

Lindström and Rabinowicz on relational belief revision

Erik J. Olsson

Lund University and University of Constance

Abstract: It is usually taken for granted that a theory of belief revision should describe justified changes from one belief state to another belief state where the output state is uniquely determined given the input state and the new information. This uniqueness assumption has been questioned by Lindström and Rabinowicz whose theory of relational belief revision allows for the result of belief revision to be indeterminate in the sense that there may be many possible end states that are equally rational. The main aim of the paper is to inquire into the possible motives behind this generalization of the standard functional setting.

1. Introduction

It is a great pleasure to contribute to this Festschrift on the occasion of Wlodek's 60th birthday. Wlodek is a first rate philosopher and we Swedes should consider ourselves lucky that he has decided to pursue his career in Sweden and Lund rather than, as one might have expected and feared, under more glamorous circumstances elsewhere. I am personally indebted to Wlodek in many respects. He was a valued teacher during my student years in Uppsala in the 1980's and when I went on to write my PhD thesis I benefited a lot from our discussions. In the 1990's, he was appointed the Chair in Practical Philosophy in Lund, whereas I went to Constance in Germany to work as a post-doc. This meant that I didn't meet Wlodek for several years, except at some conferences that we both attended, including one in Constance that I organized (with Volker Halbach). When I returned to Sweden, in 2003, I was happy to have Wlodek as a colleague in Lund. If asked to mention just one thing I value about Wlodek, as an intellectual, I would point to his unique ability to quickly grasp a new philosophical theory, even an intricate one. Not only will he be able understand the new proposal; he will also immediately either prove or – what would be more often the case – disprove it, usually triumphantly and with an argument that sounds absolutely compelling. Discussing with Wlodek is

therefore both rewarding and intimidating at the same time, rewarding because of the many insights thus gained, intimidating because one becomes painfully aware of one's own philosophical limitations. Even so, as my own thinking has matured I have come to believe, rightly or wrongly, that even Wlodek can be mistaken at times. In this paper, I will point to some claims of his that do not seem entirely correct, or at least not sufficiently underpinned. However, I will also argue, sometimes with the help of other authors, that some of the more important theses are nonetheless true. So, it could still be held that Wlodek is always right "in principle".

Wlodek's philosophical interests are, by Swedish standards, unusually broad, covering a wide range of topics in both practical and theoretical philosophy. In this paper, I will focus exclusively on a small part of his work, namely, the theory of relational belief revision developed by him in collaboration with another much appreciated teacher from my Uppsala days, Sten Lindström. Technically, this theory is a generalization of the well-known AGM theory of belief revision due to Alchourrón, Gärdenfors and Makinson (1985). Whereas in the AGM theory belief revision is seen as a function from a belief state and a sentence to a new belief state, L&R propose to view the revision process as a relation, thus allowing there to be several equally reasonable revisions of a theory with a given proposition.

Much of L&R's work on relational belief revision has been concerned with formal aspects of that kind of revision. For example, they provide an axiom system for relational belief and show how this system can be related to axioms for epistemic entrenchment. They also show that in the special case of functional belief revision the axioms are sufficient to guarantee representability, but that this result cannot be extended to relational belief revision in general. These results and others are important contributions to philosophical logic. For the details, I refer to Rabinowicz and Lindström (1991) and (1994).

In this paper, however, I will not be concerned so much with formal aspects of relational belief revision as with its philosophical foundations. From a logical point of view, it is often interesting in itself to study a generalization of some other formal theory. Yet for a generalization to be of more than purely formal significance, the added generality must provide some advantage beyond generality

itself. To take a trivial example, the concept of a rational number is more general than the concept of a natural number. This observation is obviously insufficient to establish the added value of having rational numbers which rather derives from the fact that, while natural numbers are greater for counting, we need the rational numbers to be able to perform measurements, e.g., of length. So what about allowing belief revision processes that is not functional – does it solve any problems that we could not solve before? In what follows I will assume some knowledge of the standard AGM theory, although most of the discussion will be conducted on a general philosophical level.

2. Relational revision as a solution to the non-uniqueness problem

In their 1989 paper, Lindström and Rabinowicz considered the problem of relating probabilistic and non-probabilistic belief change. In the non-probabilistic model, a belief state of an ideally rational agent is modeled by a belief set, i.e., as a consistent set of sentences of some formal language closed under logical consequence. Belief revision is taken to be a process leading from one such belief set to another. In the probabilistic model, belief states are represented as probability functions and belief revision or, as L&R prefer to say, probability revision is seen as a function from probability functions to probability functions.

One way of relating probability and non-probabilistic models would be to regard probability revision as the more fundamental concept and to see belief revision as derived from probability revision. The most straightforward idea would be to identify the agent's belief set with the "top" of her probability function, where by the top is meant the set of all sentences assigned probability one. As L&R note, however, it is not as clear how to proceed if we choose instead to take the agent's belief set as our starting point. L&R's propose tentatively to define the revision $A * x$ of a belief set A with a proposition x by taking the probability function associated with A , revising it with x and letting $A * x$ be the top of the result. The problem with this suggestion, however, is that it is unclear how to identify "the probability function associated with A ". There are of course many probability functions that have A as their top. If we take any two such functions, there is no

guarantee that revising them by x will yield the same belief set. Hence, the “non-uniqueness problem”.

In their paper, L&R consider a number of possible responses to this problem, none of which they find absolutely convincing. The alternatives that they find clearly unsatisfactory are: (1) imposing a suitable condition on probability revision, (2) letting the revised belief set be the common part of the different possible candidates to this title, (3) introducing an explicit mapping that associates a specific probability function with each belief set. They find two other strategies more plausible. One involves introducing the notion of a belief state as a new primitive and defining both belief and probability revision in terms of a revision operation on belief states. The other, finally, is letting belief revision be a relation rather than a function. The conclusion of the paper is that one of the latter strategies should be pursued, or possibly a combination of both.

Let us consider how the relational approach is supposed to solve the non-uniqueness problem. First of all, what does this approach involve, more precisely? A belief revision relation is defined as a ternary operation $R \subseteq \mathbf{K} \times \text{Con} \times \mathbf{K}$ satisfying a number of axioms. Thus it is assumed that, if $A \mathbf{R}_x B$, then $x \in B$. This means that, if B is a possible result (for a given agent) of revising A by sentence x , then x is an element of B . Those acquainted with the AGM revision axioms will recognize this as the relational equivalent of the success postulate.

According to L&R once we have the relational approach, “[t]here is no longer any problem of defining belief revision from probability revision” (1989). The reason is we need not assume anymore that there should be a unique belief set that is the result of revision A by x . Hence there is no pressure to specify a unique probability function that has A as its top and that can be used for arriving at a unique new belief set by taking the top of that distribution revised by x . Rather, we can allow for their being a number of possible revision results depending on what probability distribution is chosen. Formally,

Definition: $A \mathbf{R}_x B$ iff, for some P in \mathbf{P} , $A = t(P)$ and $B = t(P * x)$.

L&R conclude that, once the relational view is adopted, the non-uniqueness problem “disappears”.

How convincing is this proposal? The non-uniqueness problem would indeed disappear if it could be shown that one of its central presuppositions is false or at least unfounded. Has this been shown with respect to the presupposition that belief revision should yield a unique belief set as output? I cannot see that it has. All that has been shown is that formal sense can be made of belief revision without that presupposition. This is a weak argument for the falsity of that presupposition. It could still very well be the case that belief revision does or should yield a unique result. Those who have advanced a functional approach have not done so, I submit, because they have been convinced that a relational approach would be formally intractable. Rather, they have simply thought that a theory of belief revision should give a *definite* description (or prescription) of the belief state of a rational agent after that agent has received new information. After all, the agent must, unless she suffers from schizophrenia, eventually adopt a single state of belief upon learning something new. The problem, it is generally agreed, is to say what that uniquely determined state of belief is (or should be).

Hence, for those who insist that we need to define non-probabilistic belief revision with the help of probabilistic belief revision without taking a probability function as given, there is still a uniqueness problem. *Pace* L&R, introducing the relational setting does very little, if anything, to make that problem go away. Moreover, one could flatly deny that there is a problem to start with, i.e., that we need to define non-probabilistic in the way envisaged by L&R manner at all. On closer examination, L&R present no detailed case for why this would be an urgent project. In conclusion, the amount of support for the relational framework derivable from the first L&R paper is not very impressive.

3. Relational revision as arising from entrenchment incomparability

In a second paper, presented at a workshop in 1989 and published in a proceedings volume in 1991, L&R provide a more compelling motivation for the relational framework. Their point of departure is Adam Grove’s paper from 1988 where two related models of functional belief revision are presented, one in terms of a family

of spheres around the agent's theory G , viewed as a set of possible worlds, and the other in terms of an epistemic entrenchment ordering of propositions. Grove's spheres may be thought of possible fallback theories relative to the agent's original theory. By a fallback theory is meant a theory that may be reached by deleting propositions that are not sufficiently entrenched. In other words: fallback theories are theories that are closed upwards under entrenchment so that, if T is a fallback, A belongs to T and B is at least as entrenched as A , then B also belongs to T . Figure 1 illustrates Grove's family of spheres around a given theory G .




Figure 1: A theory and its family of spheres.

We notice that the spheres around a theory are nested, i.e., simply ordered. For any two spheres, one is included in the other.

The next picture illustrates how revision is supposed to work in the Grove model.




Figure 2: Revision in the Grove model.

The area labeled H in Figure 2 represents the revision of G with a proposition A. The result of revising G by A is taken to be the strongest A-permitting fallback theory of G expanded by A. This corresponds to the taking intersection of A with the smallest sphere around G that is compatible with A. This clearly gives a unique result. (If A is inconsistent, the revision by A is taken to be the inconsistent theory, i.e. the empty set of worlds.)

But assume now that some propositions may be incomparable with respect to entrenchment. Two propositions are incomparable if neither is at least as entrenched as the other. Hence, allowing for incomparability means relaxing the assumption that the entrenchment ordering is connected. As a result, the family of fallbacks around a given theory no long has to be nested. It will no long be a family of spheres but, to use L&R's term, rather a family of "ellipses". Allowing for incomparability vis-à-vis entrenchment means opening up for the possibility that there may be several different ways to revise a theory with a given proposition. See Figure 3 for an illustration.




Figure 3: Revision in Grove model with incomparability.

In the picture, the two ellipses represent two different fallback theories for G . Each of them is a strongest A -permitting fallback. Hence, both H and K is the intersection of A with a strongest A -permitting fallback. It is natural, therefore, to say that both are possible revisions of G by A .

Still, so far this is merely a hypothetical defense of relational belief revision. What L&R have argued is that a case could be made for relational belief revision *if propositions can plausibly be incomparable with respect to entrenchment*. The question remains as to whether propositions *can* be incomparable.

L&R's view is that they can. We can, they say, be unable to compare propositions "perhaps because the propositions are so different from each other, or perhaps because they are totally unrelated" (1991, p. 106). The vagueness of this short account of the roots of incomparability makes it difficult to assess. In particular, it is unclear what L&R mean propositions being "different". However, their reference to "totally unrelated" propositions does suggest unrelatedness with respect to topic. Still, there are many unrelated propositions that are easily comparable. For instance, I consider my belief that the earth is round much more entrenched than my topically unrelated belief that we will have pork for lunch today. Hence, topic-difference cannot be a *source* of entrenchment incomparability.

On this reading of their proposal, L&R still owe us an explanation of why some cases of topic-difference lead to incomparability and some don't.

4. Other accounts of the roots of incomparability and relationality

Isaac Levi has provided what I take to be a more compelling defense of entrenchment incomparability. Without going into any technical details, incomparability results, says Levi, not because we are comparing propositions that are different content-wise, but “due to conflict or indeterminacy in the agent's values and goals” (Levi, 2004, p. 206). This yields indeterminacy in the sense that the agent's assessment of informational value needs to be represented not as a single measure (a so-called M-measure) but as a (convex) set of such measures. Each such measure, as Levi shows via his concept of damped informational value, gives rise to a permissible entrenchment ordering of the agent's beliefs. Levi notes that each permissible entrenchment ordering yields a nested system of spheres in the sense of L&R, and so “[i]f we consider all unions of the sets of fallbacks associated with each permissible ordering, we have a system of fallbacks of the sort considered by Lindström and Rabinowicz with an associated entrenchment ordering that allows for incomparabilities” (ibid., p. 211).

Agents with conflicting theoretical goals and values may end up in a situation where there is no unique way to order belief with respect to epistemic entrenchment but several equally admissible orderings. This means that a given change in belief may give rise to an indeterminate result: there can be several equally rational ways to change beliefs based on considerations of entrenchment. A relational theory can accommodate the indeterminacy. A functional theory cannot. Hence, the relational approach is, philosophically speaking, a justified generalization after all, or so it would seem.

However, it is one thing to say that belief revision *based on considerations of entrenchment* can yield an indeterminate result; it is quite another to say that belief revision *all things considered* can yield an indeterminate result. It may well be that, when we construct a new belief state by deleting, according to some recipe, beliefs that are not sufficiently entrenched, we don't end up with a unique rationally prescribed belief state. But many researchers would be unhappy with letting this be

the end story. In such cases, they would say, rationality dictates that we invoke a rule for ties as a secondary criterion. To be specific, the result of belief revision all things considered should be a belief state containing all and only the beliefs that are common to all admissible belief states. This new belief state will be unique. Hence, while entrenchment-based belief revision may require a relational representation, belief revision all things considered is still a functional matter.

An advocate of the relational framework would have to reject the proposal that the end result of belief revision should be a compromise between admissible belief states. She would have to dismiss the claim that such a tie-breaking rule is rationally prescribed. In her view, all that rationality can deliver is a set of admissible beliefs states among which the inquirer is ultimately free to choose arbitrarily. This is the philosophical position that the advocate of the relational framework is forced to take. (Actually the relational theorist could also argue that relational belief revision accounts for only the first part of the belief revision process – the part preceding the invocation of a rule for ties. However, this proposal does not rhyme very well with the relational theorist’s typical way of presenting the relational theory as an alternative to the AGM theory which, as we know, purports to account for the whole revision process. I will return to this point at the very end of the paper.)

Yet this is a severely problematic position. In the domain of practical actions, it can sometimes be rationally permissible to choose arbitrarily. If, for instance, the agent takes “eating an apple” and “eating a banana” to be equally attractive options, and better than all other alternatives, he could not be faulted if he decided, eventually, to pick, say, the apple. Rationality does not forbid arbitrary choices among possible practical actions that tie for optimality. The picture changes as we move over to the theoretical domain. Someone finding evolutionary theory and creationism equally promising explanations of why nature is as it is, and better than all other explanations, would be irrational to pick one of the two theories arbitrarily as his definite view on the matter. Rather, all he could definitely believe in is what evolutionary theory and creationism have in common (which, I presume, is not very much). Note that we are here exclusively concerned with full beliefs. For full beliefs, arbitrary choices among options that tie for optimality are rationally

forbidden. Obviously, this claim does not conflict with the fact that a researcher may legitimately make an arbitrary choice among optimal options concerning what *research program* to pursue or what *hypothesis* to test.¹

Here is another proposal for how to justify the relational framework.² The relational model may not be so plausible so long as we confine ourselves to the consideration of small changes of specific beliefs within a system of belief where the system itself does not undergo any dramatic changes. But consider a case of *bone fide* scientific theory change. We have a theory which must be changed in the light of the outcome of one or more experiments. We also have criteria for rational theory choice, such as empirical adequacy, simplicity, fruitfulness and the like. Given all this, there is no guarantee that one single unique theory will satisfy our adequacy criteria optimally. Rather, we should not be surprised to find that several theories tie for optimality. In such cases, moreover, the alternative to choose the common part of all admissible theories does not seem very promising. After all, the intersection of a number of good theories need not itself be a good theory. Hence, modeling real scientific change requires that we allow for indeterminacy.

In response, I agree that our criteria are usually not strong enough to point uniquely at one theory as the one to go for, and that this may leave us with several theories that tie for optimality. But I reject the notion that we are at that point free to make an arbitrary choice of what to believe. We must be careful here, though, about what we are talking about. “Choosing a theory” can mean either choosing to believe in the theory, i.e. accepting it as true, or choosing to pursue it as a plausible research program. Again, arbitrary choices are forbidden on the first but not on the second interpretation.

It is true, also, that the intersection of a number of good theories need not qualify as a good theory.³ So, if we follow the advice to believe in what the best theories have in common, what we may end up with a suboptimal theory. In defense of the intersection strategy, however, I would like to propose that the

¹ Similar criticisms of the relational approach are offered in Levi (2004), p. 213.

² This proposal was made by Sten Lindström in personal communication.

³ Isaac Levi has argued, with much ingenuity, that the meet of two optimal theories is also optimal. For his latest attempt, see Levi (2003). However, I believe – for reasons similar to those put forward by Rott (2006) – that he has been unsuccessful so far and that the prospects of making progress in this direction in the future are severely limited.

intersection theory should not be chosen as the *end-point* of inquiry but as an *intermediate result in an on-going investigation*. Usually, an agent who has come to the conclusion that what the best theories have in common is what can be assumed to be the case at the present stage of inquiry would continue asking which one of those best theories should eventually be accepted. Still on the agenda's agenda, this question would serve to motivate further inquiry and deliberation. To choose a suboptimal theory in this sense is, in my view, not objectionable. Suppose, for instance, that the inquirer starts out by considering the alternative theories T_1 , T_2 and T_3 . Subsequent inquiry and deliberation reveals that T_1 and T_2 are equally good under the circumstances and that they are better than T_3 . In other words, T_1 and T_2 tie for optimality. On the present view, the inquirer may now be justified in adding to her background beliefs everything that T_1 and T_2 have in common, provided that she retains on the agenda the question which one of T_1 or T_2 is ultimately to be chosen. The distinction between the two different senses in which the disjunction of two or more theories can be accepted cannot be made in standard belief revision theory, whether relational or not. The reason is that theories in the AGM-inspired tradition do not take the agent's research agenda seriously in the sense of including it as a part of the formal model. For a sketch of a theory that does, see Olsson and Westlund (2006).⁴

⁴ Another possibility comes to mind for how the relational framework could be defended: From the point of view of group rationality, it may be advantageous that different researchers pursue different lines of research. Intellectual diversity is a vaccine against the group painting itself into a corner by focusing too much collective attention on a limited number of options that may eventually prove to be dead ends. Now researchers are more inclined to pursue their own line of research and work out the details of their theory if they firmly believe that the theory they are working on is true. Moreover, since people have a natural tendency to tell (what they believe to be) the truth this also makes them more inclined to communicate their results. Hence from the standpoint of the group, the optimal setting may be one in which researchers choose their beliefs to some extent arbitrarily in line with the relational approach to belief change. So the argument goes. Against this proposal one could point out that, if a researcher fully believes in the theory she is pursuing, she will be more dogmatic and less inclined to search for possible counterevidence than if she merely entertained the theory as a promising, but possibly false, research hypothesis. (On the other hand, she will be more motivated to find faults in other researchers' competing theories. And, of course, these other researchers will, if they are similarly biased, be quite happy to prove her wrong in turn.)

5. Methodological and systematic benefits of the relational approach

The conclusion so far is that none of the attempts that have been offered in justification of the relational framework is entirely successful. For all we have seen so far, the relational framework does not seem to add anything of substance to our understanding of the belief revision process. However interesting from a formal point of view, the extra generality offered by the relational approach has, it would appear, no philosophical relevance. But this conclusion may be premature. The relational approach has recently been revived by Neil Tennant who claims to have provided it with “new foundations” (Tennant, 2006). Tennant’s rather disappointing starting point is a claim that echoes the view that was criticized in the previous section:

AGM-theory provides an account of expansion, contraction, and revision of theories with respect to sentences. But it does so by treating the ‘operations’ of contraction and revision as though they were functional, with uniquely defined values, for any given rational agent, on all possible inputs $\langle K, A \rangle$. An alternative and arguably more reasonable approach would be to treat contracting and revising as non-deterministic processes that can produce a variety of possible values on any given input $\langle K, A \rangle$. A mark of rationality, on the part of any agent, would be to countenance such variety rather than to insist on uniquely defined outcomes. Hence a theory of relational theory-change should be able to furnish such variety, by treating contraction and revision more generally as relational, not functional, notions (Tennant, 2006, p. 490).

Unfortunately, what Tennant identifies as a “mark of rationality” is actually quite the opposite. If there really are “a variety of possible values on any given input” then the end result of belief revision should be what these possible values have in common. This will yield a uniquely determined belief set.

But this is not Tennant’s only motivation for exploring the relational setting. He also claims that reformulating the AGM theory in relational terms has the “methodological advantage” (p. 493) of helping us to “identify certain inadequacies of AGM-theory that might more easily escape attention in the functional setting” (ibid.). Tennant is here referring to his “degeneration” theorems (see below). The reformulation has the additional systematic benefit of making it possible to derive, rather than stipulate, one of the central AGM principles, namely, that of extensionality. According to the extensionality principle, the result of revising the same belief state with A or with B is the same, if A and B are logically

equivalent. In other words, the result of revision depends only on the logical content of the input sentence and not on its syntactic structure. Tennant argues – correctly as far as I can see – that stipulating rather than deriving the extensionality principle is essential to the functional framework.

Tennant mentions this treatment of extensionality as one of the main differences between his relational framework and that of L&R. In the latter, extensionality identifies as a postulate and not as a derived principle. There are other significant differences as well. Thus Tennant, in his analysis of contraction and revision, focuses on the principal cases. In revision, for instance, the principal case is when the new information is actually inconsistent with the consistent theory that is to be revised. Tennant is right in thinking that in the AGM theory too much attention is paid to extremal cases. This criticism carries over, in some degree, to L&R's relational formulation of AGM. A further point of divergence concerns the formulation of the belief change postulates. Where L&R formulate their system as a set of axioms, Tennant chooses a formulation in terms of introduction and elimination rules similar to those employed in systems of natural deduction. Tennant, finally, makes much out of the fact that L&R's treatment of revision is agent relative, whereas his own goal is to explicate a notion of revision that is agent independent. It is true that L&R make references to an agent in their informal explanations. Still, since the agent does not figure explicitly in the formalism, I don't think there is substantial disagreement on this point. Tennant is nonetheless right in saying that, all in all, his treatment of relational revision "has significantly different aims and methods from the 'relational' treatment of revision given by Lindström and Rabinowicz" (p. 495). Even so, I object to Tennant's insinuation that L&R's approach is not fully relational. While there are, as we just saw, substantial differences between the two approaches, these dissimilarities do not concern the theories' degrees of relationality. L&R's theory is every bit as relational as Tennant's own, although Tennant has a point in saying that, in stipulating rather than deriving the extensionality principle, L&R do not make full use of their relational resources.

Tennant goes on to prove his main results which are a number of "degeneration" theorems for his relational reformulation of the AGM theory.

According to one theorem, the mere logical closure of a new (consistent) input A qualifies as a revision of any theory K inconsistent with A . It does not matter with other beliefs might have been in K . The revision result can be, in an uncomfortable sense, independent of the original belief state. The upshot of Tennant's illuminating relational analysis is that the constraints on rational revision imposed by the AGM postulates are weak enough to allow all kinds of strange "revisions". The postulates look much more substantial than they really are.

6. Conclusion

Studying relational belief revision is a natural generalization of the usual functional approach. Lindström and Rabinowicz should be credited for being the first to make this observation, as well as for investigating, in impressive formal detail, its consequences. The philosophical problem they confront is to make likely that the added generality thus gained can be of use in our attempts to understand the revision process. For reasons already stated, I don't think that L&R have been entirely successful in that particular regard. Moreover, the relational approach, as they present it, seems to assume the legitimacy of arbitrary theoretical choices among options that tie for optimality. There are strong philosophical reasons to reject this assumption. This deficiency is shared by Tennant's relational treatment. Tennant, however, actually succeeds, I think, in showing that the relational setting adds, methodologically as well as systematically, to our understanding of belief revision, at least if the latter is understood roughly in the sense of the AGM theory. His degeneration theorems and treatment of the extensionality "postulate" are cases in point.

What I would like suggest, in conclusion, is that the belief revision process be divided into two stages. In the first stage, we have entrenchment-based revision along the lines drawn by L&R. The result may be, as L&R correctly note, that several potential belief states tie for optimality. This necessitates, as Levi has insisted, that a rule for ties is invoked in a second stage. The end result will be a unique belief state that represents what all the candidate states have in common. The relational framework is adequate, I propose, for analyzing the first part of this process. Yet, since the final result will be a unique belief state, belief revision is, in

the end, a functional process. Relational theorists have in effect weakened their case by boldly presenting their theory as an account of the complete revision process when in fact it is valid only for its first phase.

References

Alchourrón, C., Gärdenfors, P., and Makinson, D. (1985), “On the logic of theory change: partial meet contractions and revision functions”, *Journal of Symbolic Logic* 40: 510-530.

Gärdenfors, P. (1988), *Knowledge in flux: modeling the dynamics of epistemic states*, MIT Press: Cambridge, Mass..

Levi, I. (2004), *Mild contraction: evaluating loss of information due to loss of belief*, Clarendon Press: Oxford.

Lindström, S., and Rabinowicz, W. (1989), “On probabilistic representation of non-probabilistic belief revision”, *Journal of Philosophical Logic*, 89, 18: 69-101.

Lindström, S., and Rabinowicz, W. (1991), “Epistemic entrenchment with incomparabilities and relational belief revision”, in *The Logic of Theory Change*, Fuhrmann, A, and Morreau, M. (eds.): 93-126.

Olsson, E. J., and Westlund, D. (2006), “On the Role of the Research Agenda in Epistemic Change”, to appear in *Erkenntnis*, electronically published on the *Erkenntnis* website in 2006.

Rabinowicz, W. and Lindström, S. (1994), “How to model relational belief revision”, in D. Prawitz and D. Westerståhl (eds.), *Logic and Philosophy of Science in Uppsala*, Kluwer, 69-84.

Rott, H. (2006), “The value of truth and the value of information: on Isaac Levi’s epistemology”, in *Knowledge and Inquiry: Essays on the Pragmatism of Isaac Levi*, Olsson, Erik J. (ed), Cambridge University Press: New York, pp. 179-200.

Tennant, N. (2006), “New Foundations for a Relational Theory of Theory-Revision”, *Journal of Philosophical Logic* 35 (5): 489-528.