Unified Semantics for Revision and Update or the Theory of Lazy Update

Verónica Becher*
e-mail: postmaster@becher.uba.ar
Universidad de Buenos Aires, FCEyN
Depto. de Computación
Ciudad Universitaria
(1428) Buenos Aires, ARGENTINA

March 1995

Abstract

This paper deals with the connection between two formalisms for theory change: the AGM model of revision and the Katsuno-Mendelzon model of update. Revision and update have been considered orthogonal operations serving different purposes. However, there have been no attempts of recasting one function in terms of the other one, and this is what this work addresses. We semantically recast the AGM revision function as a special case of update. The strategy is to define a global notion of distance (from the theory under update) in terms of the indexical orderings of comparative similarity. We show that the change-operation entailed by this global distance is precisely an AGM revision function. This result allows for conclusions about the two fundamentally different forms of theory change.

Keywords: theory change, belief revision, belief update,nonmonotonic reasoning

1 Introduction

It has been recognized that there are two fundamentally different ways of changing knowledge. The change that results of regarding the new information as contributing to our knowledge or correcting our theory of the world when regarded as static has been dubbed *revision*. We take the AGM theory [6] as the representative model for revision. In contrast, Katsuno and Mendelzon have considered the change in our knowledge in response to an evolving world, referring to it as an *update* [2]. Revision dictates that in case of

*Much of this study has been performed during my visit to The University of British Columbia in January 1995.

contradiction as little as possible of the old theory should be changed in order to accommodate the new facts; it is in this sense that change should be minimal. Usually we have an incomplete theory of the world; namely, our knowledge is incomplete. When doing update one should consider each possible complete description of the world separately and take as the result of the update operation each of the possible outcomes. Each separate "update" determines a set of possibilities that are most similar to the particular description under update but contain the new evidence. It is in this sense that the update operation involves maximal similarity or minimal difference. Let us illustrate the two different forms of change with the following example. Suppose I'm assigned a new office, for which I ignore whether there are ants nor have been ants. I was told that: *There are ants if and only if there have been ants.*

Suppose that I move now to my new office and encounter an ant. What should I conclude? I could accept the evidence as confirming that, effectively, there have been ants and there are ants in my new office. The new fact allowed me to expand my knowledge. This is what *revision* sanctions for this example.

A different way to look at it, is to perform a case-analysis over what we know. I knew that *there are ants if and only there have been ants.* There are two possibilities: either (1) there used to be ants, and there are ants, or (2) there used to be no ants and there are no ants. Suppose (1). In this context, finding an ant is perfectly reasonable. Now suppose (2). Since there must be no ants, my finding of an ant is conflicting. The evidence could be telling that ants *came in* just before me, say, due to the leftovers of yesterday's party. We conclude that definitely there *are* ants now in my office but nothing could be said about ants in the past. This is the type of change dictated by *update*.

In this study we recast the revision function in terms of the semantic apparatus of update and shed light into the relation between them. Our strategy is to define a new operation, a variant of the standard update, that we call *lazy update*, and show that it is a *revision* operation. To calculate the result of a standard update one should perform a case analysis, and take as the result the outcomes sanctioned by all the cases. The central idea behind a lazy update is the ability to compare among the possible outcomes, with respect to their closeness to the theory under update. A lazy update selects just the *least distant* outcomes.

So far researchers who studied the relationship between revision and update have considered them as orthogonal operations serving different purposes. Examples of these studies are [2], [5], [7]. A different view is that of Boutilier in [3] who provides unifying semantics for revision and update in a framework that combines them. Boutilier defines a model for his *generalized update* operation that incorporates observations, events and beliefs. However, there have been no attempts of recasting one operation in terms of the other one, and this is precisely what this work addresses.

This paper is organized as follows. In section 3 and 4 we briefly present the AGM model for revision and the KM model for update. Section 5 reveals the connection between the two formalisms in semantic terms. Section 6 gives some concluding remarks and proposes future research.

2 Logical Preliminaries

We will consider a finitary classical propositional language L and denote with $\mathbf{P}$ the set of all propositional letters. We take Cn to be the consequence operation. The classical connectives will be denoted by $\wedge, \vee, \supset, \neg$. Capital letters $A, B, C, \ldots$ will be used to denote arbitrary formulae. The identically true and false propositions are denoted with $\top, \bot$, respectively.

An interpretation of L is a function from $\mathbf{P}$ to {true, false}; namely, an assignment of truth values to the propositional variables. In this note we may use the terms world, valuation, maximal consistent set of L, and interpretation interchangeably. We take W as the set of all interpretations of L. Given $A \in L$ we denote by $\|A\|$, the proposition for A or set of A-worlds, the elements of W satisfying A.

A preorder $\leq$ over W is a reflexive and transitive relation on W. A preorder is *total* if for every $v, w \in W$, either $w \leq w$ or $w \leq v$. We take $w < v$ as $w \leq v$ but $v \not\leq w$.

3 Belief Revision

The AGM model of revision assumes an ideal agent who possesses a deductively closed *belief set* K, a set of sentences of some logical language, reflecting the agent's beliefs about the current state of the world. In case the agent has to incorporate information A to his state of belief, K has to be *revised* by A. If A is consistent with K, the operation is simply the addition of A to K, dubbed *expansion*: $K^+_A = Cn(K \cup \{A\})$. However, if A conflicts with K, in order to avoid contradiction and yet incorporate A, some beliefs from K must be given up before A is adopted. The AGM theory calls this operation the *revision* of K by A: K^*_A, such that $*$ is a function defined over $\mathcal{K} \times L \rightarrow \mathcal{K}$, where $\mathcal{K}$ is the set of all belief sets. The following eight postulates constrain what a revision function can be.

(K*1) K^*_A is a belief set.

(K*2) $A \in K^*_A$

(K*3) $K^*_A \subseteq K^+_A$

(K*4) If $\neg A \notin K$ then $K^+_A \subseteq K^*_A$

(K*5) $K^*_A = Cn(\bot)$ iff $Cn(\neg A) = Cn(\emptyset)$

(K*6) If $Cn(A) = Cn(B)$ then $K^*_A = K^*_B$

(K*7) $K^*_{A \wedge B} \subseteq (K^*_A)^+_B$

(K*8) If $B \notin K^*_A$ then $(K^*_A)^+_B \subseteq K^*_{A \wedge B}$

As A. Grove in [4] has showed, the AGM theory of revision can be modeled semantically using a total ordering (preorder) on possible worlds reflecting their relative plausibility (or closeness to theory K). In the same spirit as Grove's work, Katsuno and Mendelzon in [2] describe the semantics for revision as follows. They pose a knowledge base as a finite cover of the deductively closed set, and consider a function that assigns to each propositional formula φ a preorder $\leq_\varphi$ over W. This assignment is *faithful* if the following three conditions hold:[1]

1. For any $v, w \in \|\varphi\|$ then $v \leq_\varphi w$.
2. If $w \in \|\varphi\|$ but $v \notin \|\varphi\|$ then $w \leq_\varphi v$ holds.
3. If $\varphi \equiv \mu$ then $\leq_\varphi = \leq_\mu$.

The first constraint says that all φ-worlds are equally ranked, while the second says that φ-worlds are ranked as more plausible than non-φ-worlds. The third constraint asserts that the assignment $\leq_\varphi$ is syntax irrelevant. KM show that $*$ is a revision function if and only if there is a faithful assignment that maps each knowledge base φ to a *total* preorder $\leq_\varphi$ such that

$$\|\varphi^*_A\| = \{\min(\|A\|, \leq_\varphi)\}$$

The function min returns the set of all A-minimal elements v with respect to $\leq_\varphi$, where v is minimal in $V \subseteq W$ if $v \in V$ and there is no $u \in V$ such that $u <_\varphi v$.

Instead of the qualitative ranking relation one can adopt the presentation of W. Spohn's ordinal conditional functions [8]. This richer construction will be of use later when defining the connection between revision and update, since it naturally introduces a metric. An ordinal conditional function k is a function from W, a given set of possible worlds into the class of ordinals, such that some possible worlds are assigned the smallest ordinal 0. k reflects a plausibility ranking of the possible worlds. The worlds that are assigned the smallest ordinals are the most plausible. If $k(w) < k(v)$ then w is more plausible than v or "more consistent" with the current state of belief.

The plausibility ranking of worlds can be naturally extended to an ordering of propositions (sets of possible worlds), by requiring that the ordinal assigned to proposition X be the smallest ordinal assigned to the worlds included in X. That is,

$$k(X) = min\{k(w) : w \in X\}$$

Hence, for all propositions $X \subseteq W$, either $k(X) = 0$ or $k(\neg X) = 0$. Also, for all non-empty propositions $X, Y \subseteq W$, $k(X \cup Y) = min\{k(X), k(Y)\}$. The proposition denoted by the most plausible beliefs is:

$$\|K\| = \{w : k(w) = 0\}$$

It is clear how the function k induces a revision function $*$. To revise K by A means to adopt the most plausible beliefs in A as dictated by function k.

$$\|K^*_A\| = \{w \in \|A\| : k(w) = k(\|A\|)\}$$

[1] In general there is yet another condition required. It is the *smoothness* condition or, as D. Lewis [1] calls it, the *Limit Assumption*. It requires that every subset of W (proposition) to have some minimum in $\leq$. This condition is always satisfied if the set of worlds is finite, or, if $\leq$ is well founded (no infinite descending chains).

4 Belief Update

As briefly discussed in the Introduction, Katsuno and Mendelzon (KM) proposed a change-operation that mirrors the change in knowledge in a dynamic setting. They called this operation an *update*. KM represent a knowledge base by a propositional formula φ, as a finite cover of the deductively closed set. $\varphi o A$ denotes the result of *updating* the knowledge base φ with the sentence A. The eight postulates governing the update operation are:

(U1) $\varphi o A$ implies A

(U2) If φ implies A then $\varphi o A$ is equivalent to φ.

(U3) If both φ and A are satisfiable then $\varphi o A$ is also satisfiable.

(U4) If $\varphi \equiv \mu$ and $A \equiv B$ then $\varphi o A \equiv \mu o B$.

(U5) $(\varphi o A) \wedge B$ implies $\varphi o(A \wedge B)$

(U6) If $\varphi o A$ implies B and $\varphi o B$ implies A then $\varphi o A \equiv \varphi o B$

(U7) If φ is complete then $(\varphi o A) \wedge (\varphi o B)$ implies $\varphi o(A \vee B)$

(U8) $(\varphi \vee \mu) o A \equiv (\varphi o A) \vee (\mu o A)$

KM remark that the update operator behaves monotonically, in the sense that if a knowledge base φ implies μ (meaning that φ contains at least the same information as μ does), then $\varphi o A$ implies $\mu o A$ (meaning that the update of φ will be a new theory at least as specific as the update of theory μ). Monotonicity is an important difference between revision and update. The AGM revision operator is non-monotonic in the following sense:

$$\text{If } H \subseteq K \text{ then, not necessarily } H * A \subseteq K * A$$

It is not the case that if one theory is more specific than another, the revision of the first is necessarily more specific than the revision of the other. We illustrate it with an example. Let a theory H denoting that all we know about Tweety is that it is a pet, so that we have no opinion of its flying abilities. After being informed that Tweety does not fly, we shall revise theory H, concluding that Tweety is a cat or a dog. However, suppose we start with a theory K containing more specific knowledge than Tweety just being a pet. Let pet Tweety a normal bird; hence, it flies. If we get to know that in fact Tweety does *not* fly, we should revise our beliefs, concluding that, say, it must have a broken wing. As we can observe, the two resulting revisions disjoint.

KM formalize a notion of closeness between interpretations to semantically characterize the update operation. Instead of associating each knowledge base φ with an ordering (as done in revision), they consider a function that maps each world w to a partial pre-order $\leq_w$. The meaning of this ordering is a measure of closeness or comparative similarity: $v \leq_w u$ if and only if world v is as close to world w as u is. The comparative similarity ordering dictates that the most plausible changes to w resulting in A lead to those A-worlds that are minimal in $\leq_w$. KM require for the assignment to be faithful to satisfy the following *centering* condition, which says that for every w, no world is as similar to w as w itself:

$$\text{If } v \leq_w w \text{ then } v = w$$

The following characterization result holds for the update operation.

Theorem 1 (Katsuno and Mendelzon, [2]) *Let o be an update operator. The following conditions are equivalent:*
(i) The update operator o satisfies Conditions (U1)-(U8)
(ii) There exists a faithful assignment that maps each interpretation w to a partial pre-order $\leq_w$ such that $\|\varphi oA\| = \bigcup_{w\in\|\varphi\|} min(\|A\|, \leq_w)$

KM show that is also possible to design a class of update operators based on *total* pre-orders (so that all elements of W are comparable in terms of their relative closeness) by adding postulate (U9):

(U9) If φ is complete and $(\varphi oA) \wedge B$ is satisfiable then $\varphi o(A \wedge B)$ implies $(\varphi oA) \wedge B$.

A total preorder associated with world w is what D. Lewis calls a sphere centered at w. Total preorders are of central importance for us, since while revealing the connection between revision and update we will define a new operator dubbed *lazy update* based on total preorders. It is clear how a total ordering can be recast in terms of an ordinal conditional function or a ranking: To each world $w \in W$ we associate a total pre-order as an ordinal conditional function k_w, such that all the information encoded in $\leq_w$ is placed in k_w. For instance, if $v \leq_w u$ then $k_w(v) \leq k_w(u)$. The centering condition becomes:

$$k_w(w) = 0 \text{ and for every } v \in W \text{ such that } v \neq w,\ k_w(v) > 0.$$

We obtained an update model based on ordinal conditional functions $M = \langle W, k_i \rangle$, where each k_i satisfies the centering condition. The update of φ by A is defined as:

$$\|\varphi oA\| = \bigcup_{i\in\|\varphi\|} \{w \in \|A\| : k_i(w) \text{is minimal}\}$$

5 Relating Revision and Update

Our aim in this section is to relate revision and update in the same semantic framework. We are to propose that the revision function can be understood in terms of the update semantic apparatus. We plan to define a special kind of update that responds to the following desideratum: Based on the underlying information of comparative similarity, take the evidence as being itself minimally bizarre, or maximally plausible. We are confined to a set of possibilities of what the world looks like. When we observe some evidence we analyze its impact over what we know. Update dictates that to calculate the result of accepting the new information, we must perform a case analysis over our incomplete picture of the world, find out what is the most plausible outcome(s) for each case, and take as the result *all* the possible outcomes. Even though for each case we selected the closest outcome entailing the change, some outcomes could be relatively implausible. Could we have a *measure* to determine when one outcome is *more plausible* than another? What is a sensible notion to compare outcomes? We suggest that one outcome is more plausible than another when it is at a *closer distance* to the knowledge base. Hence, we could define a new operation that picks as a result of the change just the outcomes that are minimally distant. It remains to define a sense to measure distance.

In this section we assume finitely axiomatizable propositional belief sets. For ease of presentation, we will make a notation abuse and denote both, a belief set and the propositional formula associated to it, with letter K. Let us assume the update model based on ordinal conditional functions $M = \langle W, k_i \rangle$. We define the distance from a particular world v to world w as $distance_v(w)$ as the rank number of w in k_v. Since each function k obeys the centering condition the following definition leads to a very natural notion of distance.

Definition 2 $distance_v(w) = k_v(w)$

Notice that this conception of distance is not symmetrical; namely, the distance from w to v may not be the same as the distance from v to w.

Now it becomes possible to compare the distance from world w to its closest A-world along the relation k_w, with respect to the distance from world v to its closest A-world along the relation k_v. However our desideratum is to compare distances from propositions to propositions. Let us first define the distance from a proposition (let's use $\|K\|$) to a particular world w. It is natural to let the distance from a set to a particular point as the minimum of the distances from each element of the set to the point. To be precise we should also give the definition for the limiting case of distance from an empty proposition. Sort of arbitrary, we define the distance from an empty proposition to any world as the smallest ordinal 0.

Definition 3

$$distance_{\|K\|}(w) = \begin{cases} min\{k_v(w) \ for \ \ each v \in \|K\|\} & ,if \ \|K\| \neq \emptyset \\ 0 & ,otherwise \end{cases}$$

This definition naturally generalizes for distances between propositions.

Definition 4 $distance_{\|K\|}(\|H\|) = min\{distance_{\|K\|}(h), for \ each \ h \in \|H\|\}$

As before, notice that in general, $\text{distance}_{\|K\|}(\|H\|)$ is different from $\text{distance}_{\|H\|}(\|K\|)$. Let us recognize that definition 3 induces for each proposition a rank of worlds, according to the distance from the proposition to each single world.

Corollary 5 *Let $\|K\| \subseteq W$. $distance_{\|K\|}$ is an ordinal conditional function, satisfying the centering condition with respect to all the worlds in $\|K\|$.*

Based on the notion of distance from a proposition we have obtained an ordinal conditional function centered in that proposition. This motivates the following definition.

Definition 6 $k_K(w) = distance_{\|K\|}(w)$.

We are ready to define the lazy update function that we will notate with the symbol $\bullet$. The lazy-update of K by A is a new theory $K^{\bullet}_A$ denoted by the A-worlds minimally distant to $\|K\|$.

Definition 7 $\|K_A^\bullet\| = \{w \in \|A\| \subseteq W : \textit{distance}_{\|K\|}(w) = \textit{distance}_{\|K\|}(\|A\|)\}$

equivalently, $\|K_A^\bullet\| = \{w \in \|A\| : k_K(w) \text{is minimal}\}$
The following two theorems relate lazy-update with standard update. The first one asserts that when the knowledge base is complete the two kinds of update coincide.

Theorem 8 *If K is complete then $\|K_A^\bullet\| = \|K_A^o\|$.*

The proof is quite trivial. If K is complete, it is characterized by a unitary proposition $\|K\| = \{u\}$. So, $\|KoA\| = \bigcup_{i\in\|K\|}\{w \in \|A\| : k_i(w) \text{ is minimal}\} = \{w \in \|A\| : k_u(w)\text{is minimal}\} = \{w \in \|A\| : k_K(w)\text{is minimal}\} = \|K_A^\bullet\|$.

The second theorem indicates that a a lazy update yields a more informed theory than standard update. In other words, a lazy update neglects some possibilities -because it regards them as implausible- that the standard update considers. As a result the lazy update is more opinionated than the standard update.

Theorem 9 *If $\|K\| \neq \emptyset$, $\|K_A^\bullet\| \subseteq \|K_A^o\|$.*

The proof is trivial since $\{w \in \|A\| : k_K(w)\text{is minimal}\} \subseteq \bigcup_{i\in\|K\|}\{w \in \|A\| : k_i(w)\text{is minimal}\}$. The reason for this theorem to be conditioned on theory K being consistent is that the standard update function will not reestablish inconsistency (once the theory K becomes inconsistent, no standard update operation will eliminate the inconsistency). In contrast, the lazy update, has been strategically defined to overcome inconsistent belief sets.

The crucial difference between the two forms of update relies in that standard update uses the indexical ranks k_w (the comparative similarity orderings), while the lazy update uses the globalized rank k_K. What have we gained by generalizing the notion of distance between worlds to distance from a proposition to a world? We have naturally spread the indexical rank of comparative closeness (or similarity) obtaining a *global rank* associated with a proposition. As a direct consequence of this "globalization", lazy update ignores some of the possible outcomes that standard update would consider. Lazy update relies on only those possible worlds that regard the change as maximally plausible. The following immediate conclusion is obtained: *If possible, a lazy update will understand a change as having caused no change at all.* The next theorem is called for.

Theorem 10 *If A is consistent with K, then $\|K_A^\bullet\| = (\|K\| \cap \|A\|)$*

The proof is as follows. Assume A is consistent with K. Then $\|K\| \cap \|A\| \neq \emptyset$, so there is some $w \in (\|K\| \cap \|A\|)$. Then for those w in $(\|K\| \cap \|A\|)$, their distance to theory K is minimal, because $\text{distance}_K(w) = 0$ and 0 is the smallest ordinal. Clearly, for any other $v \in \|A\|$ such that $v \notin \|K\|$ $\text{distance}_K(v) > 0$. By definition, $K_A^\bullet = \{w \in \|A\| : k_K(w)\text{is minimal}\}$. Then $K_A^\bullet = (\|K\| \cap \|A\|)$.

The theorem expresses that whenever the change is consistent with the knowledge base the result of lazy-updating it is just to add the evidence to the knowledge base. We interpret this behaviour as no actual change in the real world: the evidence confirms what it was already a possibility in our picture of the world. This theorem is crucial since it closes the gap between revision and update. It becomes not surprising that the lazy update function $\bullet$ satisfies the AGM revision postulates $K*1 - K*8$.

Theorem 11 $\bullet$ *is a revision operator, satisfying* $K*1 - K*8$.

Distance from theory K becomes the standard global plausibility ordering used in revision. A world w is as *plausible* as v with respect to a theory K if and only if the distance from $\|K\|$ to w is not greater than the distance from $\|K\|$ to v.

We have recast the problem of relating revision and update into finding a meaningful *global* binary relation between worlds as a function of the rankings k_w (orderings of comparative similarity $\leq_w$). A candidate global ranking k could have been any arbitrary k_w. But it is evident that the revision function that would induce does not satisfy the AGM revision properties.

6 Concluding Remarks

Along the lines of Grahne's axiomatization of updates [5] we could axiomatize lazy updates in a metric logic. The standard update and the lazy update would be connectives "o" and "$\bullet$" in the object language, both being binary non-truth functional connectives. Theory K updated by A would be written as KoA. If $B \in KoA$, we could write $(KoA) \supset B$ (where$\supset$ is material implication) Similarly, if B belongs to the lazy-update of K by A, we would write $(K \bullet A) \supset B$.

In our presentation of revision as a lazy update in model $M = \langle W, k_i \rangle$, a belief set K is denoted by a proposition which possesses a plausibility ordering k_K as a function of the k_i. The revision of K by A is a new belief set K^*_A, also a proposition in the *same* model M, which has automatically defined its plausibility ordering $k_{K^*_A}$. Whether this resulting plausibility ordering is meaningful requires some thinking. A straight consequence of having the plausibility ordering automatically defined for every proposition is that we can easily allow for iterated revisions. It is clear that the standard update operation accepts iteration [5]. In other words, it is clear how to calculate sequences of updates (module that the underlying relations of comparative similarity remain intact, which is not a minor assumption). However, as it has been extensively analyzed, it is not clear what a revision sequence should entail. Since we have proved that lazy update is a revision operation we may wonder what kind of iterated revision it embraces. Apparently, given our definition of distance as it stands, the notion of iterated revision it yields is *not* intuitive. The problem is that the iteration of lazy-updates does not preserve much of the plausibility rank associated to theory K[2]. A new scheme for "distances", or a way of obtaining a unique total ordering as a combination of the ranks of comparative similarity k_w might be required to obtain a reasonable lazy update for iteration.

C. Boutilier[3] has pointed out that the semantic connection that we found for revision and update relies almost exclusively on the fact that the standard KM update possess a model that is centered. However, this centering assumption is not necessarily very intuitive, since embodies the idea that every world is most similar to itself and no other

[2] A different way to say this is that the conditional assertions that illustrate the underlying plausibility rank associated to theory K are not preserved through a lazy update. The conditionals that seem to be preserved when revising by a sentence A are just those whose antecedent is strictly less plausible than A.

[3] Personal communication

world can be as similar to it. He suggested a sensible variation of the classical update: a weakly centered model; that is, every world can have many worlds most similar to itself but including itself. Under this variation, the lazy update operation is not any more a revision function.

7 Acknowledgments

Thanks to Craig Boutilier and Carlos Alchourrón for their valuable comments on early ideas on this paper. I'm also grateful to Carlos Areces for commenting a previous draft of this work. This research has been supported by a fellowship of the Informatics Program of the University of Buenos Aires.

References

[1] D. Lewis. 1973. *Counterfactuals*, Blackwell, Oxford.

[2] H. Katsuno and A. Mendelzon. 1991. On the Difference between Updating a Knowledge Base and revising it. In *Proceedings of Second International Conference on Principles of Knowledge Representation and reasoning*, pages 387-394, Cambridge.

[3] C. Boutilier. 1994. Generalized Update: Belief Change in Dynamic Settings. ?????

[4] A. Grove. 1988. Two modellings of Theory Change. *Journal of Philosophical Logic*. 17:157:170.

[5] G. Grahne. 1991. Updates and Counterfactuals. In *Proceedings of Second International Conference on Principles of Knowledge Representation and reasoning*, pages 269-276, Cambridge.

[6] . P. Gärdenfors. 1988. *Knowledge in Flux.*MIT Press, Cambridge.

[7] S. Lindström and W. Rabinowicz. The Ramsey Test revisited. In *Conditionals and Artificial intelligence* edited by Crocco, Fariñas de Cerro, Herzig. Toulouse, France.

[8] W. Spohn. 1987. Ordinal Conditional Functions. In Harper and Skyrms, editors, *Causationin Decision, belief Change and Statistics*, D. Reidel, Dordrecht.