Jan-Willem Romeijn

Conditioning and Interpretation Shifts

Abstract. This paper develops a probabilistic model of belief change under interpretation shifts, in the context of a problem case from dynamic epistemic logic. Van Benthem [4] has shown that a particular kind of belief change, typical for dynamic epistemic logic, cannot be modelled by standard Bayesian conditioning. I argue that the problems described by van Benthem come about because the belief change alters the semantics in which the change is supposed to be modelled: the new information induces a shift in the interpretation of the sentences. In this paper I show that interpretation shifts can be modeled in terms of updating by conditioning. The model derives from the knowledge structures developed by Fagin *et al* [8], and hinges on a distinction between the propositional and informational content of sentences. Finally, I show that Dempster-Shafer theory provides the appropriate probability kinematics for the model.

Keywords: Probabilistic epistemology, Bayesian updating, Dynamic epistemic logic, Knowledge structures, Dempster-Shafer theory.

1. Introduction

In the Bayesian model, beliefs over sentences from a language $\mathcal{L}$ are represented with probability assignments over an algebra generated by possible worlds, $w = \{w_1, w_2, \ldots, w_n\}$. The sets of worlds, or propositions, are associated with the sentences by an interpretation function, $I : \mathcal{L} \times w \to \{0, 1\}$. A sentence s is associated with a proposition $[\![s]\!]$ as follows:

$$[\![s]\!] = \{w \in w : \ I(w, s) = 1\}.$$

The probability measures over sets of possible worlds $[\![r]\!]$ and $[\![s]\!]$ are understood as degrees of belief in the sentences r and s. A belief change due to the acceptance of a sentence is reflected in the probability assignment by updating according to Bayes' rule, that is, by conditioning the assignment on the proposition associated with the sentence:

$$P_s([\![r]\!]) = P([\![r]\!] \mid [\![s]\!]).$$

Presented by **Hannes Leitgeb**; *Received* July 24, 2009

Studia Logica (2012) 100: 583–606
DOI: 10.1007/s11225-012-9415-x

Note that the expression on the right can be rewritten by means of Bayes' theorem. The theorem is distinct from Bayes' rule, which is expressed here. Bayes' rule is an epistemological principle linking the new probability of $[\![r]\!]$ after learning s to the old probability of $[\![r]\!]$ conditional on $[\![s]\!]$.

The idea of conditioning is that within the set of possible worlds consistent with the sentence that is learnt, the probability is kept unchanged. This conservativity of belief change is usually referred to as *rigidity* (cf. Jeffrey [15]). It seems rather natural to impose this condition if we interpret sentences in terms of propositions: after learning the sentence s, we do not consider possible those worlds in which $\neg s$ holds, i.e., we zoom in on the worlds within the set $[\![s]\!]$. But apart from that there is no reason to nudge up the probability of one world in which s holds, at the cost of the probability of another such world. These intuitions, pleasingly illustrated by van Fraassen's [9] muddy Venn diagrams, can be given a further underpinning by the so-called dynamic Dutch book arguments.

Since we use propositions, or sets of possible worlds, to interpret sentences, a shift in the interpretation of a sentence can be represented by a change in the function that determines the set of possible worlds associated with a sentence. Now imagine that the interpretation of a sentence is partly determined by the set of sentences that have been accepted thus far. Then it may happen that the acceptance of s causes the sentence r to obtain a different interpretation. We move from I to I_s: before the acceptance of s the sentence r is associated with one set of possible worlds, $[\![r]\!]$, but after the acceptance of s it is associated with a different set of possible worlds, $[\![r]\!]_s$ say.

Shifts in interpretation of the above kind cause problems for the Bayesian model. Imagine that all possible worlds associated with the sentence s are worlds that satisfy the sentence r, so that $[\![s]\!] \subseteq [\![r]\!]$. By Bayes' rule, accepting s entails that we assign r a probability $P_s([\![r]\!]) = P([\![r]\!] \,|\, [\![s]\!]) = 1$. However, accepting s may change the interpretation of r. After the change in interpretation, some worlds that were previously included in both $[\![s]\!]$ and $[\![r]\!]$, and that had nonzero probability before the update, may fail to be included in $[\![r]\!]_s$. In that case the probability for $[\![r]\!]$ must be smaller than 1 after a complete update with s. In other words, interpretation shifts lead to a conflict between conditional and updated probability.

Various conceptual problems in formal approaches to epistemology can be viewed as examples of such interpretation shifts. One example is the violation of the principle of reflection, as discussed in van Fraassen [9] and Maher [18]. Let s be that I drink a bottle of whiskey tonight, and let r be that I am fit to drive home early tomorrow. Of course, pondering over these

sentences in the afternoon, I will assign a probability close to zero to the second conditional on the first. But after having drunk the bottle of whiskey, it may be that I assign a larger probability to being fit to drive home in the morning than this conditional probability. One way to understand this apparent inconsistency is by saying that the interpretation of the phrase "being fit to drive in the morning" has changed. Worlds in which I risk several lives including my own by driving home, previously excluded from the proposition associated with "I am fit to drive home", are now included in this proposition.

Further examples stem from the literature on belief revision. Rott [20] and Arlo-Costa and Pedersen [1] discuss cases in which a department needs to choose the most suitable candidate for a given position. Of course, if we learn that the hiring committee has narrowed down the choice to a shortlist that excludes the candidate first thought to be most suitable, then this will invite us to revise our beliefs about which candidate will win. But learning about the shortlist may also give cause to revise our ideas about the criteria the committee is using. As a result, we might also revise our beliefs on who will win if the candidate first thought to be most suitable is still on the shortlist, simply because the shortlist contains some surprising choices and thus reveals new information about the selection criteria. Such belief changes can be understood as changes in the interpretation of "most suitable candidate".

In what follows I will enter a different domain of application. Interpretation shifts of the above kind also occur in contexts where people reason about knowledge, and perform so-called epistemic actions: the domain of dynamic epistemic logic. Apart from this being a convincing domain of application, there is an independent reason to investigate interpretation shifts here. Probabilistic epistemology and dynamic epistemic logic are both concerned with representing the dynamics of belief but to date there have been few attempts at a rapprochement. From the side of dynamic epistemic logic there are attempts to accommodate probabilistic updating (e.g. Kooi [16], Baltag and Smets [3], van Benthem, Gerbrandy and Kooi [5]). But within probabilistic epistemology there has been scarce attention for the peculiarities of dynamic and epistemic settings. The present paper is aimed at filling this lacuna. Admittedly it targets one specific aspect of dynamic epistemic logic, and only from the angle of Bayesian epistemology. Moreover, it deals primarily with an example. A fully general characterisation of belief change for interpretation shifts is not given. Nevertheless I hope that this paper makes the modest beginnings of a fruitful exchange between these two research fields.

2. An example from dynamic epistemic logic

Van Benthem [4] provides an example showing that some instances of belief change cannot be modelled by Bayesian conditioning. In this section I rehearse the problem case of van Benthem to show that it exemplifies the kind of interpretation shift at issue in this paper. The example is set against the background of Kripke models for dynamic epistemic logic.[1]

Imagine Alice and Bob, who are both investigating the state of the world. The Kripke model of Figure 1 summarises the possible worlds that they consider, and also expresses their epistemic perspectives on these worlds. There are three dots which represent worlds: world 1 in which $s \wedge \neg r$, world 2 in which $s \wedge r$, and world 3 in which $\neg s \wedge r$. The double arrows indicate the so-called epistemic accessibility relations between worlds: from each world there are arrows expressing what other worlds they consider possible. These arrows are labelled for Alice and Bob separately: they both have their own set of relations, meaning that they each have their own epistemic perspective. Notably, at each world Alice and Bob both have epistemic access to this world itself. The corresponding reflexive arrows are omitted for the sake of simplicity.

Let us consider the epistemic perspectives of Alice and Bob in some more detail. If world 1 is actual, Bob knows that $s \wedge \neg r$ because he only considers it possible that he is in world 1: there is no arrow for Bob that connects world 1 to any other world. On the other hand, if world 1 is actual Alice knows that s but she does not know whether r, because following her arrows she also considers world 2 possible, and while $\neg r$ holds in world 1, r holds in world 2. This also means that at world 1 Alice does not know that Bob knows that $s \wedge \neg r$, because she considers world 2 possible and if this is the actual world, then Bob does not know $s \wedge \neg r$. More precisely, if world 2 is actual, the situation looks much the same for Alice. She considers world 1 possible as well, and so she does not know whether r holds. Bob, on the other hand, knows that r but he does not know whether s, because at world 2 he also considers world 3 possible. Accordingly, if world 2 is actual Alice and Bob both do not know of each other's lack of knowledge. In world 3, finally, Alice knows that $\neg s \wedge r$ since she only considers world 3 possible, while for Bob everything looks much the same as in world 2.

Note that sentences like r and s are all associated with possible worlds, given by the dots, and not with the epistemic relations as summarised by the double arrows. Sentences about knowledge, on the other hand, are concerned

[1] See van Ditmarsch, van der Hoek, and Kooi [7] for a general introduction.

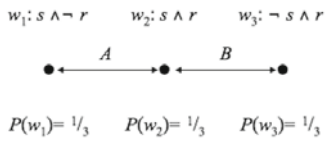

Figure 1. A Kripke diagram depicting the worlds and epistemic accessibility relations for Alice and Bob. Reflexive arrows are omitted.

also with the arrows. Note further that, since epistemic accessibility relations can be concatenated, the Kripke diagram above summarises all the higher-order beliefs that Alice and Bob have concerning each other's beliefs, their beliefs about each other's beliefs, and so on. Finally, notice that in the model introspective realisations are instantaneous. If world 2 is actual, then Bob does not know whether s but he also knows effortlessly that he does not know this, that he knows he does not know, and so on. As will be seen, the model of the present paper diverges from this: such introspections are understood as a separate stage in the process of belief change.

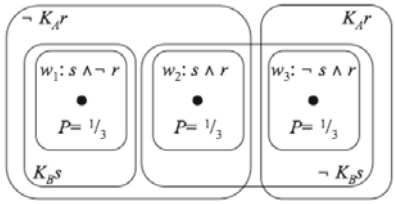

Figure 2. A diagram showing the propositions associated with the sentences $K_A s$, $K_B r$ and their negations.

As indicted, the worlds have a prior probability $P(w_i) = 1/3$. This prior expresses that we, the modelers, are uncertain about which world is actual, and hence that we are uncertain about the epistemic perspectives that Alice and Bob have. Consider the sentences "Alice does not know that r", written as $\neg K_A r$, and "Bob does not know that s", denoted $\neg K_B s$. As depicted in Figure 2, the interpretation of the sentence $\neg K_A r$ is the proposition or, equivalently, the set of worlds $[\![\neg K_A r]\!] = \{w_1, w_2\}$, while $[\![\neg K_B s]\!] = \{w_2, w_3\}$. Now imagine that it is publically announced that Alice does not know r, $\neg K_A r$, and that Alice and Bob revise their epistemic perspective accordingly. With this information we can conclude that world 3 cannot be actual, because at that world Alice knows that r. According to a naive Bayesian model, and as suggested by Figure 3, the new

probability assignment $P_{\neg K_A r}$ for the proposition $[\![\neg K_B s]\!]$ must therefore be $P_{\neg K_A r}([\![\neg K_B s]\!]) = P([\![\neg K_B s]\!] \,|\, [\![\neg K_A r]\!]) = P(\{w_2, w_3\}|\{w_1, w_2\}) = 1/2$.

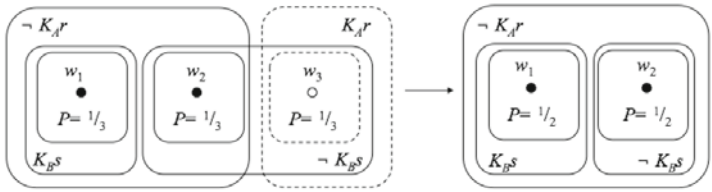

Figure 3. A Kripke diagram depicting the situation for Alice and Bob after naive conditioning on $\neg K_B r$.

But this is not the whole story about the update. In the account provided by dynamic epistemic logic, the belief change rightly involves a revision of the epistemic perspectives of Alice and Bob, as expressed by the Kripke model. Apart from ruling out w_3 a complete update requires that Alice and Bob operate on the epistemic relations between the remaining worlds. Specifically, because w_3 is ruled out, neither of them will include epistemic relations with w_3 after the update. So at world 2 Bob cannot access world 3 anymore. The proper representation of the situation after the update is therefore given on the right side of Figure 4. Importantly, in the new epistemic situation Bob is not in doubt about the sentence s in any of the two remaining worlds. After the complete update, there are no possible worlds associated with the sentence $\neg K_B s$. The new probability must therefore be $P_{\neg K_A r}([\![\neg K_B s]\!]) = P(\emptyset|\{w_2, w_3\}) = 0$, and not the half derived earlier.

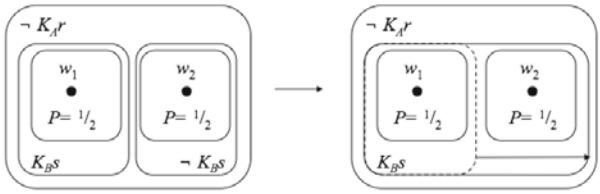

Figure 4. A Kripke diagram depicting the situation for Alice and Bob after a full update on $\neg K_A s$. The diagram on the right suggests how the full update induces a shift in interpretation.

This shows that simply conditioning on the proposition $\{w_1, w_2\}$ leads to the wrong probability assignment for the sentence $\neg K_B s$. After learning $\neg K_A r$ we must re-evaluate of each world whether or not it is a member of the proposition $[\![\neg K_B s]\!]$, and adapt the probabilities accordingly. As it turns out, the proposition changes from $[\![\neg K_B s]\!] = \{w_2, w_3\}$ to $[\![\neg K_B s]\!]_{\neg K_A r} = \emptyset$.

In short, due to the acceptance of $\neg K_A r$, the interpretation of $\neg K_B s$ has shifted.

3. Epistemic update as conditioning

Van Benthem is right to identify this problem in the Bayesian model of belief change. The above case of belief change, here portrayed in terms of a shift in interpretation, creates a conflict between conditional and updated probability. One may therefore conclude, with van Benthem, that Bayesian updating has its limitations, and subsequently turn to a different formal model of belief change. But the reaction of this paper is less dismissive of the Bayesian model. It is to look for a characterisation of possible worlds and a generalisation of Bayesian updating that together allow us to incorporate interpretation shifts. In this section I make the beginnings of such a characterisation by representing the example in such a way that the belief change is a conditioning operation after all.

The guiding idea is that all aspects of the information that play a part in the learning event must somehow be made explicit in the possible worlds semantics.[2] Dynamic epistemic logic employs what may be termed a thin notion of possible world, characterised only by sentences like r and s being true or false. The epistemic structure is superimposed on the worlds by means of accessibility relations between worlds, and the update rules operate on the worlds as well as on this epistemic structure. The following, by contrast, employs a thick notion of possible worlds, characterised also by epistemic sentences such as $\neg K_A r$ and $\neg K_B s$. This complicates the notion of possible world somewhat, but as we will see the update rule can be kept conveniently simple. The upshot is a better fit with the Bayesian approach.

The framework required here is that of so-called knowledge structures: each possible world is separately furnished with an internal epistemic structure that expresses which other worlds the agents consider possible at that world. Fagin, Halpern and Vardi [8] give a detailed treatment of this framework, and prove that knowledge structures are fully isomorphic to Kripke models. But there are nevertheless good reasons to employ knowledge structures. Most importantly, knowledge structures seem better suited for a treatment of the belief change along Bayesian lines. A Bayesian update is the imposition of a constraint on the set of possibilities, and as such it is blind

[2]This is inspired by Halpern [14] who, among other things, discusses extensions of the semantics by a probabilistic representation of the protocol used by the information source. However, this paper will ignore the question what protocol the information source is using.

to any structure other than set inclusion. But a knowledge structure is cast entirely in terms of sets of possibilities. Moreover, as will become apparent, knowledge structures allow for a more explicit account of acts of introspection, during which the agent works out the consequences of the information they received.[3]

The present section only gives an illustration of knowledge structures in the context of van Benthem's example.[4] Knowledge structures unravel the epistemic accessibility relations between worlds that we know from Kripke structures, and capture these relations in terms of possibilities associated with worlds. Fortunately we need not unpack the Kripke structure of Figure 1 very far to arrive at a semantics that can accommodate the learning events described in the foregoing. Instead of distinguishing only the possible worlds 1, 2 and 3, we now make a further distinction between the possible epistemic states, or states of mind for Alice and Bob at each of the possible worlds. At world 1, for instance, Bob can only conceive of world 1, while Alice considers world 1 and world 2 possible. Accordingly, we distinguish two epistemic states that both belong to world 1, one in which Alice and Bob both think world 1 is actual, and one in which Bob thinks world 1 is actual while Alice thinks world 2 is. Similarly, we can distinguish four epistemic states at world 2: Alice considers world 1 and 2 possible, and Bob considers world 2 and 3 possible. World 3 has two epistemic states again, corresponding to Bob considering world 2 and 3 possible. In analogy to Figure 1, Figure 5 summarises the new set of possible worlds and epistemic states.

We can now model the inclusion of the information expressed by the sentence that Alice does not know r, $\neg K_A r$, by means of a conditioning operation. We say that world i is one in which Alice does not know r if at this world we find epistemic states in which Alice considers r to be false. Updating with the sentence $\neg K_A r$ means, first of all, that we zoom in on those worlds in which Alice does not know r. In this case, as illustrated by Figure 6, we eliminate world 3 from the possible world semantics. This first stage of the update corresponds to the simple Bayesian update discussed in the foregoing: after the update we have only world 2 left in which Bob does not know s, so that presumably $P_{\neg K_A r}([\![\neg K_B s]\!]) = 1/2$. This update is equivalent to the update depicted in Figure 3.

---

[3]The two-step update procedure detailed below echoes the distinction between straightforward and conscious updates as discussed in Groeneveld [12], especially sections 4.4 to 4.6, and Gerbrandy [11], in particular sections 4.2 to 4.4.

[4]More detail on knowledge structures is provided in the next section.

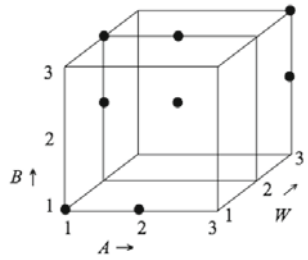

Figure 5. A possible worlds semantics depicting the epistemic states for Alice and Bob before the update. Along axis W we find the three worlds, and within each world we have an axis A coding for epistemic states of Alice, and an axis B coding for the states of Bob. For example, the foremost square represents world 1, at which Bob only considers world 1 possible while Alice considers world 1 and world 2 possible, so that the epistemic stats at world 1 have the coordinates $(1, 1)$ and $(2, 1)$ for Alice and Bob respectively.

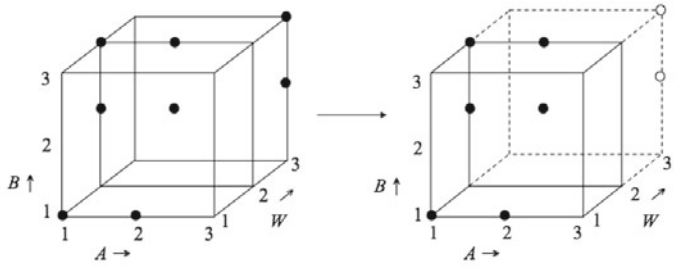

Figure 6. The possible worlds semantics depicting the epistemic states for Alice and Bob after an update with the information that world 3 is not actual, but without anyone having thought this information through.

In the foregoing it was shown that the update operation of eliminating world 3 is not the full story, because this elimination interfered with the epistemic accessibility relations, leading to an interpretation shift. This second stage of the update operation, in which Alice and Bob work out the epistemic consequences of what they learnt in the first stage, can also be modeled as a conditioning operation.[5] In the example, Alice and Bob do not just learn that world 3 is not actual, they also learn that they should not consider world 3 possible. Because of this we exclude the whole of world 3, but we also exclude those epistemic states at other worlds in which Alice and Bob consider world 3 possible.

The update that corresponds to this stage is depicted in Figure 7, which is analogous to the update of Figure 4. In dynamic epistemic logic this part of

[5]Cf. Groeneveld [12], pp. 161–162, and Gerbrandy [11], pp. 87–88.

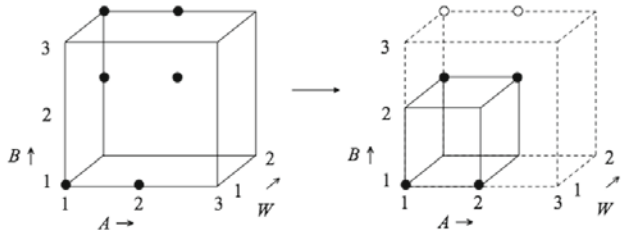

Figure 7. The possible worlds semantics depicting the epistemic states for Alice and Bob after updating with all the information included in $\neg K_A r$.

the update is sometimes called the conscious part. Crucially, in the semantics of Figure 5, this second stage of the update is a conditioning operation much like the first part: we eliminate the two epistemic states in world 2 that pertain to world 3. In the resulting possible worlds semantics there are no worlds left in which Bob does not know s, because there are no worlds left at which Bob considers a world for which $\neg s$ holds to be possible. So we have that $[\![\neg K_B s]\!]_{\neg K_A r} = \emptyset$ and hence that $P_{\neg K_A r}([\![\neg K_B s]\!]) = P_{\neg K_A r}(\emptyset) = 0$. The public announcement of $\neg K_A r$ effects a shift in the interpretation of $\neg K_B s$.

Summing up, we have sketched a representation of the belief change that occurs in the epistemic example. According to this representation, the belief change involves a shift in interpretation, but it can nevertheless be modeled in terms of conditioning. We conditioned on a particular set of epistemic states, namely all states that somehow pertain to world 3. Or in terms of the cube of Figure 5, we have removed the layer labeled 3 in all three dimensions. We might say that the set of states conditioned upon is the information contained in the sentence $\neg K_A r$, and that this information is richer than just the proposition associated with the sentence. In what follows, I will show what we have gained by this representation of the belief change. The key is that the belief change, including the interpretation shift, consists in the elimination of states, so that the change is amenable to a more or less Bayesian treatment. But before we can fill in the details of the interpretation shift and the corresponding probability kinematics, we need to provide some more detail on knowledge structures, and on how they are related to the present model.

4. Knowledge structures

In this section we provide a framework for the epistemic states and updates introduced in the foregoing. The eventual goal is to come up with a

probability kinematics for the interpretation shifts exemplified in Section 2. But we can only define the probability kinematics properly once we have a clearly defined structure on which the probability assignments rest. As indicated, the knowledge structures of Fagin *et al* [8] fulfill this role. I introduce the general idea of these structures below, and then focus on the aspects of knowledge structures that are relevant to our present concerns.

The general idea of Fagin *et al* is that possible worlds can be associated with further structure that expresses the epistemic perspective, or the set of possible epistemic states, of the agents at that world. The complete specification of both the world and the perspectives of the agents is termed a *knowledge structure*. It can be built up inductively, as follows.

Definition 1 (Knowledge structure). Let W be a set of worlds. At each world w, define the 0-th order knowledge assignment $f_0(X)$ for the agent X as the world w_i, and call $\langle f_0 \rangle$ the 1-world. Now assume inductively that k-worlds have been defined. The k-th order knowledge assignment $f_k(X)$ for agent X is the subset of the k-worlds that agent X considers possible. Then a $(k+1)$-world is a sequence of knowledge assignments $\langle f_0, f_1, \ldots f_k \rangle$ if it satisfies the following criteria:

Correctness: The knowledge assignment of order $k-1$ is among the possibilities in the k-th order knowledge assignment for each agent, so for all X we have $\langle f_0, \ldots f_{k-1} \rangle \in f_k(X)$.

Introspection: Every possibility in the k-th order knowledge assignment of any agent X corresponds with the knowledge assignment at order $k-1$, so for all X we have that if $\langle g_0, \ldots g_{k-1} \rangle \in f_k(X)$ then $g_{k-1}(i) = f_{k-1}(X)$.

Extension: Every possibility in the $(k-1)$-th order knowledge assignment for any agent X has an extension to a k-th order knowledge assignment, so for all X we have that if $\langle g_0, \ldots g_{k-2} \rangle \in f_{k-1}(X)$ then $\exists g_{k-1} : \langle g_0, \ldots g_{k-1} \rangle \in f_k(X)$. Similarly, every possibility in the k-th order knowledge assignment of X is an extension of some $(k-1)$-th order knowledge assignment.[6]

A knowledge structure f is an infinite sequence $\langle f_0, f_1, \ldots \rangle$.

At each world, a knowledge structure describes the *epistemic perspective* of the agents up to arbitrary order. The structures provide a representation that is equivalent with Kripke models, but they are also different in important ways. Whereas a Kripke models covers all orders of knowledge in

[6] This direction of the equivalence is already entailed by the combination of Correctness and Introspection.

a single representation, knowledge structures give us a separate handle on the different orders in the epistemic perspectives of the agents. As a result, we can closely follow how information that an agent receives percolates up in the hierarchy of knowledge, and how this causes the agent to adapt her beliefs at ever higher orders. For this reason knowledge structures can be tailored to represent epistemic states of finite order. Rather than employing the complete structure, we may use the first k elements in the structure to represent knowledge up to order k.

For the purpose of modeling interpretation shifts in the epistemic example, we only need a very shallow knowledge structure. It suffices to specify the knowledge of Alice and Bob up to order 1, on the basis of the given Kripke model. That is, we need only determine the epistemic perspectives of Alice and Bob on the worlds, and we can ignore any further introspective and intersubjective knowledge. The epistemic perspectives of agent X can be characterised as follows:

$$\langle f_0(X), f_1(X)\rangle = \langle w_i, \{w_{i'} : \langle w_i, w_{i'}\rangle \in R_X\}\rangle\,.$$

Agents Alice and Bob are here indexed with X, and their epistemic accessibility relations are R_A and R_B respectively. So at world 1 Alice's epistemic perspective is $\langle w_1, \{w_1, w_2\}\rangle$ while Bob's perspective is $\langle w_1, \{w_1\}\rangle$. At world 2, Alice has $\langle w_2, \{w_1, w_2\}\rangle$ while Bob has $\langle w_2, \{w_2, w_3\}\rangle$ and at world 3 Alice and Bob have $\langle w_3, \{w_3\}\rangle$ and $\langle w_3, \{w_2, w_3\}\rangle$ respectively.

The most salient important difference between Kripke models and knowledge structures is that in the latter, the epistemic perspective is located at a particular world. In what follows, I will spell out worlds and epistemic perspectives as sets of possible *epistemic states* or states of mind, to contrast them with states of affairs. Consider an epistemic perspective at world w_i of order k, described by the prefix of a knowledge structure of length $k+1$, namely $\langle f_0^i, \ldots, f_k^i\rangle$. For every world i and agent X, and for every $j \leq k$, every j-th order knowledge assignment $f_j^i(X)$ is some set of j-worlds. We define the set of epistemic states of order k at a particular world w_i as

$$w_i \times \prod_X f_1^i(X) \times \ldots \times f_k^i(X). \tag{1}$$

An epistemic state of order k is a complete description of a particular possibility, regarding the world and regarding the knowledge of both agents up to order k. Notice that the knowledge assignments f_j^i carry a superscript i to denote the world with which they are associated. The full set of epistemic states at world i determines the prefix of length k of the knowledge structure, and hence the epistemic perspective.

Let me illustrate the set of epistemic states for the example of van Benthem. The perspective in the example is of order 1, so at world w_i the epistemic states are simply the elements generated by a Cartesian product of $f_1^i(A)$ and $f_1^i(B)$. Following Equation (1) we find:

$$\begin{aligned} W_1 &= \{111, 121\}, \\ W_2 &= \{212, 213, 222, 223\}, \\ W_3 &= \{332, 333\}, \end{aligned} \tag{2}$$

where each 3-tuple iab is shorthand for $\langle w_i, w_a, w_b \rangle$. The first element in the tuple expresses the world itself, the second expresses the possibility for Alice, and the third expresses the possibility for Bob. Notably, these epistemic states are exactly the states used in the representation of Section 3. In other words, the representation of Section 3 is now seen to be a particular fragment of the more systematically developed notion of knowledge structure.

5. Propositions, information, and interpretation

The next step in this exposition is to clarify the notion of update used in Section 3. Using the framework of epistemic states, we distinguish the propositional and informational content of a sentence, we define a notion of updating by informational content, and finally we specify how such updates give rise to interpretation shifts.

The basic update mechanism for knowledge structures is conditioning. Upon receiving new information, we eliminate the worlds that are inconsistent with the information, and we also eliminate these worlds from the remaining knowledge structures. However, there are several ways of making the latter process of elimination precise, and the theory of updating knowledge structures has unfortunately not converged onto a clear set of procedures for it.[7] Therefore, rather than reviewing the update procedures in general terms before applying them to the case at hand, I will define an elimination process that directly fits the needs of the epistemic example. When restricted to this highly specific context, the proposed update coincides with the updates considered in Gerbrandy and Groeneveld [10] and Renardel de Lavalette [19], and in this context it is also equivalent with the update rule described in the public announcement logic of Baltag, Moss, and Solecki [2].

[7] It seems that updating in knowledge structures has received relatively little attention, especially after the groundbreaking work of Baltag, Moss and Solecki [2]. Gerbrandy and Groeneveld [10] and Renardel de Lavalette [19] are noteworthy exceptions.

The structure over which the update is defined is the set of epistemic states. We denote individual states with ω, and the set of all states with Ω. Worlds are sets of states, denoted W_i, and these worlds form a partition of the entire set of states, $\mathcal{W} = \{W_1, W_2, \ldots, W_n\}$. In the epistemic example, the states only concern the worlds that the agents can conceive of in each of the worlds. But even with this limited notion of epistemic state, the algebra of states will be richer than the algebra of possible worlds. Some sets of states will not coincide with an element from the algebra of worlds, and this allows us to distinguish between the propositional and the informational content of a sentence.

The *propositional content* of the sentence u, written $\llbracket u \rrbracket \in \mathscr{P}(\mathcal{W})$, consists of all the worlds W_i for which the sentence u is true. It is determined by the interpretation function $I(u, w_i)$. The propositional content $\llbracket r \rrbracket$, for example, is the set of worlds W_i at which r is true, so $\llbracket r \rrbracket = \{W_2, W_3\}$. Notice that the propositional content $\llbracket u \rrbracket$ is different from $[\![u]\!]$, because the latter consists in a set of worlds $w_i \in \boldsymbol{w}$ while the former consists of $W_i \in \mathcal{W}$. Importantly, a world may belong to the propositional content of a sentence in virtue of states within the world W_i. For example, the propositional content of the sentence $\neg K_A r$ is given by $\llbracket \neg K_A r \rrbracket = \{W_1, W_2\}$, because in both of these worlds we find states in which Alice thinks $\neg r$. As explained below, such dependencies between states and propositional content drive the shifts in interpretation.

The *informational content* of a sentence u, written as $\lceil u \rfloor$, is a set of epistemic states that may cut across the worlds W_i. In the epistemic setting, we can derive the informational content from the propositional content. Using only the knowledge structure up to order 1, the propositional and informational content of a sentence u are related according to

$$\lceil u \rfloor \;=\; \{\langle w_i, w_a, w_b \rangle : W_i, W_a, W_b \in \llbracket u \rrbracket\}. \tag{3}$$

In other words, the states $\langle w_i, w_a, w_b \rangle$ in $\lceil u \rfloor$ are those for which all elements are included in $\llbracket u \rrbracket$. So, for example, the information contained in the sentence $\neg K_A r$ is the set of states in which none of the elements is w_3. This is because the propositional content $\llbracket \neg K_A r \rrbracket$ rules out world 3, so that neither Alice nor Bob can conceive of themselves as being in world 3 anymore.

The proposed *update rule* is to condition the set of epistemic states on $\lceil u \rfloor$ when learning the sentence u. The point of this rule is that we do not merely eliminate the worlds inconsistent with u, but also the epistemic states that involve worlds inconsistent with u. The latter stage in the update

corresponds to the revision of the knowledge structures or, in other words, the revision of the epistemic accessibility relations. Notice that, strictly speaking, the informational content of a sentence depends on the cognitive ability and diligence of the agents. In the epistemic example, it is conceivable that Alice and Bob do not think the sentence $\neg K_A r$ through, in which case it presents them with the information of all epistemic situations within worlds 1 and 2, eliminating only the situations in world 3. I take it as an advantage of knowledge structures that they can accommodate such failures of logical omniscience. However, in this paper I assume that the agents think the sentences through, at least up to first order.

It may now seem that we can solve all our problems by using epistemic states as the units in which we interpret the sentences. Why employ a separate notion of possible world if we have a more fine-grained space of epistemic states available? It turns out that a model in which the epistemic states are taken as the units of analysis comes out wrong. In the Kripke models, whether or not a world is included in a certain proposition may depend on the relations this world has with other worlds. Similarly, in the algebra of states the inclusion of a world in a proposition will depend on the states included in the world. If we interpret sentences in terms of epistemic states only, we fail to capture the interplay between propositional content and epistemic states, and thereby loose sight of the interpretation shifts.

We can finally make precise the idea of an *interpretation shift*: the interpretation of a sentence, and hence its propositional content, may change if we update by conditioning on the informational content of another sentence. Recall the interpretation function $I : \mathcal{L} \times \mathcal{W} \to \{0, 1\}$. Now for some sentences, the interpretation will depend on the states included in the world:

$$I(u, w_i) = G_u\left(W_i\right),$$

where G_u is an indicator function whose value may depend on the composition of W_i. After learning a sentence v, the elimination of states in W_i may cause $G_u(W_i)$ to change in value, and hence change the interpretation of u:

$$I_v(u, w_i) = G_u\left(W_i \cap \lceil v \rfloor\right).$$

For sentences u depending on the composition of W_i, the old interpretation I and the new interpretation I_v may differ, thus effecting a shift in interpretation.

Let me illustrate this in terms of the epistemic example. It is because of the inclusion of a state of mind in which Bob thinks that $\neg s$ that we say that world W_2 belongs to $[\![\neg K_B s]\!]$. When learning the sentence $\neg K_A r$,

we condition on the information $\lceil \neg K_A r \rfloor$. This induces a change to the composition of world W_2 because the states 223 and 213 are removed, leaving no state of mind in which Bob thinks $\neg s$. As a result, the propositional content $[\![\neg K_B s]\!]$, which included world W_2 before the information $\lceil \neg K_A r \rfloor$ came in, does not include this world after the information $\lceil \neg K_A r \rfloor$ has been processed.

Summing up, we express beliefs by a probability assignment over sets of worlds, but we frame the information we receive in terms of sets of states that may cut across worlds. Information change is modelled by means of conditioning on information, and this may involve changing the composition of the worlds. As a result, conditioning on new information may require us to redraw the map of propositions. Conditioning thereby captures the way in which the interpretations of sentences change.

6. Probability kinematics for interpretation shifts

The foregoing provides an analysis of how sentences receive different interpretations depending on what we learn. We now provide an account, in Bayesian spirit, of how a modeler adapts her beliefs in response to such interpretation shifts. It turns out that our model of the belief change requires a generalization of the standard probabilistic expressions of belief to Dempster-Shafer belief functions.

We first define these belief functions, as a measure over the space of states Ω, partitioned by the worlds W_i.

Definition 2 (Belief functions). Let $M : \mathscr{P}(\mathcal{W}) \to [0,1]$ be a mass function, $\sum_i M(W_i) = 1$. Belief in the set $\lceil v \rfloor$ is expressed by an interval-valued probability,

$$Bel(\lceil v \rfloor) = \left[\underline{P}(\lceil v \rfloor), \overline{P}(\lceil v \rfloor)\right].$$

where the so-called upper and lower probability are defined as

$$\overline{P}(\lceil v \rfloor) = \min_{u:\, \lceil v \rfloor \subseteq [\![u]\!]} M([\![u]\!]), \tag{4}$$

$$\underline{P}(\lceil v \rfloor) = \max_{u:\, [\![u]\!] \subseteq \lceil v \rfloor} M([\![u]\!]). \tag{5}$$

Notice that $[\![u]\!] \in \mathscr{P}(\mathcal{W})$ is here used as a variable. Of course, if $\lceil v \rfloor \in \mathscr{P}(\mathcal{W})$ as well, then the interval-valued belief collapses onto a point.

We can incorporate new information into the mass function by Dempster's rule of combination.

Definition 3 (Dempster's rule). Let the new information $M^\star$ be a mass function over a partition $\mathcal{V} = \{V_1, \ldots, V_m\}$ of Ω. Then the mass function

after the update is given by

$$M_{\mathcal{V}}(W_i \cap V_j) \;=\; \frac{M(W_i)M^{\star}(V_j)}{\sum_{ij} M(W_i)M^{\star}(V_j)}.$$

The new mass function can be used to generate the new belief function *Bel*. In the case of simple conditioning on the set V, the above partition is simply $\mathcal{V} = \{V, V^c\}$, where V^c is the complement of V, and the mass function has $M^{\star}(V) = 1$. In such cases we write Bel_V for the new belief function.

It will be clear that we can readily apply this update rule to the algebra of epistemic states. To see how this leads to updates in Bayesian spirit, I now explore a number of properties of belief functions and Dempster's rule, focusing on the case that we move from *Bel* to $Bel_{\lceil v \rfloor}$.[8]

First we discuss the relation between belief functions and probabilistic expressions of belief. Recall that the Bayesian model of belief employs sharp probability assignments. Because the probability of epistemic states within worlds W_i is not determined, any probability assignment of the epistemic states is admitted, as long as they sum to the mass assigned to the world as a whole. Each mass function thus corresponds to a set of probability assignments over Ω, denoted $\mathbb{P}$, whose members each comply to the restrictions set by Equations (5) and (4). The restrictions are simply that $P(W_i) \;=\; M(W_i)$ for all i and for each $P \in \mathbb{P}$. Belief functions can be taken as a generalisation of the Bayesian model towards sets of probability assignments.[9]

Secondly, much like belief functions generalize sharp probability functions, Dempster's rule of combination is in some sense a refinement of Bayesian updating. To see the connection, we represent a mass function by a set of probability assignments $\mathbb{P}$ over epistemic states. Now imagine that we learn the information corresponding to the set V. What operation must we perform on the members of the set $\mathbb{P}$ such that the set of updated probability assignments, $P_V \in \mathbb{P}_V$, is a representation of the mass function M_V?

First consider the case in which the new information coincides with a set of worlds, $V \in \mathscr{P}\mathcal{W}$. Let H_i be a function that indicates whether V overlaps with W_i, so $H_i(V) = 0$ if $W_i \cap V = \emptyset$ and $I_i(\lceil v \rfloor) = 1$ otherwise. We can

[8] For a more extensive discussion of Dempster-Shafer belief functions, see Shafer [21] and Halpern [14], pp. 32–40.

[9] Apart from the fact that such sets of probability assignments can be used in a probabilistic model of interpretation shifts, there are independent reasons for expressing belief in these terms. A recent overview of approaches that use sets of probability functions is Haenni *et al* [13]. For more detail on the precise formal relation between belief functions, inner and outer measures, and sets of probability assignments I refer to Halpern [14] and the references therein.

define the set

$$V_{\mathcal{W}} = \{W_i : \, H_i(V) = 1\}.$$

Because in this case $V \in \mathcal{W}$ we have $V = V_{\mathcal{W}}$. The key observation is that the new mass function M_V, derived by Dempster's rule, will correspond to a set of probability assignments $\mathbb{P}_V$ that we can arrive at by performing a Bayesian update with the set $V_{\mathcal{W}}$ on each of the members of $\mathbb{P}$. For all $W_i \notin V_{\mathcal{W}}$ we have $M_V(W_i) = 0$ so that indeed $P_V(W_i) = P(W_i|V_{\mathcal{W}}) = 0$. And for all $W_i \in V_{\mathcal{W}}$, the new mass will be

$$M_V(W_i) = \frac{M(W_i)}{M(V_{\mathcal{W}})}.$$

It follows that $P_V(W_i) = P(W_i|V_{\mathcal{W}})$ as well. For every $P \in \mathbb{P}$, the update with V is thus identical to the update we would have had if the worlds W_i had been the units of analysis.

Next consider the case in which the information V cuts across worlds. For propositions $[\![u]\!] \in \mathcal{W}$ that do not depend on the epistemic states included in the worlds, and hence are not susceptible to interpretation shifts, the new probability is very similar to what Bayesian updating prescribes as well. To see this, simply consider the mass and probability assignment on the level of worlds again. For the sets W_i, Dempster's rule becomes

$$M_V(W_i) \; = \; \frac{H_i(V)\, M(W_i)}{\sum_i H_i(V)\, M(W_i)}.$$

So for any world W_i not intersecting with V, we can deduce that $P_V(W_i) = 0 = P(W_i|V)$, as above. For worlds W_i intersecting with V, on the other hand, we can derive that

$$P_V(W_i) = \frac{M(W_i)}{\sum_i H_i(V) M(W_i)} = P\left(W_i \,|V_{\mathcal{W}}\right).$$

As before, the new mass function M_V, derived by Dempster's rule, will correspond to a set of probability assignments $\mathbb{P}_V$ that we can arrive at by performing a Bayesian update with the set $V_{\mathcal{W}}$ on each of the members of $\mathbb{P}$. So on the level of worlds, Dempster's rule works much the same as the rule of Bayes.

It may be insightful to point to a connection between the present application of Dempster's rule, and a particular application of Lewis' idea of updating by imaging.[10] If the set V is not a proposition, $V \notin \mathscr{P}(\mathcal{W})$,

[10] Cf. Lewis [17], p. 310.

learning the information expressed in V means that the probability of the complement V^c is set to 0 by Dempster's rule. The probability of all those W_i that do not intersect with V is redistributed proportionally over the remaining worlds. The probability of states $\omega \in V^c$ which belong to one of the worlds W_i intersecting with V, on the other hand, is projected onto the remaining states in that world, collected in $W_i \cap V$. In other words, updating by Dempster's rule is similar to updating by imaging in the sense that statesthat receive probability 0 divide their former probability equally over the remaining states in their own world, if there are any. The nearest epistemic state is always one in the same world.

Finally, a remark on sentences whose propositional content depends on the internal structure of the worlds, or in other words, on the states included in it. As explained in the foregoing, the sentence $\neg K_B s$ of the epistemic example is of this kind: a world W_i is a member of the proposition associated with it, if it includes an epistemic state in which Bob thinks that $\neg s$. The probability of the corresponding proposition is not governed by Dempster's rule alone, because the change in the interpretation of $\neg K_B s$ follows directly from the conditioning operation and the nature of the interpretation function. The advantage of Dempster's rule is rather that it allows us to update on a set of states that is not a proposition, and whose probability is consequently not defined. It determines a new probability assignment but it is versatile enough to accommodate changes in interpretation.

7. Application to the epistemic example

We now apply the model of the preceding section to the problem from epistemic logic. The example serves as a stand-in for a much larger class of cases in which interpretations shift, also when these shifts lead to non-trivial probability assignments.

In the example, the epistemic states can be partitioned into three worlds, $\mathcal{W} = \{W_1, W_2, W_3\}$. Following Equation (2), the states are given by $\omega = \langle w_i, w_a, w_b \rangle$ or iab for short. The complete set of states for the example is:

$$\Omega = \{111, 121, 212, 213, 222, 223, 332, 333\}.$$

Any set of epistemic states counts as information that we may learn. But we can make the standard package of information more precise by means of Equation (3). For instance, since $[\![\neg K_A r]\!] = \{W_1, W_2\}$ in virtue of the epistemic states inside W_1 and W_2, we have that $\lceil \neg K_A r \rfloor = \{111, 121, 212, 222\}$, i.e. , all situations that do not involve world 3. In terms of the figure, the

set $\lceil \neg K_A s \rfloor$ consists of the four possibilities in the smaller cube on the right of Figure 7. Similarly, because $[\![K_A r]\!] = \{W_3\}$. we have $\lceil K_A r \rfloor = \{333\}$.

We can now define the belief functions and apply Dempster's rule. We assign a mass function M to the worlds,

$$\begin{aligned} M(W_1) &= M(\{111, 121\}) = 1/3 \\ M(W_2) &= M(\{212, 213, 222, 223\}) = 1/3 \\ M(W_3) &= M(\{332, 333\}) = 1/3 \end{aligned}$$

Thereby we have defined a probability for all elements of $\mathscr{P}(\mathcal{W})$, for instance:

$$P([\![\neg K_A r]\!]) = P(\{W_1, W_2\}) = 2/3.$$

As indicated, the probabilities of the situations within the worlds are not given. Any probability assignment over epistemic states within a world is allowed, as long as it sums to the probability assigned to the world as a whole.

Now say that we learn the information of the sentence $\neg K_A r$, and thus condition on the set $\lceil \neg K_A r \rfloor$. How do we incorporate this into the mass function? First we determine a mass function $M^\star$ that expresses the information in $\neg K_A r$: it simply assigns mass 1 to the set $\lceil \neg K_A r \rfloor$. The mass function after the update is then defined by Dempster's rule. So, for example, after the update we have

$$M_{\lceil \neg K_A r \rfloor}([\![r]\!]) = M_{\lceil \neg K_A r \rfloor}(\{212, 222\}) = 1/2,$$

$$M_{\lceil \neg K_A r \rfloor}([\![s]\!]) = M_{\lceil \neg K_A r \rfloor}(\{111, 121, 212, 222\}) = 1.$$

These are also the resulting probability assignments. On the level of worlds, where the mass function is identical to the probability assignment, Dempster's rule gives the same results that Bayesian conditioning would have given. However, Dempster's rule is suitable for conditioning on sets of states that cut across worlds, like $\lceil \neg K_A r \rfloor$, whereas Bayesian conditioning is not defined for conditioning on such sets.

Now consider how the model deals with changes in interpretation. In the case at hand, we can compute the probability of the proposition $[\![\neg K_B s]\!]$ before and after having conditioned on the information $\lceil \neg K_A r \rfloor$. Before, we have

$$P([\![\neg K_B s]\!]) = P(\{W_2, W_3\}) = 2/3.$$

After the update, however, the sentence $\neg K_B s$ has a propositional content given by $[\![\neg K_B s]\!] = \emptyset$. This does not follow from Dempster's rule, but

rather from the way the sentence $\neg K_B s$ is associated with sets of worlds by the interpretation function. But with that new interpretation in place, we can determine a new mass function $M_{\lceil \neg K_A r \rfloor}$ by Dempster's rule, and compute

$$P_{\lceil \neg K_A r \rfloor}(\llbracket \neg K_B s \rrbracket) = P_{\lceil \neg K_A r \rfloor}(\emptyset) = 0.$$

So by the combination of Dempster's rule and the interpretation function, we manage to model the belief dynamics during an interpretation shift.

Summing up, the model of interpretation shifts in the epistemic context consists of two parts: the interpretation function and particular rules on set membership determine the propositional and informational content of sentences respectively, and Dempster-Shafer theory determines how to assign precise and imprecise probabilities to all these sets. Dempster's rule is versatile enough to accommodate updates by the informational content of a sentence, which need not correspond to a set of worlds.

8. Further research

Clearly, the model presented in this paper is not the full story on epistemic updates, nor on interpretation shifts. However, I hope to have given an idea of how an almost Bayesian model of belief change under interpretation shifts can be developed. The crucial insights are that it makes a distinction between propositional and informational content, that updating by conditioning on informational content may alter the propositional content of sentences, and finally that the probability kinematics for such updates is given by Dempster-Shafer theory.

Let me end with some further applications and extensions of the present model. A natural development in the epistemic setting concerns knowledge of higher order than what is involved in the example. In principle we can refine the set of epistemic states within each world further, to include beliefs of Alice about Bob's beliefs, and so on. The resulting model will still be equivalent to a model using knowledge structures, and thus to a Kripke model. However, it is not clear how the notion of updating used in this paper should be generalized, and it seems that the conception of an update in knowledge structures has not been developed enough to settle this matter either. Moreover, the exact relation between the propositional and informational content has not been worked out in the general case, and the same holds for the dependence of the interpretation function on epistemic states. And finally, there are many open questions on how the present model relates to yet other

approaches, for instance to so-called Harsanyi type spaces.[11] In short, this paper provides a model of interpretation shifts and an illustration thereof in the context of dynamic epistemic logic, with the intention of starting a fruitful exchange. It certainly does not establish a fully developed probabilistic alternative.

One interesting extension of the present model trades on the fact that Dempster's rule allows for updates with non-trivial mass functions $M^\star$. Instead of an update that mimics Bayesian updating based on the information in a single set $\lceil v \rfloor$, this leads to something like an update governed by Jeffrey's rule, based on a probability assignment over a partition $\mathcal{V}$. In such an update we redefine the probability over possible information sets. A curious consequence of this is that information sets whose probability was imprecise may, after the update, receive a precise probability. However, the details of such an update rule have yet to be worked out.

Another possible extension concerns the problems in formal and probabilistic epistemology referred to in Section 1. It will be interesting to see if we can use the model of this paper for clarifying these problems. Of course, once you have a hammer, every problem looks like a nail. But even so, I firmly believe that the present model for belief change is applicable to other domains than epistemic logic, and that violations of reflection, for instance, may benefit from a formal treatment along the lines suggested in this paper. The notions of perspective and state elaborated above have now been given an epistemic interpretation. But the model itself might be useful for capturing other modal interpretations, for instance those of nomic and physical possibility.

Finally, a word of caution. The foregoing might create the suggestion that there cannot be a fully Bayesian model for belief changes involving interpretation shifts. But this is certainly not true. One option, not investigated in this paper, is to take as possible worlds all evolutions of the worlds and the epistemic situations therein. A possible world may be defined by an entry determining the actual world, namely 1, 2, or 3, together with an infinite sequence of such matrices. In terms of these rather elaborate possible worlds, we can again define an algebra, over which we define probability assignments and operations such as conditioning. This algebra will surely be rich enough for accommodating any belief change over epistemic situations and worlds, because it makes the time evolution of the beliefs explicit in the algebra. In this particular case as in formal modeling more generally,

[11] Type spaces also allow for a Bayesian treatment of higher order beliefs. They are often used in game theory. Brandenburger [6] provides a very readable introduction to the idea.

the proper trade-off between the richness of the algebraic structure and the complexity of the update rule is a matter of taste.

Acknowledgements. This research was supported by the Netherlands Organization for Scientific Research (NWO) and the Spanish Ministry of Education. I would like to thank Alexandru Baltag, Barteld Kooi, Johan van Benthem, Eric Pacuit, Ronnie Hermens, Thomas Müller, and audiences at the graduate center of the City University of New York and the PCCP-seminar of the Faculty of Philosophy of the University of Groningen.

References

[1] Arlo-Costa, H., and A. P. Pedersen, Social Norms, Rational Choice and Belief Change. *Belief Revision meets Philosophy of Science*, Logic, Epistemology, and the Unity of Science Series, Vol. 21, E. J. Olsson and S. Enqvist (eds.), Springer, 2010.

[2] Baltag, A., L. S. Moss, and S. Solecki, The logic of public announcements, common knowledge, and private suspicions. *TARK '98: Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge*, Morgan Kaufman, 1998, pp. 43–56.

[3] Baltag, A., and S. Smets, A Qualitative Theory of Dynamic Interactive Belief Revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge (eds.), *Texts in Logic and Games*, Vol. 3, Amsterdam University Press, 2008, pp. 9–58.

[4] Benthem, J. F. A. K. van, Conditional Probability meets Update Logic. *Journal of Logic, Language and Information* 12:409–421, 2003.

[5] Benthem, J. F. A. K. van, J. Gerbrandy, and B. P. Kooi, Dynamic Update with Probabilities. *Studia Logica* 93(1):67–93, 2009.

[6] Brandenburger, A., The Power of Paradox. *International Journal of Game Theory* 35:465–492, 2007.

[7] Ditmarsch, H. van, W. van der Hoek, and B. P. Kooi, *Dynamic Epistemic Logic*. Springer, Dordrecht, 2007.

[8] Fagin, R., J. Y. Halpern, and M. Y. Vardi, A model-theoretic analysis of knowledge. *Journal of the ACM* 38(2):382–428, 1991.

[9] Fraassen, B. C. van, *Laws and Symmetry*. Clarendon Press, Oxford, 1989.

[10] Gerbrandy, J. and W. Groeneveld, Reasoning about information change. *Journal of Logic, Language, and Information* 6:147–169, 1997.

[11] Gerbrandy, J., *Bisimulations on Planet Kripke*. ILLC dissertation series (DS-1999-01), University of Amsterdam, 1999.

[12] Groeneveld, W., *Logical Investigations into Dynamic Semantics*. ILLC dissertation series (DS-1995-18), University of Amsterdam, 1995.

[13] Haenni, R., J. W. Romeijn, G. Wheeler, and J. Williamson, *Probabilistic Logics and Probabilistic Networks*. Synthese Library, Springer, 2010.

[14] Halpern, J. Y., *Reasoning about Uncertainty*. MIT Press, Cambridge (MA), 2003.

[15] Jeffrey, R. C., *Probability and the Art of Judgement*. Cambridge University Press, Cambridge, 1992.

[16] Kooi, B. P., Probabilistic Dynamic Epistemic Logic. *Journal of Logic, Language and Information* 12(4):381–408, 2003.
[17] Lewis, D., Probabilities of Conditionals and Conditional Probabilities. *Philosophical Review* 85(3):297–315, 1976.
[18] Maher, P., *Betting on Theories*. Cambridge University Press, Cambridge, 1993.
[19] Renardel de Lavalette, G. R., Changing modalities. *Journal of Logic and Computation* 14(2):253–278, 2004.
[20] Rott, H., A Counterexample to Six Fundamental Principles of Belief Formation. *Synthese* 139(2):225–240, 2004.
[21] Shafer, G., *A Mathematical Theory of Evidence*. Princeton University Press, Boston, 1976.

Jan-Willem Romeijn
Faculty of Philosophy
University of Groningen
Oude Boteringestraat 52
9712 gl Groningen
The Netherlands
j.w.romeijn@rug.nl