bosch, and L. Villamil, editors, *Proceedings of the 12th International Congress of Logic, Methodology, and Philosophy of Science (LMPS)*, pages 140–141. Oviedo University Press, 2003.

[vDvdHK05] H.P. van Ditmarsch, W. van der Hoek, and B.P. Kooi. Dynamic epistemic logic with assignment. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 05)*, volume 1, pages 141–148, New York, 2005. ACM Inc.

[vDvdHK06] H.P. van Ditmarsch, W. van der Hoek, and B.P. Kooi. Dynamic epistemic logic. Graduate textbook, scheduled to appear, 2006.

[vE04] J. van Eijck. Guarded actions. Technical report, Centrum voor Wiskunde en Informatica, Amsterdam, 2004. CWI Report SEN-E0425.

[vvK05] H. P. van Ditmarsch, W. van der Hoek, and B. P. Kooi. Public announcements and belief expansion. In R. Schmidt, I. Pratt-Hartmann, M. Reynolds, and H. Wansing, editors, *Advances in Modal Logic*, volume 5, pages 335–346. King's College Publications, London, 2005.