Distance Semantics for Belief Revision *

Daniel Lehmann[†] Menachem Magidor[‡] Karl Schlechta[§]

August 3, 1999

Abstract

A vast and interesting family of natural semantics for belief revision is defined. Suppose one is given a distance d between any two models. One may then define the revision of a theory K by a formula α as the theory defined by the set of all those models of α that are closest, by d, to the set of models of K. This family is characterized by a set of rationality postulates that extends the AGM postulates. The new postulates describe properties of iterated revisions.

1 Introduction

1.1 Overview and related work

The aim of this paper is to investigate semantics and logical properties of theory revisions based on an underlying notion of distance between individual models. In many situations it is indeed reasonable to assume that the agent has some natural way to evaluate the distance between any two models of the logical language of interest. The distance between model m and model m' is a measure of how far m' appears to be from the point of view of m. This distance may measure how different m' is from m under some objective measure: e.g., the number of propositional atoms on which m' differs from m if the language is propositional and finite, but it may also reflect a subjective assessment of the agent about its own capabilities such as, for example, the probability that, if m is the case, the agent (wrongly) believes that m' is the case. Any such distance between models may be used to define a procedure for theory revision: both a theory T and a formula α define a set

*This work was partially supported by the Jean and Helene Alfassa fund for research in Artificial Intelligence and by grant 136/94-1 of the Israel Science Foundation on "New Perspectives on Nonmonotonic Reasoning".

[†]Institute of Computer Science, Hebrew University, 91904 Jerusalem, Israel, lehmann@cs.huji.ac.il

[‡]Institute of Mathematics, Hebrew University, 91904 Jerusalem, Israel, menachem@math.huji.ac.il

[§]Laboratoire d'Informatique de Marseille, CNRS ESA 6077, CMI, 39 rue Joliot Curie, F-13453 Marseille Cédex 13, France, ks@gyptis.univ-mrs.fr

of models, $M(T)$ and $M(\alpha)$, respectively, and the result of revising T by α, $T * \alpha$, is the theory of the set of those models of $M(\alpha)$ that are closest to $M(T)$.
The purpose of this paper is not to suggest specific useful notions of distance. It assumes some abstract notion of a distance is given and studies the properties of the revisions defined by this distance. The distances that will be considered in this paper do not always satisfy the properties generally accepted for distances, they are really pseudo-distances. There is no obvious reason why, in particular, our distance should satisfy the triangular inequality or even be symmetric. It may be the case that, from the perspective of m, m' looks very far away, but, from the perspective of m', m looks close by. In the terms of one of our examples above: if m is the case, our agent may give a very low probability to m' being the case, but if m' is the case, our agent may well be hesitant about whether m or m' is the case: assume, for example, that m and m' differ by the value of one atomic proposition p that is tested by our agent. The test for p may well be very reliable if p is the case but quite unreliable if p is not the case.

1.1.1 Related work

In [AGM85], Alchourrón, Gärdenfors and Makinson introduced the study of theory revision. Their account of revision is indirect: they describe contractions in terms of maximal non-implying sub-theories and they go on to characterize revisions, reducing revision to contraction via the Levi identity. In [Gro88], Grove gave a direct, semantic, characterization of revision. The result of revising a theory K by a proposition α is determined by the models of α that are, individually, *closest* to the set, taken collectively, of all models of K. It thus uses a relation between individual models and sets of models. It is natural to seek to analyze such *closeness* in terms of a distance function between models. A first attempt was made by Becher in [Bec95], in view of comparing revision and update in a unified setting. Becher worked with not necessarily symmetric distances and showed that the AGM postulates hold in distance based revision, but gave no representation result. Independently, the authors presented, in [SLM96], a preliminary version of the results given in this paper. There, only the finite case of symmetric distances was treated. We deal here with the infinite case of symmetric distances and with the finite case of non-symmetric distances. We also provide here proofs and counter-examples. We present, first, our results on an abstract level, dealing with abstract sets and, then, specialize our results to the case of sets of models. In recent work (personal communication) Areces and Becher gave a representation result for the arbitrary, i.e. infinite and not necessarily symmetric case. Their conditions are different from ours, based on complete consistent theories, i.e. single models, and partly in an "existential" style, whereas our conditions are "universal" and more in the AGM style. We do not know whether there is an easy direct, i.e. not going via the semantics, proof of the equivalence of the Areces/Becher and our conditions. Thus, their approach is an alternative route, whose relation to our results is a subject of further research. The infinite case for non-symmetric distances with conditions in our style is still open.

1.1.2 Structure of this paper

After a short motivation in Section 1.1, we present and discuss the AGM framework for revision in Section 1.2 and modify it slightly. Section 1.3 introduces pseudo-distances, which are distances weakened to the properties essential in our context. In particular, pseudo-distances are not necessarily symmetric. We formalize revision based on a (pseudo-) distance, and we show that the usual AGM properties hold for such distance-based revisions, at least in the finite case. An additional property (definability preservation) guarantees them to hold in the infinite case, too. Here, we also discuss some properties of distance-based revision going beyond the AGM postulates.
Our main results are algebraic in nature, and work for arbitrary sets, not only for sets of models. The translation to logic is then straightforward.
Section 2 presents the algebraic representation results, which describe the conditions which guarantee that a binary set operator $\mid$ is representable by a pseudo-distance d, i.e. that $A \mid B$ is the set of $b \in B$ d-closest to A, formally that $A \mid B = A \mid_d B := \{b \in B : \exists a_b \in A \, \forall a' \in A \, \forall b' \in B. d(a_b, b) \leq d(a', b')\}$.
In Section 2.2, we treat the case of symmetric pseudo-distances, the result applies to sets of arbitrary cardinality. Note that this infinite case requires a limit condition: For non-empty sets A, B there is some $b \in B$ with closest distance (among the elements of B) to A. In Section 2.3, we treat the not necessarily symmetric case, but our result applies only to the finite situation.
Section 3 finally translates the results of Section 2 to logic. We there describe the conditions which guarantee that a revision operator $*$ can be represented by a pseudo-distance d between models, i.e. $T * T' := Th(M(T) \mid_d M(T'))$ - where $M(T)$ is the set of models of the theory T, and $Th(X)$ the set of formulas valid in the set of models X.
Analogously to the algebraic characterization, the logical representation results are for possibly infinite languages in the symmetric case (with some caveat about definability preservation) and for finite ones in the not necessarily symmetric case.

1.2 Belief revision

Intelligent agents must gather information about the world, elaborate theories about it and revise those theories in view of new information that, sometimes, contradicts the beliefs previously held. Belief revision is therefore a central topic in Knowledge Representation. It has been studied in different forms: numeric or symbolic, procedural or declarative, logical or probabilistic.

1.2.1 The AGM framework

One of the most successful frameworks in which belief revision has been studied has been proposed by Alchourrón, Gärdenfors and Makinson, and is known as the AGM framework. It deals with operations of revision that revise a theory (the set of previous beliefs) by a formula (the new information). It proposes a set of rationality postulates that any reasonable revision should satisfy. A large number of researchers in AI have been attracted by and have developed this approach further: both in the abstract and by devising revision procedures that satisfy the AGM rationality postulates.

Two remarks should be made immediately. First, the AGM framework presents rationality postulates for revision. It does not choose anyone specific revision among the many possible revisions that satisfy those postulates. Those postulates are justified and defended by the authors, but, recently, some doubts have been expressed as to their desirability, at least for modeling updates, see [KM92] and, more importantly for us, it is not clear that the AGM postulates are all what one would like. A number of authors, in particular [FL94], [DP94], [Leh95], have in fact argued that one would expect some additional postulates to hold. But the consideration of additional postulates has proved slippery and dangerous: the postulates proposed in [DP94] have been shown inconsistent in [Leh95] and have been modified in [DP97]. But this modification forces on us the rejection of one of the basic ontological commitments of the AGM framework, which brings us to a second remark. Secondly, one of the basic ontological commitments of the AGM approach is that what the agent is revising is a belief set. In other terms, epistemic states are belief sets. That this is the AGM position is clear from the formalism chosen: the left hand argument of the star revision operation is a belief set, from the motivation presented, and it is explicitly recognized in [Gr88], p. 47.

Over the years, a large number of researchers have moved away from this identification, sometimes without recognizing it [BG93], [Bou93], [DP97], [Wil94], [NFPS95]. Recently, conclusive evidence has been put forward [Leh95], [FH96] to the effect that this identification of epistemic states to belief sets is not welcome in many AI applications. When iterated revisions are considered, it is reasonable to assume that the agent's epistemic state includes information related to its history of revisions and that all this history, not only the agent's current belief set, may influence future revisions.

This paper keeps the AGM commitment to identifying epistemic states with belief sets, but proposes additional rationality postulates. Those additional postulates characterize exactly the revisions that are defined by pseudo-distances. They constrain revisions in the way they treat their left argument, the theory to be revised (in this respect the AGM postulates are extremely, probably excessively, liberal) and they imply highly non-trivial properties for iterated revisions. This paper therefore treats iterated revisions within the original AGM commitment to the identification of epistemic states and belief sets. Results related to the ideas of this paper may be found in [BGHPSW97]. Similar ideas in a context in which epistemic states are not belief sets may be found in [BLS99].

This work provides a semantics for theory revision à la AGM, or for a sub-family of such revisions. It is the first such effort to describe semantically the whole revision operation $*$ in a unified way. Previous attempts [Gro88], [GM88] describe the revision of each theory K by a different structure without any *glue* relating the different structures: sphere systems or epistemic entrenchment relations, corresponding to different K's. In this paper, the revisions of the different K's are obtained from the same pseudo-distance. A tight fit (coherence) between the revisions of different K's seem crucial for a useful treatment of iterated revisions: it must be the same revision operation that executes the successive revisions for any interesting properties to appear. Our revisions are therefore defined by a polynomial (in the number of models considered) number of pseudo-distances instead of an exponential number of sphere systems or epistemic entrenchment relations (one for each theory K).
Semantics based on a more or less abstract notion of distance is not a new idea in non-classical logics. The best known example is perhaps the Stalnaker/Lewis distance semantics for counterfactual conditionals, see e.g. [Lew73].
The AGM framework, defined in [AGM85], studies revision operations, denoted $*$, that operate on two arguments: a set K of formulas closed under logical deduction on the left and a formula α on the right. Thus $K * \alpha$ is the result of revising theory K by formula α, using revision method $*$.

Notation 1.1

Our base logic will be classical propositional logic, though our main results are purely algebraic in nature, and therefore carry over to other logics, too.
By abuse of language, we call a language finite whose set of propositional variables is finite.
A theory will be an arbitrary set of formulas, not necessarily deductively closed.
We use the customary notation $Cn(T)$ for the set of all logical consequences of a theory T. $Cn(T, \alpha)$ will stand for $Cn(T \cup \{\alpha\})$.
$Con(T)$ will stand for: T is (classically) consistent, $Con(T, T')$ abbreviates $Con(T \cup T')$.
$\models$ will be classical validity, and $\models T \leftrightarrow T'$ will abbreviate the obvious: $\models T \rightarrow \phi'$ for all $\phi' \in T'$ and $\models T' \rightarrow \phi$ for all $\phi \in T$.
Given a propositional language $\mathcal{L}$, $M_{\mathcal{L}}$ will be the set of its models.
$M(T)$ will be the models of a theory T (likewise $M(\phi)$ for a formula ϕ), and $Th(X)$ the set of formulas valid in a set of models X.
$T \vee T' := \{\phi \vee \phi' : \phi \in T, \phi' \in T'\}$.
$\mathcal{P}$ will be the power set operator.
The logical connectives $\wedge$ and $\vee$ and the set connectives $\cap$ and $\cup$ always have precedence over the revision and set operators $*$ and $\mid$.
Numbering of conditions: $(\mid i)$ will number conditions common to the symmetric and the not necessarily symmetric set operators $\mid$, $(\mid Si)$ and $(\mid Ai)$ conditions for respectively the symmetric and the not necessarily symmetric set operator $\mid$. $(*i)$, $(*Si)$, $(*Ai)$ will do the same for the theory revision operator $*$.

The original AGM rationality postulates are the following, for K a deductively closed set of formulas, and α, β formulas.

Definition 1.1
$(K*1)$ $K*\alpha$ is a deductively closed set of formulas.
$(K*2)$ $\alpha \in K*\alpha$.
$(K*3)$ $K*\alpha \subseteq Cn(K,\alpha)$.
$(K*4)$ If $\neg\alpha \notin K$, then $Cn(K,\alpha) \subseteq K*\alpha$.
$(K*5)$ If $K*\alpha$ is inconsistent then α is a logical contradiction.
$(K*6)$ If $\models \alpha \leftrightarrow \beta$, then $K*\alpha = K*\beta$.
$(K*7)$ $K*\alpha\wedge\beta \subseteq Cn(K*\alpha,\beta)$.
$(K*8)$ If $\neg\beta \notin K*\alpha$, then $Cn(K*\alpha,\beta) \subseteq K*\alpha\wedge\beta$.

1.2.2 Modifications of the AGM framework

We prefer to modify slightly the original AGM formalism on two accounts. First, it seems to us that the difference required in the types of the two arguments of a revision: the left argument being a theory and the right argument being a formula is not founded. The lack of symmetry is twofold: the left-hand argument, being a theory, may be inherently infinite and not representable by a single formula while the right-hand argument is always a single formula, but also the left-hand argument is not an arbitrary set of formulas, but closed under logical implication, whereas the right-hand argument is not deductively closed, thus requiring Postulate $K*6$ to assert invariance under logical equivalence. There is no serious reason for this lack of symmetry. We shall therefore prefer a formalism that is symmetric in both arguments.
Our results for symmetric pseudo-distances are valid for infinite sets, our results for not necessarily symmetric pseudo-distances have been proved only for finite sets - where sets are to be understood as sets of models. The latter are thus proved for languages based on a finite set of propositional variables only. We therefore choose a formalism in which both arguments of the revision operator are theories, which, in the not necessarily symmetric case, will be assumed to be equivalent to single formulas. We thus look at $T*T'$, the theory that is the result of revising theory T by the new information represented by theory T'.
Secondly, in the AGM formalism, each one of the theory K and the formula α may be inconsistent. There is no harm in doing so, but the interesting revisions are always revisions of consistent theories by consistent formulas and the consideration of inconsistent arguments makes the treatment unnecessarily clumsy. Therefore, we shall only revise consistent theories by consistent theories, and assume both arguments are consistent.
The AGM postulates may now be rewritten in the following way. We rewrite $K*3$ and $K*4$ in one single postulate, $(*3)$, and similarly for $K*7$ and $K*8$, in $(*4)$. $K*1$ and $K*5$ are summarized in our general prerequisite and $(*1)$.
Remember that T, T', T'', S, S' are now arbitrary consistent theories.

Definition 1.2
$(*0)$ If $\models T \leftrightarrow S$, $\models T' \leftrightarrow S'$, then $T * T' = S * S'$,
$(*1)$ $T * T'$ is a consistent, deductively closed theory,
$(*2)$ $T' \subseteq T * T'$,
$(*3)$ If $T \cup T'$ is consistent, then $T * T' = Cn(T \cup T')$,
$(*4)$ If $T * T'$ is consistent with T'', then $T * (T' \cup T'') = Cn((T * T') \cup T'')$.

1.3 Revision based on pseudo-distances

1.3.1 Pseudo-distances

We will base our semantics for revision on pseudo-distances between models. Pseudo-distances differ from distances in that their values are not necessarily reals, no addition of values has to be defined, and symmetry need not hold. All we need is a totally ordered set of values. If there is a minimal element 0 such that $d(x,y) = 0$ iff $x = y$, we say that d respects identity. Pseudo-distances which do not respect identity have their interest in situations where staying the same requires effort.
We first recollect:

Definition 1.3
A binary relation $\leq$ on X is a preorder, iff $\leq$ is reflexive and transitive. If $\leq$ is in addition total, i.e. iff $\forall x, y \in X$ $x \leq y$ or $y \leq x$, then $\leq$ is a total preorder.
A binary relation $<$ on X is a total order, iff $<$ is transitive, irreflexive, i.e. $x \not< x$ for all $x \in X$, and for all $x, y \in X$ $x < y$ or $y < x$ or $x = y$.

Note 1.1:
If $\leq$ is a total preorder on X, $\approx$ the corresponding equivalence relation defined by $x \approx y$ iff $x \leq y$ and $y \leq x$, [x] the $\approx$ -equivalence class of x, and we define $[x] < [y]$ iff $x \leq y$, but not $y \leq x$, then $<$ is a total order on $\{[x] : x \in X\}$.

Definition 1.4
$d : U \times U \rightarrow Z$ is called a pseudo-distance on U iff (d1) holds:
(d1) Z is totally ordered by a relation $<$.
If, in addition, Z has a $<$ -smallest element 0, and (d2) holds, we say that d respects identity:
(d2) $d(a,b) = 0$ iff $a = b$.
If, in addition, (d3) holds, then d is called symmetric:
(d3) $d(a,b) = d(b,a)$.
(For any $a, b \in U$.)
Let $\leq$ stand for $< \cup =$.

Note that we can force the triangle inequality to hold trivially (if we can choose the values in the real numbers): It suffices to choose the values in the set $\{0\} \cup [0.5, 1]$, i.e. in the interval from 0.5 to 1, or as 0.
Recall that our main representation results are purely algebraic, and apply to arbitrary sets U, which need not necessarily be sets of models. Intuitively however, U is to be understood as the set of models for some language $\mathcal{L}$, and the distance from m to n, $d(m, n)$ represents the "cost" or the "difficulty" of a change from the situation represented by m to the situation represented by n. M. Dalal [Dal88] has considered one such distance: the distance between two propositional worlds is the number of atomic propositions on which they differ, i.e., the Hamming distance between worlds considered as binary k-dimensional vectors, where k is the number of atomic propositional variables. A. Borgida [Bor85] considered a similar but different distance, based on set inclusion. His distances are not totally ordered and therefore the framework presented here does not fit his work. Another example of such a distance is the trivial distance: $d(m, n)$ is 0 if $m = n$ and 1 otherwise.
Both those distances satisfy the triangular inequality. In applications dealing with reasoning about actions and change, one may want to consider the distance between two models to represent how difficult, or unexpected, the transition is. In such a case, a natural pseudo-distance may well not be symmetric.
We give the formal definition of the elements of B d-closest to A:

Definition 1.5
Given a pseudo-distance $d : U \times U \to Z$, let for $A, B \subseteq U$ $A \mid_d B := \{b \in B : \exists a_b \in A \forall a' \in A \forall b' \in B . d(a_b, b) \leq d(a', b')\}$.

Definition 1.5 may be presented in a slightly different light. Put $(a, b) < (a', b')$ iff $d(a, b) < d(a', b')$. Let $min_<(A \times B)$ be the set of all minimal elements (under $<$) of the set $A \times B$. Then, $A \mid_d B$ is nothing else than the right projection (on B) of $A \times B$.
Thus, $A \mid_d B$ is the subset of B consisting of all $b \in B$ that are closest to A. Note that, if A or B is infinite, $A \mid_d B$ may be empty, even if A and B are not empty. A condition assuring non-emptiness will be imposed when necessary.
The aim of Section 2 of this article is to characterize those operators $\mid : \mathcal{P}(U) \times \mathcal{P}(U) \to \mathcal{P}(U)$, for which there is a pseudo-distance d, such that $A \mid B = A \mid_d B$. We call such $\mid$ representable:

Definition 1.6
An operation $\mid$ is representable iff there is a pseudo-distance $d : U \times U \to Z$ such that
(1) $A \mid B = A \mid_d B := \{b \in B : \exists a_b \in A \forall a' \in A \forall b' \in B (d(a_b, b) \leq d(a', b'))\}$.

1.3.2 Revision based on pseudo-distances

The representation results of [AGM95], the semantics of Grove [Gro88] and the very close connection with the rational relations of [LM92], showed in [GM94], all leave essentially

unanswered the question of the nature of the dependence of the revision $T * T'$ on its left argument, T. Since we, like most researchers in Artificial Intelligence, are mostly interested in iterated revisions, proper understanding, and semantics, for this dependence is crucial. The purpose of this paper is to answer the question by proposing a suitable semantics. We completely characterize the semantics by a set of postulates. We do not claim that the semantics proposed are the most general ones, we present one family of reasonable semantics, based on pseudo-distances between models.
The following is the central definition, it describes the way a revision $*_d$ is attached to a pseudo-distance d on the set of models.

Definition 1.7
$T *_d T' = Th(M(T) \mid_d M(T'))$.
$*$ is called representable iff there is a pseudo-distance d on the set of models s.t. $T * T' = Th(M(T) \mid_d M(T'))$.

The main goal of this work is to characterize the properties, i.e., rationality postulates satisfied by revisions representable by pseudo-distances.

1.3.3 Revision based on pseudo-distances and the AGM postulates

The AGM postulates hold for revision based on pseudo-distances in the finite case

Definition 1.8
An operation $\mid$ on the sets of models of some logic is called definability preserving iff $M(T) \mid M(T')$ is again the set of models of some theory S for all theories T, T'.

Abstractly, definability preservation strongly couples proof theory and semantics. To obtain the same kind of results without definability preservation, we would have to allow a "decoupling" on a "small" set of exceptions. This is illustrated e.g. by the results in [Sch97] for the definability preservation case, and in [Sch98] for the unrestricted case of representation results for preferential structures. A similar problem arose already in a finite situation in [ALS98] in the context of partial and total orders, and is treated there by an inductive process.
A first easy result is: any such revision defined for a *finite* language satisfies the AGM postulates $(*0) - (*4)$, if d respects identity. (We use $\mid$ to abbreviate $\mid_d$.)
The same proof shows that the AGM postulates also hold in the infinite case, if the operation $\mid$ is definability preserving, and if we impose a limit condition for postulate $(*1)$.
$(*0)$ is evident, as we work with models.
$(*1)$ holds in the finite case, we will impose it, i.e. a limit condition, in the infinite case.
$(*2)$ trivial by definition.

$(*3)$ this holds, as $d(a,a)$ is minimal for all a, by respect of identity.
$(*4)$ Note that $M(S \cup S') = M(S) \cap M(S')$, and that $M(S * S') = M(S) \mid M(S')$. By prerequisite, $M(T * T') \cap M(T'') \neq \emptyset$, so $(M(T) \mid M(T')) \cap M(T'') \neq \emptyset$. Let $A := M(T)$, $B := M(T')$, $C := M(T'')$. " $\subseteq$": Let $b \in A \mid (B \cap C)$. By prerequisite, there is $b' \in (A \mid B) \cap C$. Thus $d(A,b') \geq d(A, B \cap C) = d(A,b)$. As $b \in B$, $b \in A \mid B$, but $b \in C$, too. " $\supseteq$" : Let $b' \in (A \mid B) \cap C$. Thus $d(A,b') = d(A,B) \leq d(A, B \cap C)$, so by $b' \in B \cap C$ $b' \in A \mid (B \cap C)$. We conclude $M(T) \mid (M(T') \cap M(T'')) = (M(T) \mid M(T')) \cap M(T'')$, thus that $T * (T' \cup T'') = Cn((T * T') \cup T'')$.

The AGM postulate $(*4)$ may fail in the infinite not definability preserving case

The importance of definability preservation is illustrated by the following example, which shows that already the AGM properties may fail when the distance between models does not preserve definability. Essentially the same example will show in Section 3 (Example 3.1 there) that our Loop Condition $(*S1)$ may fail when the distance is not definability preserving. We see here that this is not related to our stronger conditions, but happens already in the general AGM framework.

Example 1.1

Consider an infinite propositional language $\mathcal{L}$.
Let T, T_1, T_2 be complete (consistent) theories, T' a theory with infinitely many models, $M(T) = \{m\}$, $M(T_1) = \{m_1\}$, $M(T_2) = \{m_2\}$, $M(T') = X \cup \{m_1, m_2\}$, $M(T'') = \{m_1, m_2\}$. Assume further $Th(X) = T'$, so X is not definable by a theory.
Arrange the models of $\mathcal{L}$ in the real plane s.t. all $x \in X$ have the same distance < 2 (in the real plane) from m, m_2 has distance 2 from m, and m_1 has distance 3 from m.
(See Figure 1.1.)
Then $M(T) \mid M(T') = X$, but $T * T' = T'$, so $T * T'$ is consistent with T'', and $Cn((T * T') \cup T'') = T''$. But $T' \cup T'' = T''$, and $T * (T' \cup T'') = T_2 \neq T''$. $\square$

AGM revisions are not all definable by pseudo-distances

But any revision defined by a pseudo-distance d also satisfies some properties that do not follow from the AGM postulates. We note again $\mid$ for $\mid_d$.
Consider, for example, the set $C = (A_1 \cup A_2) \mid B$, where A_i and B are finite sets. $d(A,B)$ will be $min\{d(a,b) : a \in A, b \in B\}$.
If $d(A_1,B) < d(A_2,B)$, then $C = A_1 \mid B$. If $d(A_2,B) < d(A_1,B)$, we have $C = A_2 \mid B$. If $d(A_1,B) = d(A_2,B)$, then we have $C = (A_1 \mid B) \cup (A_2 \mid B)$. It follows that any revision defined by a pseudo-distance satisfies (for a finite language): $(\alpha_1 \vee \alpha_2) * \beta$ is equal to $(\alpha_1 * \beta) \cap (\alpha_2 * \beta)$, to $\alpha_1 * \beta$, or to $\alpha_2 * \beta$.

Figure 1.1

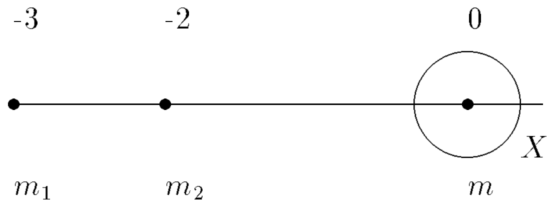

This property does not follow from the AGM postulates, as will be shown below, but seems a very natural property. Indeed, when revising a disjunction $\alpha_1 \vee \alpha_2$ by a formula β, there are two possibilities. First, it may be the case that our indecision concerning α_1 or α_2 persists after the revision, and, in this case, the revised theory is naturally the disjunction of the revisions. But it may also be the case that the new information β makes us revise backwards and conclude that it must be the case that α_1 or, respectively, α_2 was (before the new information) the better theory and, in this case, the revised theory should be $\alpha_1 * \beta$ or $\alpha_2 * \beta$. Notice that this last property of revisions generated by pseudo-distances is the left argument analogue of AGM's Ventilation Principle which concerns the argument on the right. The Ventilation Principle follows from the AGM postulates and states that: $\alpha * (\beta_1 \vee \beta_2)$ is equal to $(\alpha * \beta_1) \cap (\alpha * \beta_2)$, to $\alpha * \beta_1$ or to $\alpha * \beta_2$.
One can conclude that any revision defined by a pseudo-distance satisfies the following properties, that deal with iterated revisions:
if $\delta \in (K * \alpha) * \gamma$ and $\delta \in (K * \beta) * \gamma$, then $\delta \in (K * (\alpha \vee \beta)) * \gamma$
and
if $\delta \in (K * (\alpha \vee \beta)) * \gamma$, then, either $\delta \in (K * \alpha) * \gamma$ or $\delta \in (K * \beta) * \gamma$.
Those properties seem intuitively right. If after any one of two sequences of revisions that differ only at step i (step i being α in one case and β in the other), one would conclude that δ holds, then one should conclude δ after the sequence of revisions that differ from the two revisions only in that step i is a revision by the disjunction $\alpha \vee \beta$, since knowing which of α or β is true cannot be crucial. This property is an analogue for the left argument of the Or property of [KLM90]. Similarly, if one concludes δ from a revision by a disjunction, one should conclude it from at least one of the disjuncts. This property is an analogue for the left argument of the Disjunctive Rationality property of [KLM90], studied in [Fre93]. It is easy to see that the property (C1) of Darwiche and Pearl [DP94], i.e., $(K * \alpha) * (\alpha \wedge \beta) = K * (\alpha \wedge \beta)$ is not satisfied by all revisions defined by pseudo-distances. Section 2 will precisely characterize those revisions that are defined by pseudo-distances.

Notice that in each of the AGM postulates, the left-hand side argument of the revision operation $(*)$ is the same all along: all revisions have the form $K*$. Since, as has been shown above, every revision defined by a pseudo-distance satisfies the AGM postulates, if, for each theory K we define $K*$ by some pseudo-distance, then the revision defined will satisfy the AGM postulates, even if we use different pseudo-distances for different K's.
Consider, for a simple example, 4 points in the real plane, a, b, c, d, to be interpreted as the models of a propositional language of two variables. Let a have the coordinates (0,1), b (0,-1), c (1,0), d (2,0), and define by the natural distance the revisions with any $X \neq \emptyset$ except for $X := \{a, b\}$ on the left hand side. As seen above, they will satisfy the AGM postulates. To define the revisions with $\{a, b\}$ on the left hand side, interchange the positions of c and d. This, too, satisfies the AGM postulates. As the AGM postulates say nothing about coherence between different K's, all these revisions together satisfy the AGM postulates.
But we will then have $\{a\} \mid \{c, d\} = \{b\} \mid \{c, d\} = \{c\}$, but $\{a, b\} \mid \{c, d\} = \{d\}$, so such a system of revisions cannot be defined by a pseudo-distance.

2 The algebraic representation results

2.1 Introduction

First, a generalized abstract nonsense result. This result is certainly well-known and we claim no priority. It will be used repeatedly below to extend a relation R to a relation S. The equivalence classes under S will be used to define the abstract distances.

Lemma 2.1
Given a set X and a binary relation R on X, there exists a total preorder S on X that extends R such that
(2) $\forall x, y \in X(xSy, ySx \Rightarrow xR^*y)$
where R^* is the reflexive and transitive closure of R.

Proof:
Define $x \equiv y$ iff xR^*y and yR^*x. The relation $\equiv$ is an equivalence relation. Let $[x]$ be the equivalence class of x under $\equiv$. Define $[x] \preceq [y]$ iff xR^*y. The definition of $\preceq$ does not depend on the representatives x and y chosen. The relation $\preceq$ on equivalence classes is a partial order: reflexive, antisymmetric and transitive. A partial order may always be extended to a total order. Let $\leq$ be any total order on these equivalence classes that extends $\preceq$. Define xSy iff $[x] \leq [y]$. The relation S is total (since $\leq$ is total) and transitive (since $\leq$ is transitive): it is a total preorder. It extends R by the definition of $\preceq$ and the fact that $\leq$ extends $\preceq$. Let us show that it satisfies Equation (2) of Lemma 2.1. Suppose xSy and ySx. We have $[x] \leq [y]$ and $[y] \leq [x]$ and therefore $[x] = [y]$ by antisymmetry ($\leq$ is an order relation). Therefore $x \equiv y$ and xR^*y. $\Box$

The algebraic representation results we are going to demonstrate in this Section 2 are independent of logic, and work for arbitrary sets U, not only for sets of models. On the other hand, if the (propositional) language $\mathcal{L}$ is defined from infinitely many propositional variables, not all sets of models are definable by a theory: There are $X \subseteq M_{\mathcal{L}}$ s.t. there is no T with $X = M(T)$. Moreover, we will consider only consistent theories. This motivates the following:
Let $U \neq \emptyset$, and let $\mathcal{Y} \subseteq \mathcal{P}(U)$ contain all singletons, be closed under finite non-empty $\cap$ and finite $\cup$, $\emptyset \notin \mathcal{Y}$ and consider an operation $|: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathcal{Y}$. (For our representation results, finite $\cap$ suffices.)
We are looking for a characterization of representable operators. We first characterize those operations $|$ which can be represented by a symmetric pseudo-distance in Section 2.2, and then those representable by a not necessarily symmetric pseudo-distance in Section 2.3.

Notation 2.1
For $a \in U$, $X \in \mathcal{Y}$ $a \mid X$ will stand for $\{a\} \mid X$ etc.

2.2 The result for symmetric pseudo-distances

We work here with possibly infinite, but nonempty U.
Both Example 2.1 and Example 2.2 show that revision operators are relatively coarse instruments to investigate distances. The same revision operation can be based on many different distances. Consequently, in the construction of the distance from the revision operation, one still has a lot of freedom left. Example 2.2 will show that, in the case one does not require symmetric distances, the freedom is even greater. The reader should note that the situation described in Example 2.2 corresponds to the remark in the proof of Proposition 2.5, that the constructed distance d does not necessarily satisfy $d(A, B) = min\{d(a, B) : a \in A\}$, i.e., may behave strangely on the left hand side. But even when the pseudo-distance is a real distance, the resulting revision operator $|_d$ does not always permit reconstructing the relations of the distances.
Distances with common start (or end, by symmetry) can always be compared by looking at the result of revision:
$a \mid_d \{b, b'\} = b$ iff $d(a, b) < d(a, b')$,
$a \mid_d \{b, b'\} = b'$ iff $d(a, b) > d(a, b')$,
$a \mid_d \{b, b'\} = \{b, b'\}$ iff $d(a, b) = d(a, b')$.
This is not the case with arbitrary distances $d(x, y)$ and $d(a, b)$, as the following example will show.

Example 2.1
We work in the real plane, with the standard distance, the angles have 120 degrees. a' is closer to y than x is to y, a is closer to b than x is to y, but a' is farther away from b'

Figure 2.1

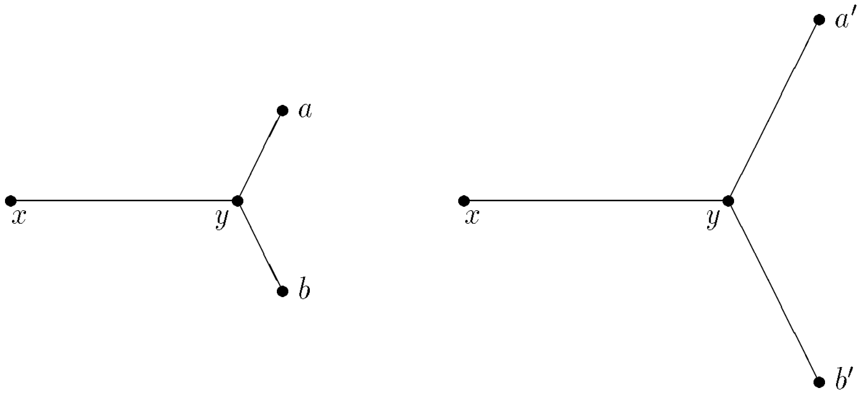

than x is from y. Similarly for b,b'. But we cannot distinguish the situation $\{a, b, x, y\}$ and the situation $\{a', b', x, y\}$ through $|_d$. (See Figure 2.1.)

Proof:

Seen from a, the distances are in that order: y, b, x.
Seen from a', the distances are in that order: y, b', x.
Seen from b, the distances are in that order: y, a, x.
Seen from b', the distances are in that order: y, a', x.
Seen from y, the distances are in that order: $a/b, x$.
Seen from y, the distances are in that order: $a'/b', x$.
Seen from x, the distances are in that order: $y, a/b$.
Seen from x, the distances are in that order: $y, a'/b'$.
Thus, any $c \,|_d\, C$ will be the same in both situations (with a interchanged with a', b with b'). The same holds for any $X \,|_d\, C$ where X has two elements.
Thus, any $C \,|_d\, D$ will be the same in both situations, when we interchange a with a', and b with b'. So we cannot determine by $|_d$ whether $d(x, y) > d(a, b)$ or not. $\square$

Proposition 2.2

Let $U \neq \emptyset$, $\mathcal{Y} \subseteq \mathcal{P}(U)$ be closed under finite non-empty $\cap$ and finite $\cup$, $\emptyset \notin \mathcal{Y}$.
Let $A, B, X_i \in \mathcal{Y}$.
Let $|: \mathcal{Y} \times \mathcal{Y} \to \mathcal{Y}$, and consider the conditions
$(\mid 1)$ $A \mid B \subseteq B$
$(\mid 2)$ $A \cap B \neq \emptyset \to A \mid B = A \cap B$
$(\mid S1)$ (Loop): $(X_1 \mid (X_0 \cup X_2)) \cap X_0 \neq \emptyset$, $(X_2 \mid (X_1 \cup X_3)) \cap X_1 \neq \emptyset$, $(X_3 \mid (X_2 \cup X_4)) \cap X_2 \neq \emptyset$, ... $(X_k \mid (X_{k-1} \cup X_0)) \cap X_{k-1} \neq \emptyset$ imply $(X_0 \mid (X_k \cup X_1)) \cap X_1 \neq \emptyset$.
(a) $\mid$ is representable by a symmetric pseudo-distance $d : U \times U \to Z$ iff $\mid$ satisfies $(\mid 1)$ and $(\mid S1)$.
(b) $\mid$ is representable by an identity-respecting symmetric pseudo-distance $d : U \times U \to Z$ iff $\mid$ satisfies $(\mid 1)$, $(\mid 2)$, and $(\mid S1)$.

Note that $(\mid 1)$ corresponds to $(*2)$, $(\mid 2)$ to $(*3)$, $(*0)$ will hold trivially, $(*1)$ holds by definition of $\mathcal{Y}$ and $\mid$, $(*4)$ will be a consequence of representation. $(\mid S1)$ corresponds to: $d(X_1, X_0) \leq d(X_1, X_2)$, $d(X_2, X_1) \leq d(X_2, X_3)$, $d(X_3, X_2) \leq d(X_3, X_4) \leq \ldots \leq d(X_k, X_{k-1}) \leq d(X_k, X_0) \to d(X_0, X_1) \leq d(X_0, X_k)$, and, by symmetry, $d(X_0, X_1) \leq d(X_1, X_2) \leq \ldots \leq d(X_0, X_k) \to d(X_0, X_1) \leq d(X_0, X_k)$, i.e. transitivity of $\leq$, or to absence of loops involving $<$.
We first show the hard direction via a number of auxiliary definitions and lemmas (up to Fact 2.4). We assume all A,B etc. to be in $\mathcal{Y}$, and $(\mid 1)$, $(\mid S1)$ to hold from now on.
We first define a precursor $\parallel A, B \parallel$ to the pseudo-distance between A and B, and a relation $\leq$ on these $\parallel A, B \parallel' s$. We then prove some elementary facts about $\parallel \ldots \parallel$ and $\leq$ in Fact 2.3. We extend $\leq$ to a total preorder S using Lemma 2.1, the pseudo-distances will be S-equivalence classes of the $\parallel A, B \parallel' s$. It remains to show that the revision operation $|_d$ defined by this pseudo-distance is the same as the operation we started with, this is shown in Fact 2.4.

Definition 2.1
Set $\parallel A, B \parallel \leq \parallel A, B' \parallel$ iff $(A \mid B \cup B') \cap B \neq \emptyset$,
set $\parallel A, B \parallel < \parallel A, B' \parallel$ iff $\parallel A, B \parallel \leq \parallel A, B' \parallel$, but not $\parallel A, B \parallel \geq \parallel A, B' \parallel$.

$\parallel A, B \parallel$ is to be read as the pseudo-distance between A and B or between B and A. Recall that the pseudo-distance will be symmetric, so $\parallel \ldots \parallel$ operates on the unordered pair $\{A, B\}$. Note that $A \mid B \neq \emptyset$, by definition of the function $\mid$.
Let $\leq^*$ be the transitive closure of $\leq$, we write also $<^*$ if it involves $<$. Write $\parallel a, B \parallel$ for $\parallel \{a\}, B \parallel$ etc.
The loop condition reads in the $\parallel$-notation as follows: $\parallel X_0, X_1 \parallel \leq \parallel X_2, X_1 \parallel \leq \parallel X_2, X_3 \parallel \leq \parallel X_4, X_3 \parallel \leq \ldots \leq \parallel X_k, X_{k-1} \parallel \leq \parallel X_k, X_0 \parallel \to \parallel X_0, X_1 \parallel \leq \parallel X_0, X_k \parallel$

Fact 2.3
(1) $\parallel A, B \parallel \not\leq \parallel A, B' \parallel$ iff $\parallel A, B' \parallel < \parallel A, B \parallel$
(2) $B' \subseteq B \to \parallel A, B \parallel \leq \parallel A, B' \parallel$

(3) There are no cycles of the forms $\| A,B \| \leq \| A,B' \| \leq \ldots \leq \| A,B'' \| \leq \| A,B \|$ or $\| A,B \| \leq \| A,B' \| \leq \ldots \leq \| A'',B \| \leq \| A,B \|$ involving $<$. (The difference between the two cycles is that the first contains possibly only variations on one side, of the form $\| A,B'' \| \leq \| A,B \| \leq \| A,B' \|$, the second one possibly only alternating variations, of the form $\| A'',B \| \leq \| A,B \| \leq \| A,B' \|$.)
(4) $b \in A \mid B \rightarrow \| A,b \| \leq \| A,B \|$
(5) $b \notin A \mid B$, $b \in B \rightarrow \| A,B \| < \| A,b \|$
(6) $\| A,b \| \leq^* \| A,B \|$, $b \in B \rightarrow b \in A \mid B$
(7) $b \in A \mid B$, $a_b \in b \mid A$, $a_b \in A' \subseteq A$ implies (a) $b \in A' \mid B$, (b) $A' \mid B \subseteq A \mid B$.
(8) $b \in A \mid B$, $a_b \in b \mid A$, $a' \in A$, $b' \in B \rightarrow \| a_b, b \| \leq^* \| a', b' \|$
(9) $b \in B$, $b \notin A \mid B$, $b' \in A \mid B$, $a_{b'} \in b' \mid A$, $a \in A$. Then $\| a_{b'}, b' \| <^* \| a, b \|$.
If ($\mid 2$) holds, then
(10) $A \cap B \neq \emptyset \rightarrow \| A,B \| \leq^* \| A',B' \|$
(11) $A \cap B \neq \emptyset$, $A' \cap B' = \emptyset \rightarrow \| A,B \| <^* \| A',B' \|$

Proof:
(1) and (2) are trivial.
(3) We prove both variants simultaneously. Case 1, length of *cycle* $= 1$: $\| A,B \| < \| A,B \|$, so $(A \mid B) \cap B = \emptyset$, contradiction. Case 2: length > 1 : Let e.g. $\| A_0,B_0 \| \leq \| A_0,B_1 \| \leq \ldots \leq \| A_0,B_k \| < \| A_0,B_0 \|$ be such a cycle. If the cycle is not yet in the form of the loop condition, we can build a loop as in the loop condition by repeating elements, if necessary. E.g.: $\| A_0,B_0 \| \leq \| A_0,B_1 \| \leq \| A_0,B_2 \|$ can be transformed to $\| A_0,B_0 \| \leq \| A_0,B_1 \| \leq_{by\ (2)} \| A_0,B_1 \| \leq \| A_0,B_2 \|$. By Loop, we conclude $\| A_0,B_0 \| \leq \| A_0,B_k \|$, contradicting (1).
(4) and (5) are trivial.
(6) $b \notin A \mid B \rightarrow_{by\ (5)} \| A,B \| < \| A,b \|$, contradicting $\| A,b \| \leq^* \| A,B \|$ by (3).
(7) (a) By (6), it suffices to show that $\| A',b \| \leq^* \| A',B \|$. But $\| A',b \| \leq_{by\ (2)} \| a_b,b \| \leq^*_{(4)\ twice} \| A,B \| \leq_{by\ (2)} \| A',B \|$. (b) Let $b' \in A' \mid B$, we show $b' \in A \mid B$. By (6), it suffices to show $\| A,b' \| \leq^* \| A,B \|$: $\| A,b' \| \leq_{(2)} \| A',b' \| \leq_{(4)} \| A',B \| \leq^*_{(2)\ twice} \| a_b,b \| \leq^*_{(4)\ twice} \| A,B \|$.
(8) $\| a_b,b \| \leq^* \| A,B \| \leq^* \| a',b' \|$.
(9) $\| a_{b'},b' \| \leq^*_{(4)\ twice} \| A,B \| <_{(5)} \| A,b \| \leq_{(2)} \| a,b \|$.
(10) $\| A,B \| \leq \| A,B \cup B' \|$, as $(A \mid B \cup B') \cap B \neq \emptyset$, by $A \cap B \subseteq A \mid B \cup B'$. Likewise $\| A,B \cup B' \| \leq \| A \cup A',B \cup B' \|$. Moreover, $\| A \cup A',B \cup B' \| \leq \| A',B' \|$ by (2).
(11) We show first that $A \cap B \neq \emptyset$, $A \cap B' = \emptyset$ implies $\| A,B \| < \| A,B' \|$: $A \mid B \cup B' = A \cap (B \cup B') = A \cap B \subseteq A$, so $(A \mid B \cup B') \cap B' = \emptyset$. Thus, $\| A,B \| \leq^*_{by\ (10)} \| A',A' \| < \| A',B' \|$. □

We define:

Definition 2.2
Let S, by Lemma 2.1, be a total preorder on $\{\parallel A,B \parallel: A,B \in \mathcal{Y}\}$ extending $\leq$ s.t. $\parallel A,B \parallel S \parallel A',B' \parallel$ and $\parallel A',B' \parallel S \parallel A,B \parallel$ imply $\parallel A,B \parallel \leq^* \parallel A',B' \parallel$.
Let $\parallel A,B \parallel \approx \parallel A',B' \parallel$ iff $\parallel A,B \parallel S \parallel A',B' \parallel$ and $\parallel A',B' \parallel S \parallel A,B \parallel$, and $[\parallel A,B \parallel]$ be the set of $\approx$-equivalence classes and define $[\parallel A,B \parallel] < [\parallel A',B' \parallel]$ iff $\parallel A,B \parallel S \parallel A',B' \parallel$ but not $\parallel A',B' \parallel S \parallel A,B \parallel$. This is a total order on $\{[\parallel A,B \parallel]: A,B \in \mathcal{Y}\}$. Define $d(A,B) := [\parallel A,B \parallel]$ for $A,B \in \mathcal{Y}$.
If $(\mid 2)$ holds, let $0 := [\parallel A,A \parallel]$ for any A. This is then well-defined by Fact 2.3, (10).
Note that by abuse of notation, we use $\leq$ also between equivalence classes.

Fact 2.4
(1) The restriction to singletons of d as just defined is a symmetric pseudo-distance; if $(\mid 2)$ holds, then d respects identity.
(2) $A \mid B = A \mid_d B$.

Proof:
(1)
(d1) Trivial. If $[\parallel b,c \parallel] < [\parallel a,a \parallel]$, then $\parallel b,c \parallel \leq^* \parallel a,a \parallel$, but not $\parallel a,a \parallel \leq^* \parallel b,c \parallel$, contradicting Fact 2.3, (10).
(d2) $d(a,b) = d(a,a)$ iff $\parallel a,b \parallel \leq^* \parallel a,a \parallel$ iff $a = b$ by Fact 2.3, (10) and (11).
(d3) $[\parallel a,b \parallel] \leq [\parallel b,a \parallel]$ is trivial.
(2)
"$\subseteq$": Let $b \in A \mid B$. Then there is $a_b \in b \mid A$. By Fact 2.3, (8), $\parallel a_b, b \parallel \leq^* \parallel a',b' \parallel$ for all $a' \in A$, $b' \in B$. So $d(a_b,b) \leq d(a',b')$ for all $a' \in A$, $b' \in B$ and $b \in A \mid_d B$.
"$\supseteq$": Let $b \in B$, $b \notin A \mid B$. Take $b' \in A \mid B$, $a_{b'} \in b' \mid A$, $a \in A$. Then by Fact 2.3, (9) $\parallel a_{b'}, b' \parallel <^* \parallel a,b \parallel$, so $b \notin A \mid_d B$.
□

It remains to show the easy direction of Proposition 2.2.
All conditions but $(\mid S1)$ are trivial. Define for two sets $A,B \neq \emptyset$ $d(A,B) := d(a_b,b)$, where $b \in A \mid_d B$, and $a_b \in b \mid_d A$. Then $d(A,B) = d(B,A)$ by $d(a,b) = d(b,a)$ for all a,b. Loop amounts thus to $d(X_1,X_0) \leq \ldots \leq d(X_k,X_0) \rightarrow d(X_0,X_1) \leq d(X_0,X_k)$, which is now obvious. □ (Proposition 2.2)

2.3 The result for not necessarily symmetric pseudo-distances

Note that we work here with finite U only, $\mathcal{Y}$ will be $\mathcal{P}(U) - \{\emptyset\}$.

Figure 2.2

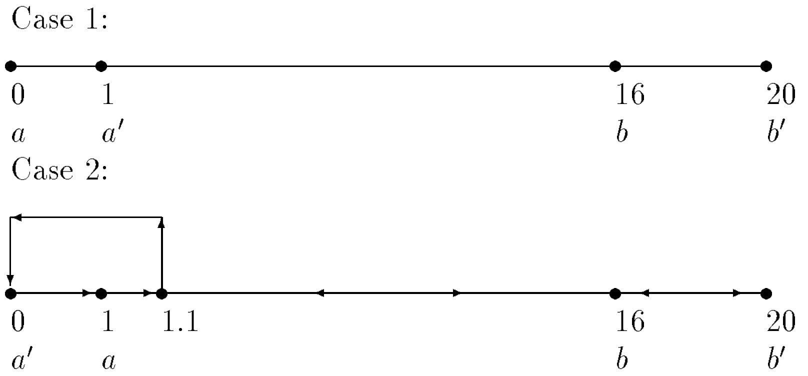

We first give an Example, which illustrates the expressive weakness of a not necessarily symmetric distance.

Example 2.2
This example, illustrated in Figure 2.2, shows that we cannot find out, in the non symmetric case, which of the elements a, a' is closest to the the set $\{b, b'\}$ (we look from a/a' to $\{b, b'\}$). In the first case, it is a', in the second case a. Yet all results about revision stay the same.
In the first case, we can take the "road" in both directions, in the second case, we have to follow the arrows. (For simplicity, the vertical parts have length 0.) Otherwise, distances are as indicated by the numbers, so e.g. in the second case, from a' to a it is 1, from a to a' 1.2. For any $X, Y \subseteq \{a, a', b, b'\}$ $X \mid Y$ will be the same in both cases, but, seen from a or a', the distance to $\{b, b'\}$ is closer from a' in the first case, closer from a in the second.

The characterization of the not necessarily symmetric case presented in the following perhaps does not seem very elegant at first sight, but it is straightforward and very useful in the search for more elegant characterizations of similar operations. For our characterization a definition is necessary. It associates a binary relation between pairs of non-empty subsets of U : intuitively, $(A, B)R_{|}(A', B')$ may be understood as meaning that the revision $|$ requires the pseudo-distance between A and B to be smaller than or equal to that between A' and B'. The main idea of the representation theorem is to define a relation (the relation $R_{|}$ of Definition 2.3) that describes all inequalities we know must hold between pseudo-distances, and require that the consequences of those inequalities are upheld (conditions $(\mid A2)$ and $(\mid A3)$ of Proposition 2.5). The proof of the theorem shows that Definition 2.3 was comprehensive enough.

Definition 2.3
Given an operation $\mid$, one defines a relation $R_{\mid}$ on pairs of non-empty subsets of U by: $(A,B)R_{\mid}(A',B')$ iff one of the following two cases obtains:
(1) $A = A'$ and $(A \mid B \cup B') \cap B \neq \emptyset$,
(2) $B = B'$ and $(A \cup A' \mid B) \neq (A' \mid B)$,
If the pseudo-distance is to respect identity, we also consider a third case:
(3) $A \cap B \neq \emptyset$.

Definition 2.3 can be written as:
(1) $(A \mid B \cup B') \cap B \neq \emptyset \Rightarrow (A,B)R_{\mid}(A,B')$,
(2) $(A \cup A' \mid B) \neq (A' \mid B) \Rightarrow (A,B)R_{\mid}(A',B)$,
(3) $A \cap B \neq \emptyset \Rightarrow (A,B)R_{\mid}(A',B')$.
In the sequel we shall write R instead of $R_{\mid}$. As usual, we shall denote by R^* the reflexive and transitive closure of R.
Notice also that we do not require that the pseudo-distance between A and B be less or equal than that between A' and B' if $A' \subseteq A$ and $B' \subseteq B$, as one could expect. In fact, a theorem similar to Proposition 2.5 below may be proved with a definition of R that includes a fourth case: $(A,B)R(A',B')$ if $A' \subseteq A$ and $B' \subseteq B$, and its proof is slightly easier, but we prefer to prove the stronger theorem. Notice also that, in order to avoid the fourth case just mentioned, the conclusion of case (2) is $(A,B)R_{\mid}(A',B)$, and not the seemingly stronger but in fact weaker in the absence of the fourth case mentioned above: $(A,B)R_{\mid}(A \cup A',B)$.
We may now formulate our main technical result. Condition $(\mid A1)$ expresses a property of Disjunctive Rationality ([KLM90], [LM92], [Fre93]) for the left-hand-side argument of the operation $\mid$.

Proposition 2.5
Consider the following conditions:
$(\mid 1)$ $(A \mid B) \subseteq B$,
$(\mid A1)$ $(A \cup A' \mid B) \subseteq (A \mid B) \cup (A' \mid B)$,
$(\mid A2)$ If $(A,B)R^*(A,B')$, then $(A \mid B) \subseteq (A \mid B \cup B')$,
$(\mid A3)$ If $(A,B)R^*(A',B)$, then $(A \mid B) \subseteq (A \cup A' \mid B)$,
$(\mid 2)$ If $A \cap B \neq \emptyset$, then $A \mid B = A \cap B$,
$(\mid A4)$ If $(A,B)R^*(A',B')$ and $A' \cap B' \neq \emptyset$, then $A \cap B \neq \emptyset$.
(a) An operation $\mid: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathcal{Y}$ is representable by a pseudo-distance iff it satisfies the conditions $(\mid 1)$, $(\mid A1) - (\mid A3)$ for any non-empty sets $A, B \subseteq U$, where the relation R is generated by cases (1) and (2) of Definition 2.3.
(b) An operation $\mid: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathcal{Y}$ is representable by an identity-respecting pseudo-distance iff it satisfies the conditions $(\mid 1)$, $(\mid 2)$, $(\mid A1) - (\mid A4)$ for any non-empty sets $A, B \subseteq U$, where the relation R is generated by cases (1) - (3) of Definition 2.3.

Proof:

First, we shall deal with the soundness part of the theorem, and then with the more challenging completeness part. We prove (a) and (b) together.

Suppose, then, that $\mid$ is representable by a pseudo-distance. The function d acts on pairs of elements of U, and it may be extended to a function on pairs of non-empty subsets of U in the usual way: $d(A,B) = min\{d(a,b) : a \in A, b \in B\}$.

Then Equation (1) in Definition 1.6, defining representability, may be written as:

(3) $A \mid B = \{b \in B \mid d(A,b) = d(A,B)\}$.

We must now show that the conditions of Proposition 2.5 hold.

Condition ($\mid 1$) is obvious.

Condition ($\mid A1$) holds since $d(A \cup A', B) = min\{d(A,B), d(A',B)\}$.

Considering Definition 1.6 and the different cases of Definition 2.3, we shall see that $(A,B)R(A',B')$ implies $d(A,B) \leq d(A',B')$. Case 1 is obvious. Let us treat case (2). Clearly $d(A \cup A', B) = min\{d(A',B), d(A,B)\}$. We shall show that if $d(A',B) < d(A,B)$, then $A \cup A' \mid B = A' \mid B$. Suppose $d(A',B) < d(A,B)$. Then, $d(A \cup A', B) = d(A',B) < d(A,B)$. Therefore $A \cup A' \mid B = A' \mid B$. Case 2 has been taken care of. If d respects identity, Case 3 is obvious. We conclude that $(A,B)R^*(A',B')$ implies that $d(A,B) \leq d(A',B')$. Condition ($\mid A2$) holds because $d(A,B) \leq d(A,B')$ implies $d(A, B \cup B') = d(A,B)$. Condition ($\mid A3$) holds because $d(A,B) \leq d(A',B)$ implies $d(A \cup A', B) = d(A,B)$.

It remains to show that ($\mid 2$) and ($\mid A4$) follow from respect of identity:

Condition ($\mid 2$) holds because $A \mid B = \{b \in B : d(A,b) = d(A,B) = 0\}$ if $A \cap B \neq \emptyset$. Condition ($\mid A4$) holds because $d(A,B) \leq d(A',B') = 0$ implies $d(A,B) = 0$.

For the other direction, we work, unless stated otherwise, in the base situation, i.e. where at least conditions ($\mid 1$), ($\mid A1$) $-$ ($\mid A3$) hold, and the relation R is generated by at least cases (1) and (2) of Definition 2.3.

In our proof, a number of lemmas will be needed. These lemmas will be presented when needed, and their proof inserted in the midst of the proof of Proposition 2.5. Again, we first extend the relation R on pairs (A,B) to a total preorder S using Lemma 2.1, and use the S-equivalence classes as pseudo-distances. Recall that the d thus defined will behave nicely on the right hand side, but not necessarily on the left hand side: $d(A,B) = min\{d(A,b) : b \in B\}$ will hold, but not necessarily $d(A,B) = min\{d(a,B) : a \in A\}$. Again, it remains to show that the revision operation $\mid_d$ defined by this pseudo-distance is the same as the operation we started with, this is done in the rest of the proof.

First, a simple result, analogous to the Or rule of [KLM90].

Lemma 2.6

For any sets A, A', B, $(A \mid B) \cap (A' \mid B) \subseteq A \cup A' \mid B$.

Proof:

Without loss of generality we may assume that $A \mid B \neq A \cup A' \mid B$. Then $(A',B)R(A,B)$ by case (2) of Definition 2.3, and $A' \mid B \subseteq A \cup A' \mid B$ by condition ($\mid A3$) of Proposition

2.5. $\square$

We consider the set $\mathcal{Y}\times\mathcal{Y}$ and the binary relation R on this set defined from $\mid$ by Definition 2.3. By Lemma 2.1, R may be extended to a total preorder S satisfying:
(4) $xSy, ySx \Rightarrow xR^*y$.
Let Z be the totally ordered set of equivalence classes of $\mathcal{Y}\times\mathcal{Y}$ defined by the total preorder S. The function d sends a pair of subsets A, B to its equivalence class under S. We shall define $d(a,b)$ as $d(\{a\},\{b\})$. Notice that we have first defined a pseudo-distance between subsets of U, and then a pseudo-distance between elements of U. It is only the pseudo-distance between elements that is required by the definition of representability. The pseudo-distance between subsets just defined must be used with caution because it does not satisfy the property: $d(A,B) = min\{d(a,b) : a \in A, b \in B\}$. It satisfies half of it, as stated in Lemma 2.7 below.
Clearly, $(A,B)R(A',B')$ implies $d(A,B) \leq d(A',B')$. Equation (4) also implies that if $d(A,B) = d(A',B')$, then $(A,B)R^*(A',B')$.
The following argument prepares respect of identity. Suppose that $\mid$ satisfies ($\mid$ 2) and ($\mid A4$) too, and that R was defined including case (3) of Definition 2.3. Defining $0 := d(A,A)$ for any $A \in \mathcal{Y}$, we see that (a) 0 is well-defined: By definition, $(A,A)R(B,B)$ for any $A,B \in \mathcal{Y}$. (b) there is no $d(B,C) < 0$: By definition again, $(A,A)R(B,C)$. (c) $d(A,B) = 0$ iff $A\cap B \neq \emptyset$: $A\cap B \neq \emptyset$ implies $(A,B)R(A,A)$, so $d(A,B) = 0$. $d(A,B) = 0$ implies $(A,B)S(A,A)S(A,B)$, so $(A,B)R^*(A,A)$, so $A\cap B \neq \emptyset$ by ($\mid A4$).
The next lemma shows that our pseudo-distance d behaves nicely as far as its second argument is concerned.

Lemma 2.7
For any A, B $d(A,B) = min\{d(A,b) : b \in B\}$
and
(5) $A \mid B = \{b \in B \mid d(A,b) = d(A,B)\}$.

Proof:
(Remember the elements of $\mathcal{Y}$ are non-empty.) Suppose $b \in B$. Since $(A \mid B\cup\{b\})\cap B \neq \emptyset$ by condition ($\mid$ 1) of Proposition 2.5, $(A,B)R(A,b)$ by case (1) of Definition 2.3, and therefore $d(A,B) \leq min\{d(A,b) : b \in B\}$. If $b \in A \mid B$, then $(A \mid B) \cap \{b\} \neq \emptyset$ and, by Definition 2.3, case (1), $(A,b)R(A,B)$ and therefore $d(A,b) = d(A,B)$. We have shown that the left hand side of Equation (5) is a subset of the right hand side. Since $A \mid B$ is not empty there is a $b \in A \mid B$ and, by the previous remark, $d(A,B) = d(A,b)$ and therefore we conclude that $d(A,B) = min\{d(A,b) : b \in B\}$.
To see that the right hand side of Equation (5) is a subset of the left hand side, notice that $d(A,B) = d(A,b)$ implies $(A,b)R^*(A,B)$ and therefore, by condition ($\mid A2$) of Proposition 2.5, $A \mid b \subseteq A \mid B$ and $b \in A \mid B$. $\square$

To conclude the proof of (a), we must show that Equation (1) of Definition 1.6 holds. Suppose, first, that $b \in B$, $a \in A$ and $d(a,b) \leq d(a',b')$ for any $a' \in A$, $b' \in B$. By Lemma 2.7, $b \in a \mid B$ and $d(a,B) \leq d(a',B)$, for any $a' \in A$.
We want to show now that $b \in A \mid B$. We will show that, for any $a' \in A$, $b \in \{a,a'\} \mid B$. One, then, concludes that $b \in A \mid B$ by Lemma 2.6, remembering that U is finite. Since $b \in a \mid B$, we may, without loss of generality, assume that $a \mid B \neq \{a,a'\} \mid B$. By case (2) of Definition 2.3, $d(a',B) \leq d(a,B)$. But we already noticed that $d(a,B) \leq d(a',B)$. We can therefore conclude that $d(a,B) = d(a',B)$, so $(a,B)R^*(a',B)$, $a \mid B \subseteq \{a,a'\} \mid B$ and finally that $b \in \{a,a'\} \mid B$. We have shown that the right hand side of Equation (1) is a subset of the left hand side.
We proceed to show that the left hand side of Equation (1) is a subset of its right hand side.
Suppose that $b \in A \mid B$. By condition ($\mid$ 1) of Proposition 2.5, $b \in B$. We want to show that there exists an $a \in A$ such that $d(a,b) \leq d(a',b')$ for any $a' \in A$, $b' \in B$. Since the set U is finite, it is enough to prove that, changing the order of the quantifiers:
(6) $\forall a' \in A, b' \in B, \exists a \in A$ such that $d(a,b) \leq d(a',b')$.
Indeed, if Equation (6) holds, we get some $a \in A$ for every pair a', b', and we may take the a for which $d(a,b)$ is minimal: it satisfies the required condition. Since $A = \bigcup\{\{a',x\} : x \in A\}$ (the right-hand side is a finite union) and $b \in A \mid B$, by condition ($\mid A1$) of Proposition 2.5, there is some $x \in A$ such that $b \in \{a',x\} \mid B$. We distinguish two cases. First, if $b \in a' \mid B$, by Lemma 2.7, $d(a',b) \leq d(a',b')$ and we may take $a = a'$. Second, suppose that $b \notin a' \mid B$. We notice that, since $b \in \{a',x\} \mid B$, condition ($\mid A1$) of Proposition 2.5 implies that $b \in x \mid B$. But $b \notin a' \mid B$ also implies that $\{a',x\} \mid B \neq a' \mid B$. By Definition 2.3, case (2), $(x,B)R(a',B)$ and $d(x,B) \leq d(a',B)$. But, by Lemma 2.7, we have $d(x,b) \leq d(x,B)$ (since $b \in x \mid B$) and $d(a',B) \leq d(a',b')$. We conclude that $d(x,b) \leq d(a',b')$, and we can take $a = x$. This concludes the proof of (a).
It remains to show the rest of (b), respect of identity, i.e. that $A \cap B \neq \emptyset$ implies $A \mid B = A \cap B$, under the stronger prerequisites. Let $A \cap B \neq \emptyset$. Then for $b \in B$ $d(A,B) = 0 = d(A,b)$ iff $b \in A$. So by Equation (5) $A \mid B = A \cap B$. □

3 The logical representation results

3.1 Introduction

We turn to (propositional) logic.

Definition 3.1
A pseudo-distance d between models is called definability preserving iff $\mid_d$ is.

d is called consistency preserving iff $M(T) \mid_d M(T') \neq \emptyset$ for consistent T, T'.

The role of definability preservation in the context of preferential models is discussed e.g. in [Sch92], [ALS98] discusses a similar problem in the revision of preferential databases, and its solution. This solution requires much more complicated conditions, for this reason, we have not adopted it here.
Note that $\models T \leftrightarrow Th(M(T))$, and $T = Th(M(T))$ if T is deductively closed. Moreover, $X = M(Th(X))$ if there is some T s.t. $X = M(T)$, so if the operation $\mid$ is definability preserving, and $T * T' = Th(M(T) \mid M(T'))$, then $M(T * T') = M(T) \mid M(T')$.

3.2 The symmetric case

We consider the following conditions for a revision function $*$ defined for arbitrary consistent theories on both sides.
$(*0)$ If $\models T \leftrightarrow S$, $\models T' \leftrightarrow S'$, then $T * T' = S * S'$,
$(*1)$ $T * T'$ is a consistent, deductively closed theory,
$(*2)$ $T' \subseteq T * T'$,
$(*3)$ If $T \cup T'$ is consistent, then $T * T' = Cn(T \cup T')$,
$(*S1)$ $Con(T_0, T_1 * (T_0 \vee T_2)), Con(T_1, T_2 * (T_1 \vee T_3)), Con(T_2, T_3 * (T_2 \vee T_4)) \ldots Con(T_{k-1}, T_k * (T_{k-1} \vee T_0))$ imply $Con(T_1, T_0 * (T_k \vee T_1))$.

Note: $(*4)$ of Definition 1.2 is for free, i.e. a consequence of $(*S1)$ and the other conditions.

The following Example 3.1, very similar to Example 1.1, shows that, in general, a revision operation defined on models via a pseudo-distance by $T * T' := Th(M(T) \mid_d M(T'))$ will not satisfy $(*S1)$, unless we require $\mid_d$ to preserve definability.

Example 3.1
Consider an infinite propositional language $\mathcal{L}$. We reinterpret the models of Example 1.1 as follows:
Let T, T_1, T_2 be complete (consistent) theories, T' a theory with infinitely many models, T, T', T_2 pairwise inconsistent. Let $M(T) := \{m\}$, $M(T_1) := \{m_1\}$, $M(T_2) := \{m_2\}$, $M(T') = X \cup \{m_1\}$ and $Th(X) = T'$. (See Figure 1.1.)
Then $M(T) \mid M(T') = X$, but $T * T' := Th(X) = T'$. We easily verify $Con(T, T_2 * (T \vee T))$, $Con(T_2, T * (T_2 \vee T_1))$, $Con(T, T_1 * (T \vee T))$, $Con(T_1, T * (T_1 \vee T'))$, $Con(T, T' * (T \vee T))$, and conclude by Loop (i.e. $(*S1)$) $Con(T_2, T * (T' \vee T_2))$, which is wrong. $\Box$

We finally have

Proposition 3.1

Let $\mathcal{L}$ be a propositional language.

(a) A revision operation $*$ is representable by a symmetric consistency and definability preserving pseudo-distance iff $*$ satisfies $(*0)-(*2)$, $(*S1)$.

(b) A revision operation $*$ is representable by a symmetric consistency and definability preserving, identity-respecting pseudo-distance iff $*$ satisfies $(*0)-(*3)$, $(*S1)$.

Proof:

We prove (a) and (b) together.

For the first direction, let $\mathcal{Y} := \{M(T) : T$ a consistent $\mathcal{L}$ -theory$\}$, and define $M(T) \mid M(T') := M(T * T')$.

By $(*0)$, this is well-defined, $\mid$ is obviously definability preserving, and by $(*1)$, $M(T) \mid M(T') \in \mathcal{Y}$.

We show the properties of Proposition 2.2. $(\mid 1)$ holds by $(*2)$, if $(*3)$ holds, so will $(\mid 2)$. $(\mid S1)$ holds by $(*S1)$: E.g. $(M(T_1) \mid (M(T_0) \cup M(T_2))) \cap M(T_0) \neq \emptyset$ iff $(M(T_1) \mid M(T_0 \vee T_2)) \cap M(T_0) \neq \emptyset$ iff (by definition of $\mid$) $M(T_1 * (T_0 \vee T_2)) \cap M(T_0) \neq \emptyset$ iff $Con(T_1 * (T_0 \vee T_2), T_0)$. By Proposition 2.2, $\mid$ can be represented by an - if $(\mid 2)$ holds, identity respecting - symmetric pseudo-distance d, so $M(T * T') = M(T) \mid M(T') = M(T) \mid_d M(T')$, and $Th(M(T * T')) = Th(M(T) \mid_d M(T'))$. As $T * T'$ is deductively closed, $T * T' = Th(M(T * T'))$.

Conversely, define $T * T' := Th(M(T) \mid_d M(T'))$. We use Proposition 2.2. $(*0)$ and $(*1)$ will trivially hold. By $(\mid 1)$, $(*2)$ holds, if $(\mid 2)$ holds, so will $(*3)$. As above, we see that $(*S1)$ holds by $(\mid S1)$, where now $(M(T_1) \mid_d M(T_0 \vee T_2)) \cap M(T_0) \neq \emptyset$ iff $M(T_1 * (T_0 \vee T_2)) \cap M(T_0) \neq \emptyset$ by definability preservation. $\square$

3.3 The not necessarily symmetric case

Recall that we work here with a language defined by finitely many propositional variables. For the not necessarily symmetric case, we consider the following conditions for a revision function $*$ defined for arbitrary consistent theories on both sides.

$(*0)$ If $\models T \leftrightarrow S$, $\models T' \leftrightarrow S'$, then $T * T' = S * S'$,

$(*1)$ $T * T'$ is a consistent, deductively closed theory,

$(*2)$ $T' \subseteq T * T'$,

$(*3)$ If $T \cup T'$ is consistent, then $T * T' = Cn(T \cup T')$,

$(*A1)$ $(S \vee S') * T \vdash (S * T) \vee (S' * T)$,

$(*A2)$ If $(S, T)R^*(S, T')$, then $S * T \vdash S * (T \vee T')$,

$(*A3)$ If $(S, T)R^*(S', T)$, then $S * T \vdash (S \vee S') * T$,

$(*A4)$ If $(S, T)R^*(S', T')$ and $Con(S', T')$, then $Con(S, T)$.

Where the relation R is defined by

(1) If $Con(S * (T \vee T'), T)$, then $(S,T)R(S,T')$,
(2) If $(S \vee S') * T \neq S' * T$, then $(S,T)R(S',T)$,
and, in the identity-respecting case, in addition by
(3) If $Con(S,T)$, then $(S,T)R(S',T')$.
Note that by finiteness, any pseudo-distance is automatically definability preserving. We have

Proposition 3.2
Let $\mathcal{L}$ be a finite propositional language.
(a) A revision operation $*$ is representable by a consistency preserving pseudo-distance iff $*$ satisfies $(*0)-(*2)$, $(*A1)-(*A3)$, where the relation R is defined from the first two cases.
(b) A revision operation $*$ is representable by a consistency preserving, identity-respecting pseudo-distance iff $*$ satisfies $(*0)-(*3)$, $(*A1)-(*A4)$, where the relation R is defined from all three cases.

Proof:
We show (a) and (b) together.
We first note: If $T * T' = Th(M(T) \mid M(T'))$, then by definability preservation in the finite case $M(T * T') = M(T) \mid M(T')$, so $(M(S) \mid M(T) \cup M(T')) \cap M(T) \neq \emptyset \Leftrightarrow Con(S*(T \vee T'), T)$ and $(M(S) \cup M(S')) \mid M(T) \neq M(S') \mid M(T) \Leftrightarrow (S \vee S') * T \neq S' * T$. Thus, the relation R defined in Definition 2.3 between sets of models, and the relation R as just defined between theories correspond.
For the first direction, let $\mathcal{Y} := \{M(T) : T$ a consistent $\mathcal{L}$ -theory$\}$, and define $M(T) \mid M(T') := M(T * T')$.
By $(*0)$, this is well-defined, and by $(*1)$, $M(T) \mid M(T') \in \mathcal{Y}$.
We show the properties of Proposition 2.5. $(\mid 1)$ holds by $(*2)$. $(\mid A1)$: We show $(M(S) \cup M(S')) \mid M(T) \subseteq (M(S) \mid M(T)) \cup (M(S') \mid M(T))$. By $(*A1)$, $(S \vee S') * T \vdash (S * T) \vee (S' * T)$, so $(M(S) \cup M(S')) \mid M(T) = M(S \vee S') \mid M(T) = M((S \vee S') * T) \subseteq M(S * T) \cup M(S' * T)$ $= (M(S) \mid M(T)) \cup (M(S') \mid M(T))$. For $(\mid A2)$: Let $(M(S), M(T)) R^* (M(S), M(T'))$, so by the correspondence between the relation R between sets of models, and the relation R between theories, $(S,T)R^*(S,T')$, so by $(*A2)$ $S * T \vdash S * (T \vee T')$, so $M(S) \mid M(T) \subseteq M(S) \mid (M(T) \cup M(T'))$. $(\mid A3)$: Similar, using $(*A3)$. If $*$ satisfies $(*3)$ and $(*A4)$ and R is also generated by case (3), then $(\mid 2)$ and $(\mid A4)$ will hold by similar arguments.
Thus, by Proposition 2.5, there is an (identity-respecting) pseudo-distance d representing $\mid$, $M(T*T') = M(T) \mid_d M(T')$ holds, so by deductive closure of $T*T'$, $T*T' = Th(M(T) \mid_d M(T'))$.
Conversely, define $T * T' := Th(M(T) \mid_d M(T'))$. We use again Proposition 2.5. $(*0)$ and $(*1)$ will trivially hold. The proof of the other properties closely follows the proof in the first direction. $\Box$

4 Conclusion

We proposed a pseudo-distance semantics for the AGM theory of revision. Our semantics is in line with AGM's identification of epistemic states with belief sets. It validates additional postulates that ensure coherence conditions concerning the dependence of the revision operator $*$ on its left argument. Those postulates have been exactly characterized by representation results in the case of symmetric pseudo-distances and only in the finite case for general pseudo-distances.
The question of a representation theorem for the infinite case of general pseudo-distances with conditions in our style stays open. As our main results are purely algebraic in nature, one can be quite confident that important parts of our constructions can be used in the richer situation, in which epistemic states contain more than belief sets.

5 Acknowledgement

We would like to thank a referee for his very insightful comments, which have helped to improve the paper.

References

[AGM85] C. Alchourrón, P. Gärdenfors and D. Makinson, "On the Logic of Theory Change: partial meet contraction and revision functions", Journal of Symbolic Logic, Vol. 50, pp. 510-530, 1985.

[ALS98] L. Audibert, C. Lhoussaine and K. Schlechta: "Distance based revision of preferential logics", Belief Revision Workshop of KR98, Trento, Italy, 1998 (electronic proceedings).

[Bec95] V. Becher: "Unified Semantics for Revision and Update, or the Theory of Lazy Update", 24 Jornadas Argentinas de Informatica e Investigación Operativa, Buenos Aires, pp. 641-650, SADIO, 1995.

[BG93] C. Boutilier and M. Goldszmidt, "Revision by Conditional Beliefs", in: "Proceedings of the 11th National Conference on Artificial Intelligence (AAAI)", pp. 649-654, Morgan Kaufman, 1993.

[BGHPSW97] D. Bellot, C. Godefroid, P. Han, J. P. Prost, K. Schlechta and E. Wurbel: "A semantical approach to the concept of screened revision", Theoria, Vol. 63, Part 1-2, pp. 24-33, 1997.

[BLS99] S. Berger, D. Lehmann and K. Schlechta, "Preferred History Semantics for Iterated Updates", to appear in Journal of Logic and Computation.

[Bor85] A. Borgida, "Language features for flexible handling of exceptions in information systems", ACM Transactions on Database Systems, Vol. 10, pp. 563-603, 1985.

[Bou93] C. Boutilier, "Revision Sequences and Nested Conditionals", in: "Proceedings of the 13th I.J.C.A.I.", R.Bajcsy ed., pp 519-525, Morgan Kaufman 1993.

[DP94] A. Darwiche and J. Pearl, "On the Logic of Iterated Belief Revision", in: "Proceedings of the fifth Conference on Theoretical Aspects of Reasoning about Knowledge", R.Fagin ed., pp. 5-23, Morgan Kaufman, Pacific Grove, CA, 1994.

[DP97] A. Darwiche and J. Pearl, "On the Logic of Iterated Belief Revision", Journal of Artificial Intelligence, Vol. 89, No. 1-2, pp. 1-29, 1997.

[Dal88] M. Dalal, "Investigations into a theory of knowledge base revisions: Preliminary report", Proceedings of the 7th NCAI, pp. 475-479, 1988.

[FH96] N. Friedman and J. Y. Halpern, "Belief Revision: A Critique", in Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning, KR'96, L.C.Aiello and J.Doyle and S.Shapiro eds., pp.421-431, Morgan Kaufmann, Cambridge, Mass., 1996.

[FL94] M. Freund and D. Lehmann, "Belief Revision and Rational Inference", TR 94-16, The Leibniz Center for Research in Computer Science, Institute of Computer Science,Hebrew University, 1994.

[Fre93] M. Freund, "Injective Models and Disjunctive Relations", Journal of Logic and Computation, Vol. 3, No. 3, pp. 231-247, 1993.

[GM88] P. Gärdenfors and D. Makinson, "Revisions of Knowledge Systems using Epistemic Entrenchment", in Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge, M.Y.Vardi ed., pp. 83-95, Morgan Kaufmann, Monterey, California, 1988.

[GM94] P. Gärdenfors and D.Makinson, "Nonmonotonic Inference based on Expectations", Artificial Intelligence, Vol. 65, No. 1, pp. 197-245, 1994.

[Gro88] A. Grove, "Two Modellings for Theory Change", Journal of Philosophical Logic, vol. 17, pp. 157-170, 1988.

[Gr88] P. Gärdenfors, "Knowledge in Flux", MIT Press, 1988.

[KLM90] S. Kraus, D. Lehmann, M. Magidor, "Nonmonotonic Reasoning, Preferential Models and Cumulative Logics", Artificial Intelligence, Vol. 44, No. 1-2, pp. 167-207, 1990

[KM92] H. Katsuno and A. O. Mendelzon, "On the Difference between Updating a Knowledge Base and Revising it", in "Belief Revision", P.Gärdenfors ed., Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, pp. 183–203, 1992.

[LM92] D. Lehmann and M. Magidor, "What does a Conditional Knowledge Base Entail?", Artificial Intelligence, Vol. 55, No. 1, pp. 1-60, 1992.

[Leh95] D. Lehmann, "Belief Revision, revised", in "Proceedings of 14th IJCAI", pp. 1534–1541, Morgan Kaufmann, 1995.

[Lew73] D. Lewis: "Counterfactuals", Blackwell, Oxford, 1973.

[NFPS95] A. Nayak, N. Y. Foo, M. Pagnucco and A. Sattar, "Changing Conditional Belief Unconditionally", in: "Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge", Y.Shoham ed., pp. 119-135, Morgan Kaufman, 1996.

[SLM96] K. Schlechta, D. Lehmann and M. Magidor: "Distance Semantics for Belief Revision", in Proceedings of: Theoretical Aspects of Rationality and Knowledge, Tark VI, 1996, ed. Y.Shoham, Morgan Kaufmann, San Francisco, 1996, p. 137-145.

[Sch92] K. Schlechta: "Some results on classical preferential models", Journal of Logic and Computation, Vol. 2, No. 6, pp. 675-686, 1992.

[Sch97] K. Schlechta: "New techniques and completeness results for preferential structures", to appear in Journal of Symbolic Logic.

[Sch98] K. Schlechta: "Unrestricted preferential structures", submitted.

[Wil94] M.-A. Williams, "Transmutations of Knowledge Systems", in: "Proceedings of the fourth International Conference on Principles of Knowledge Representation and Reasoning", J.Doyle, E.Sandewall, P.Torasso, eds., pp. 619-629, Morgan Kaufman, 1994.