LUND UNIVERSITY

Functional vs. Relational Approaches to Belief Revision

Olsson, Erik J

Published in:
Belief Revision Meets Philosophy of Science

2011

[Link to publication](#)

Citation for published version (APA):
Olsson, E. J. (2011). Functional vs. Relational Approaches to Belief Revision. In E. J. Olsson, & S. Enqvist (Eds.), *Belief Revision Meets Philosophy of Science* (pp. 253-268). Springer.

Total number of authors:
1

General rights

Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Functional vs. Relational Approaches to Belief Revision

Erik J. Olsson
Lund University

Olsson, E. J. (2011). Functional vs. Relational Approaches to Belief Revision. In Olsson, E. J. & Enqvist, S. (Eds.) *Belief Revision Meets Philosophy of Science*. Springer Verlag.

Abstract: In belief revision theory there is an on-going debate between functionalists and relationalists that shows little sign of ending in shared agreement. Roughly, functionalists hold that the result of revising a cognitive state with some new datum is a unique rationally determined belief state. Relationalists disagree. In their view, there may be several rationally admissible results of revising a cognitive state with a given datum. In an attempt to contribute to conceptual clarity three ways of drawing the functionalist-relationalist distinction are identified, giving rise to six non-contradictory overall positions. I proceed to consider arguments in the literature for excluding some of these positions on logical, philosophical or other grounds. Finally, I argue that part of what feeds the functionalist-relationalist controversy is a false dilemma based on an implausible conception of what it means rationally to suspend judgment. Making this precise requires a formal framework that includes a representation of the agent's research agenda.

1. Introduction

This paper addresses a central issue in the study of theoretical rationality: how much of science is governed by such rationality and how much is rather a matter of taste or arbitrary choice?¹ The issue has been the concern of a long-standing debate between a functionalist and a relationalist approach in the part of philosophical logic devoted to the logic of belief revision.² Roughly, the functionalist holds that rationality uniquely determines the outcome

¹ For an overview of the history of this subject, see Stölzner (2000). Otto Neurath seems to have been the first to suggest, in his criticism of Descartes dating back to 1913, that in practical as well as in theoretical matters decisions must be made even if thorough investigation terminates in a set of equally reasonable alternatives. Neurath suggests basing such decisions on an "auxiliary motive", that is, ultimately by tossing a coin. For his discussion of Descartes, see Neurath (1981).

² See for example Lindström and Rabinowicz (1989, 1991), Doyle (1991), Galliers (1992), Rabinowicz and Lindström (1994), Hansson (1998), Levi (2004) and Tennant (2006a, 2006b).

of a given revision process. Relationalists disagree. According to them, considerations of theoretical rationality alone do not yield a unique recommendation for how to change one's view. Such considerations can only delineate a set of possible results among which the inquirer must ultimately choose without recourse to theoretical reason.

In my view, what has driven the debate is partly a lack of conceptual clarity, a claim to be substantiated as we proceed. In an effort to improve on the present state of affairs I shall differentiate between three different ways of drawing the line between functionalism and relationalism. How these distinctions bear on the functionalist-relationalist controversy will be the subject matter of most of the present essay.

By *strong functionalism* will be meant the view that the agent's old beliefs and the new datum are together sufficient rationally to determine the agent's new belief state after revision. The denial will be referred to as *weak relationalism*. Thus weak relationalism is the view that what can be said about rationality exclusively in terms of the old beliefs and the new datum is insufficient to single out a new state of belief.

Weak relationalism is compatible with the position that the new belief state after belief change is rationally determined given the new datum, the prior belief state and other relevant features of the agent's cognitive makeup. I will refer to this position as *weak functionalism*. The denial of weak functionalism will be called *strong relationalism*. Thus a strong relationalist holds that there is no set of cognitive features that, in conjunction with the old beliefs and the new datum, suffices to determine, in all cases, a unique new belief state after revision.

Finally, a given theory of belief revision is *relation-based* if its central primitive concept is that of a belief revision relation that is not a function. In general, a belief revision relation relates an old belief set and an input to several rationally admissible revised belief sets. A belief revision theory that takes as its chief primitive notion a belief revision relation that is a function is said to be *function-based*.

Based on these distinctions, eight combinations are *prima facie* possible (see Table 1). For instance, the position wFwRR amounts to a **weakly Functional and weakly Relational Relation-based** theory, i.e. relation-based theory to the effect that while the new belief state resulting from revision is not rationally determined by the old beliefs and the new datum alone (weak relationalism), it is thus determined once other cognitive aspects are taken into account (weak functionalism).

Two of these positions can be excluded on purely logical grounds. If the belief revision output is uniquely determined by the old beliefs and the input alone, then clearly it is also uniquely determined once more aspects of the agent's cognitive states are taken into

account. In other words, strong functionalism entails weak functionalism and hence also the negation of the opposite position, viz. strong relationalism. Therefore, sRsFF and sRsFR are contradictory.

Abbreviation	Result of belief revision determined all things considered?	Result of belief revision determined by old beliefs and new input?	Primitive concept a function?
wFsFF	Yes (weak Functionalism)	Yes (strong Functionalism)	Yes (Function-based)
wFsFR	Yes (weak Functionalism)	Yes (strong Functionalism)	No (Relation-based)
wFwRF	Yes (weak Functionalism)	No (weak Relationalism)	Yes (Function-based)
wFwRR	Yes (weak Functionalism)	No (weak Relationalism)	No (Relation-based)
sRsFF	No (strong Relationalism)	Yes (strong Functionalism)	Yes (Function-based)
sRsFR	No (strong Relationalism)	Yes (strong Functionalism)	No (Relation-based)
sRwRF	No (strong Relationalism)	No (weak Relationalism)	Yes (Function-based)
sRwRR	No (strong Relationalism)	No (weak Relationalism)	No (Relation-based)

Table 1: The prima facie possible overall positions in the debate between functionalists and relationalists

Indeed, strong functionalism, the view that the old belief state and the new datum are sufficient to determine the new belief state after revision, is a philosophically dubious position. It is part of belief revision folklore that “logical” properties alone are insufficient to characterize the revision process which is taken to involve various extra-logical features of a cognitive state, such as entrenchment, plausibility, rational choice or assessments of informational value. For that reason, we can exclude all positions that involve a commitment to strong functionalism, i.e., not only sRsFF and sRsFR but also wFsFF and wFsFR.

This leaves us with four serious contenders: wFwRF, wFwRR, sRwRF and sRwRR. They have in common a commitment to weak relationalism, i.e. to the view that the old beliefs and the new datum alone are insufficient to determine rationally the revised state of belief. Positions wFwRF and wFwRR involve weak functionalism as well. Positions sRwRF and sRwRR do not. Positions wFwRF and sRwRF are function-based. Positions wFwRR and sRwRR are not.

The next task will be to assess relationalist arguments to the effect that it is incoherent to advance weak or strong relationalism within a function-based framework. This, if true, would lead to the further exclusion of wFwRF and sRwRR, leaving only wFwRR and sRwRR still open. The latter represents a “fully relational” position which is not only strongly (and hence also weakly) relational, but also relation-based.

2. Are wFwRF and sRwRF incoherent?

The most well-known advocates of a relation-based approach to belief revision are Sten Lindström and Wlodek Rabinowicz (L&R) and Neil Tennant.³ They motivate their view essentially by referring to the basic relationalist intuition that the result of belief revision is not determined by rationality alone. According to L&R, the main idea behind their theory is “to allow for their being several equally reasonable revisions of a theory with a given proposition” (1994, p. 69). Tennant also stresses the importance to “countenance variety”. As he puts it, “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” (2006a, p. 490), the implication being that function-based theories fail to countenance variety.

Neither L&R nor Tennant makes a clear distinction between strong and weak relationalism. As a consequence, it is unclear whether what they find objectionable is already the combination of function-basedness and weak relationalism or merely the combination of function-basedness and strong relationalism, i.e. whether they wish to rule out wFwRF and hence also sRwRF, or only the latter.

My first observation will be that these two positions are both at least consistent. This will be shown by example. The AGM theory of Alchourrón, Gärdenfors and Makinson (1985) will be seen to instantiate position wFwRF, and – as Tennant et al would agree – there is no reason to think that AGM should be considered inconsistent, ill-defined

³ Lindström and Rabinowicz (1989, 1991), Rabinowicz and Lindström (1994), and Tennant (2006a, 2006b).

or otherwise internally defective. A theory instantiating position sRwRF can be obtained by slightly modifying the AGM theory.

In the AGM theory, a belief state is represented as a logically closed set of sentences, called a belief set. There are three principal types of belief change: expansion, revision and contraction. In expansion a new belief is added without any old beliefs being given up. In revision, the new information is added in a way that preserves consistency. Even if the new information is inconsistent with the original belief set, revision guarantees that the new belief set is consistent, provided that the information is itself non-contradictory. Finally, to contract a belief means to remove it from the belief set.

In the AGM theory, these types of changes are conceptualized as functions or “operations” from belief states to belief states. Thus AGM theory is function-based. The simplest of the three operations, expansion of a belief set K with a sentence α , denoted $K+\alpha$, is defined as the logical closure of the union of K and $\{\alpha\}$, i.e., $K+\alpha = Cn(K \cup \{\alpha\})$. Closing under logical consequence ensures that the result of expansion is indeed a belief set.

It is less obvious how to define the more interesting notion of genuine revision. As a preliminary, the AGM trio argued that a reasonable revision operation, denoted $*$, should satisfy eight so-called rationality postulates:

- (K*1) $K*\alpha = Cn(K*\alpha)$
- (K*2) $\alpha \in K*\alpha$
- (K*3) $K*\alpha \subseteq Cn(K \cup \{\alpha\})$
- (K*4) If $\neg\alpha \notin K$, then $Cn(K \cup \{\alpha\}) \subseteq K*\alpha$
- (K*5) $K*\alpha = K_{\perp}$ if and only if $\vdash \neg\alpha$
- (K*6) If $\vdash \alpha \leftrightarrow \beta$, then $K*\alpha = K*\beta$
- (K*7) $K*\alpha \wedge \beta \subseteq (K*\alpha)+\beta$
- (K*8) If $\neg\beta \notin K*\alpha$, then $(K*\alpha)+\beta \subseteq K*\alpha \wedge \beta$

The revision postulates are intended to capture the intuition that revisions should be, in a sense, minimal changes so that information is neither lost nor gained without compelling reasons. Thus changes in belief should obey a principle of informational economy. As Gärdenfors puts it, “the main thrust of the criterion of informational economy is that the revision of a belief set not be greater than what is necessary in order to accept the epistemic input” (1988, p. 53). According to K*3, for instance, $K*\alpha$ must not contain more information than what is included in $Cn(K \cup \{\alpha\})$.

According to AGM theory, the postulates for revision together express all that can rationally be said about the revision result exclusively in terms of the old beliefs and the new datum. Technically they delimit a class of total functions from $\mathbf{K} \times \mathbf{L}$ to $\mathbf{K}$, where $\mathbf{K}$ is the set of all logically closed sets of L-sentences, representing the class of rationally admissible revision functions. As is well known, the AGM revision postulates are together indeed insufficient to single out a single belief revision function as uniquely rational: the class of revision functions that satisfy the revision postulates contains more than one element.⁴

The fact that the AGM postulates fail to determine a unique belief revision function means that AGM is weakly relationalist. For the old belief set and the new datum are together insufficient to determine rationally the new belief set, which will also depend on what revision function is being employed. For the record, this is a consequence that Gärdenfors welcomes: “[t]he postulates (K*1) – (K*8) do not uniquely characterize the revision $K*\alpha$ in terms of only K and α . This is, however, as it should be. I believe it would be a mistake to expect that only logical properties are sufficient to characterize the revision process.” (1992, p. 11, notation and spelling adapted)

It is also part of the AGM theory that a unique revision result is guaranteed provided that more features of the agent’s cognitive apparatus are taken into account than just the old beliefs and the new input. To see why this is so, first note that revision can be reduced to contraction followed by expansion. In order to revise K by α , first contract by $\neg\alpha$ and then expand the result by α , where the contraction is performed for the purpose of making room for α . This proposal is codified in the so-called Levi identity: $K*\alpha = (K \div \neg\alpha) + \alpha$ (where $\div$ denotes contraction).

Now since expansion is already functionally defined, the Levi identity reduces the problem of how to define a revision function to the problem of how to define a contraction function. The basic mechanism for this purpose in the AGM model is that of partial meet contraction, as defined by the following identity: $K \div \alpha = \cap \gamma(K \perp \alpha)$. Here $K \perp \alpha$ is the set of all inclusion-maximal subsets of K that do not imply α , and γ is a selection function such that $\gamma(K \perp \alpha)$ is a non-empty subset of $K \perp \alpha$, unless the latter is empty, in which case $\gamma(K \perp \alpha) = \{K\}$. The intuition behind the use of the selection function is that it should select the “best” elements of $K \perp \alpha$ according to the agent’s theoretical preferences, which can be taken to be part of the agent’s cognitive constitution. Partial meet contraction amounts, then, to taking as the new state of belief what is common to the best maximal subsets of K

⁴ For a formal derivation, see Theorem 3 in Tennant (2006b).

that do not imply α , i.e., to suspend judgment between these subsets. This recipe is bound to lead to a uniquely determined result.

In this way, all three types of belief change addressed by the AGM theory can be accounted for in such a way that the result of change is uniquely determined given a suitable portion of the agent's cognitive state, including not just the agent's old beliefs, but also her theoretical preferences. Thus AGM is not only a function-based weakly relational theory; it is also weakly functionalist.

The AGM theory allows for the possibility that the function selecting a set of “best” inclusion-maximal subsets of K that do not imply α returns a set with more than member. If this were the end of the story and the agent were free to choose as he pleases between those members, we would have a strongly (and hence also weakly) relational function-based account of belief change, i.e. a theory instantiating position sRwRF.

What makes AGM weakly functional is the invocation of judgment-suspension as a tie-breaking rule. This is a powerful and intuitive strategy that seems eminently rational: what we ought to believe when confronted with several alternative theories or states of belief equally worthy of choice, and more so than any other theory or state of belief, is precisely what they all have in common. Other weakly functionalist theories are also based on suspension of judgment as a tie-breaking strategy. Levi's (1991) and Hansson's (1991) approaches are cases in point. They, too, instantiate position wFwRF, i.e. they belong to the class of weakly functionalist and weakly relationalist approaches that are also function-based.

The upshot is that the internal consistency of wFwRF and sRwRF is hardly in doubt. Yet, our point of departure in this chapter was the observation that researchers with relationalist inclinations have been questioning the coherence of those positions, arguing that relationalism should, in the interest of avoiding internal tension, be combined with a relation-based rather than a function-based framework. It is time to take a closer look at their argumentation, if only to pinpoint where they err.

The relational-based approach was first introduced in Lindström and Rabinowicz (1989). In that paper, a belief revision relation is defined as a ternary relation $\mathbf{R} \subseteq \mathbf{K} \times \mathbf{Con} \times \mathbf{K}$ satisfying the following axioms for all belief sets A , B , C , and all consistent propositions α and β :

$$(R0) (\exists D \in \mathbf{K}) A \mathbf{R}_\alpha D$$

$$(R1) \text{ If } A \mathbf{R}_\alpha B, \text{ then } \alpha \in B$$

$$(R2) \text{ If } A \cup \{\alpha\} \text{ is consistent and } A \mathbf{R}_\alpha B, \text{ then } B = A + \alpha$$

(R3) If $Cn(\{\alpha\}) = Cn(\{\beta\})$, then $A \mathbf{R}_\alpha B$ if and only if $A \mathbf{R}_\beta B$

(R4) If $A \mathbf{R}_\alpha B$, $B \mathbf{R}_\beta C$ and $B \cup \{\beta\}$ is consistent, then $A \mathbf{R}_{\alpha \wedge \beta} C$

As L&R read $A \mathbf{R}_\alpha B$, it means that B is a possible result for a given agent of revising A by the addition of α as a sole piece of new information. Axiom R0 expresses that belief revision should be defined for all belief sets A and consistent propositions α . Axioms R1-R4 mirror closely the AGM postulates for revision. Thus R3, corresponding to AGM postulate K*6 usually referred to as the postulate of *extensionality*, expresses that in revision only the logical content of the input sentence is important, not its syntactic formulation. A belief revision relation $\mathbf{R}$ is said to be functional if, in addition to R0 – R4, it satisfies:

(R5) If $A \mathbf{R}_\alpha B$ and $A \mathbf{R}_\alpha C$, then $B = C$

Tennant (2006a) presents a relation-based belief revision theory much along the same lines. He studies two primitive relational notions: $\downarrow(J, K, \alpha)$, meaning that J is a contraction of K with respect to α , and $\uparrow(J, K, \alpha)$ meaning that J is a revision of K with respect to α . Tennant's theory differs in some respects from L&R's treatment. Unlike L&R, Tennant argues for a principal case analysis of the relevant belief change relations. For revision, the principal case is where a consistent belief state K is revised with respect to a non-contradictory sentence α that is inconsistent with K . (L&R's relation is defined also for inconsistent belief sets and for new input consistent with the belief set.) A second difference is that Tennant formulates AGM style axioms in the form of a number of elimination rules and one introduction rule.

According to L&R, a relation-based theory "is natural if we think that the agent's policies for belief change may not always yield a unique belief set as the result of revising a given belief set A with a proposition x " (1989, p. 25). In their 1994 paper, they write in retrospect:

Some years ago, we proposed a generalization of the well-known approach to belief revision due to Peter Gärdenfors (cf. Gärdenfors 1988). According to him, for each theory G (i.e. each set of propositions closed under logical consequence) and each proposition α , there is a unique theory, $G^*\alpha$, which would be the result of revising G with α as a new piece of information. There is a unique theory which would constitute the revision of G with α . Thus, belief revision is seen as a function. Our proposal was

to view belief revision as a relation rather than as a function on theories. The idea was to allow for there being several equally reasonable revisions of a theory with a given proposition (notation adapted).

The implication is that Gärdenfors's theory fails to "allow for there being several equally reasonable revision of a theory with a given proposition".

Similarly, Tennant motivates his choice of a relation-based rather than a function-based formal framework as follows:

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 $\langle K, \alpha \rangle$. 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 $\langle K, \alpha \rangle$. 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, notation adapted).

The implication, again, is that AGM fails to "furnish such variety".

Tennant et al apparently believe that AGM fails to be a theory that treats contracting and revising as "processes that can produce a variety of possible values on any given input $\langle K, \alpha \rangle$ ", i.e. that AGM is, in our terminology, a strongly functionalist theory. This is incorrect. As we have seen, AGM allows for a variety of possible values on any given input $\langle K, \alpha \rangle$. The values will differ depending on what contraction/revision function is being employed. AGM is therefore to be classified as a weakly relationalist theory and not as a strongly functionalist one. This observation undercuts the main relationalist motivation for introducing a relation-based theory of belief revision as an alternative to AGM style function-based theories.⁵

⁵ Tennant's reference to "non-deterministic processes" might suggest a *causal* difference between AGM and a relation-based theory. On this reading, Tennant is taking AGM to be a theory according to which the new belief state after revision is causally determined by the old belief state and the new input, and he is seeing the relationalist position as the opposite view that the most that can be said is perhaps that there is a certain probability (less than 1) that a given belief state will ensue. However, it

This is not to say that it may not be fruitful to study relation-based belief revision for methodological or systematic purposes. This is also an aspect that Tennant stresses, suggesting reformulating the AGM theory within a relational setting because this has the “methodological advantage” (Tennant 2006a, 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 which show that the AGM postulates for revision are too liberal in the principal case of revision of a non-contradictory sentence which is inconsistent with the old beliefs (Tennant 2006b).

Tennant’s reformulation has the additional systematic benefit of making it possible to derive, rather than stipulate, K^*6 , the AGM principle of extensionality, with which we have already made acquaintance. Tennant argues that stipulating rather than deriving the extensionality principle is essential to the functional framework. While there may be systematic reasons for studying belief revision in a relation-based framework, this does not bear on the internal coherence of developing some form of relationalism within a function-based framework.

3. A functionalist-relationalist dilemma

We have seen that there seems to be no good reasons to rule out either position wFwRF or sRwRF as inconsistent, incoherent or otherwise internally defective. Hence, wFwRF, wFwRR, sRwRF and sRwRR are all still in the race. My point of departure in this section will be strong relationalism, which is an ingredient in the last two positions. What can be said in favor of the view that the result of belief revision may not be rationally determined even if all cognitive aspects of the agent have been taken into account?

In a paper presented at a workshop in 1989 and later published in a proceedings volume in 1991, L&R give what I take to be an argument for strong relationalism. 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

is generally agreed that theories of belief revision in the tradition of AGM are normative theories of rationality and not descriptive theories. Such theories of belief revision do not purport to account for how people actually change their views but rather for how they ought to do it. For that reason, it is difficult to see how a *causal* distinction between determinism and indeterminism could be relevant in this field of research. Rather what is meant by claiming or implying that AGM is deterministic must be that, according to AGM, the result of belief revision is *rationally* determined, so that given an old belief state and a new datum, there is only one new belief state that is compatible with the AGM principles of rationality. This amounts to claiming or implying, falsely, that AGM is a strongly functionalist theory when in fact it is a weakly relationalist one.

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 longer has to be nested. It will no longer 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 thus 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 by propositions being "different". Their reference to "totally unrelated" propositions suggest that they have in mind unrelatedness with respect to topic. Still, there are many topic-wise 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 do not.

Isaac Levi has provided 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).

In other words, agents with conflicting theoretical goals and values may end up in a situation where there is no unique way to order beliefs with respect to epistemic entrenchment but several equally admissible orderings. This would seem to open up for the possibility that a given change in belief may give rise to an indeterminate result; that there could be several equally rational ways to change beliefs. A strongly relationalist theory can accommodate this indeterminacy. A functionalist theory, it may appear, cannot. From this perspective, the strongly relationalist approach may seem to be the philosophically more meritorious position.

The question however is whether the strong relationalist, in her argumentation, has really exhausted all the resources of theoretical rationality. Maybe it is true that there is not always a unique way to order beliefs with respect to how entrenched they are. But many researchers, among them advocates of AGM or its variants, would be unhappy with letting this be the end of the story. In such cases, they would say, theoretical rationality dictates that we invoke a rule for ties as a further criterion. 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.

But the relationalist could counter as follows: The relational view 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 a *bone fide* scientific theory change. Suppose we have a theory which must be changed in the light of the outcome of one or more experiments and, in addition, 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. And, crucially, what guarantee do we have for thinking that suspending judgment between these optimal theories will give rise to a theoretical position that is itself optimal or even qualifies as a scientific theory at all?⁶ The result of suspending judgment between the different theories that surface concerning the extinction of the dinosaurs would hardly itself pass as a theory on the matter in the scientific sense.

The problem of justifying judgment-suspension in the apparent absence of a guarantee that it leads to a good theory motivates much of Isaac Levi's later research on belief revision, from Levi (1991) to Levi (2004). His approach has been to insist that suspending judgment between optimal theories results in a theory that is not only rationally admissible but even optimal. Levi bases his initially somewhat dubious view on an intricate analysis of the theoretical values that are involved in theory choice and how they combine into his measure of informational value.

While Levi should be credited for stressing the importance of the problem of ties and for providing a detailed and sophisticated solution to it, it is difficult to avoid the impression that his measure of informational value was chosen mainly for the reason that it gives the desired result vis-à-vis judgment-suspension. For related criticism, see Rott (2006). I believe that Levi's persistent efforts in this direction have been less than convincing, and that the prospects of making any further progress are dim. But, as always, the only good objection is another theory. I will therefore proceed to propose an alternative approach to the dilemma which has been my main concern in this section. The alternative resolution is not only simpler and, I believe, initially more appealing than Levi's proposal; it also has an independent standing that the latter to some degree lacks.

4. Weak functionalism and dynamic caution: a preliminary defense

The dispute between strong relationalism and weak functionalism concerns, again, what to accept when there are several theoretical options (hypotheses) that are "best", i.e., equally good, and better than any other option (hypothesis). The bold strategy recommended by the strong relationalist is to pick one of the best theoretical options arbitrarily, or in some other way that does not appeal to the principles of theoretical rationality, as the new view on the matter. This approach, though undoubtedly representing a minimal change approach, is intuitively repugnant. It doesn't seem rational to choose to fully to believe something when

⁶ The objection raised in this paragraph is due to Sten Lindström (personal communication).

one might just as well have fully believed something else instead. The skeptical strategy to which the weak functionalist pledges allegiance involves taking as the new position what all the best theoretical options have in common, i.e., to suspend judgment between them. The drawback of this strategy is that there is, *pace* Levi, no compelling reason to think that suspending judgment between optimal positions leads to a position that is itself optimal or even satisfactory. Thus we are presented with what seems to be a forced choice between two unattractive positions.

Otherwise put: we have to face the predicament that none of positions in the functionalist-relationalist controversy that have so far survived critical scrutiny – wFwRF, wFwRR, sRwRF or sRwRR – seems acceptable. For they all involve a commitment to either weak functionalism or strong relationalism, and we have just seen that both seem unacceptable from the standpoint of pre-systematic intuition.

I believe, however, that the dilemma is based on a misconception of what it means rationally to suspend judgment. It seems true that the belief state corresponding to suspension need not itself be one of the optimal belief states; typically, it won't. It is also true that one should never settle for a suboptimal option. But this only shows that if one accepts (only) the belief state corresponding to suspension, one must do so *without settling for it*. In other words, that belief state should not be one that we rest content with. Rather, accepting it should generate or preserve, as the case may be, a commitment to settle the original issue as to which of the best options should be accepted. The goal to settle that issue should still be on the agent's *research agenda*.

What I am suggesting is that the belief state corresponding to suspension of judgment 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 nevertheless continue asking which one of those best theories should eventually be accepted. Still on the agenda, this question serves to motivate further inquiry and deliberation aiming at the ultimate acceptance of one of the theories among which judgment was suspended.

What I am proposing, more precisely, is that epistemic states be viewed as more complex objects consisting not only of a belief set, an epistemic entrenchment ordering (or some other suitable ordering or choice mechanism) but also of a research agenda. The research agenda can be represented as a set of questions (Olsson and Westlund, 2006). A

question, in turn, can be represented as a set of potential answers. In the following, the entrenchment ordering (choice mechanism etc) will be disregarded.⁷

Suppose, for instance, that as the effect of receiving new information three alternative states of belief B_1 , B_2 and B_3 present themselves as being as good as any other. Subsequent inquiry and deliberation reveals that B_1 and B_2 are equally good and better than B_3 . On the present view, the inquirer is now justified in believing all and only what B_1 and B_2 have in common, provided that she retains on her agenda the task of deciding which one of B_1 or B_2 is ultimately to be chosen. On my reconstruction, the new epistemic state should be something like $E = \langle B_1 \cap B_2, \{ \{B_1, B_2\} \} \rangle$, where $B_1 \cap B_2$ is the new state of belief and $\{ \{B_1, B_2\} \}$ the new research agenda. Intuitively, this epistemic state is better than either of the two “bold” alternatives of choosing as the new epistemic state either $E_1 = \langle B_1, \emptyset \rangle$ or $E_2 = \langle B_2, \emptyset \rangle$ or, for that matter, $E_3 = \langle B_1 \cap B_2, \emptyset \rangle$. The latter represents suspending judgment while maintaining an empty agenda.

However, it is one thing to assert, as I have done, that dynamic caution is admissible and optimal; it is quite another to supply a decision theoretic argument to the same effect. Supplying such an argument would involve, among other things, supplying a theory for how more complex epistemic states – states that involve research agendas – are to be evaluated and compared with respect to informational value. This issue still needs to be addressed in its entirety. What can be said already at this point is that $\langle B_1 \cap B_2, \{ \{B_1, B_2\} \} \rangle$ should be decision-theoretically more advantageous than $\langle B_1 \cap B_2, \emptyset \rangle$ from the point of view informational value. In the former case, but not in the latter, there are things on the research agenda, meaning that, everything else being the same, the prospects of a future increase in informational value is relatively high. It is more difficult to rationalize in decision theoretical terms why $\langle B_1 \cap B_2, \{ \{B_1, B_2\} \} \rangle$ should be considered preferable to, say, $\langle B_1, \emptyset \rangle$. Intuitively, this has to do with the increased risk of error involved in accepting the latter as opposed to the former, but – as Levi has taught us – risk of error is a problematic notion in the context of belief contraction.

The new agenda-based proposal raises another open question of some urgency. As Rott (2006), p. 191, points out, Levi’s theory can be seen as an attempt to satisfy three principles at once. Suppose the agent has decided to contract α from her corpus. First, there is what Rott calls the Decision-theoretic Rule which says that “The corpus after a contraction must be optimal, that is, it must minimize the loss of informational value among all corpora expelling the hypothesis α ” (ibid., notation adapted). Second, the Rule for Ties in

⁷ For more recent accounts of the agenda in belief revision, see Enqvist (2009) and Genot (2009).

Contraction should be obeyed: “given a set of optimal contraction strategies, one should always choose the weakest of them if it exists” (Levi, 2004, p. 119). Finally, there is the Intersection Equality: “If members of a set S of contractions from K are equal in informational value, their intersection is equal in informational value to the informational value of any element of S ” (Levi, 2004, p. 125, notation adapted). In the special case of S being the set of contractions that are optimal in the sense of minimizing loss of informational value the Intersection Equality ensures that the intersection of S , representing judgment-suspension in a sentence-based framework, is also optimal. Levi, as we saw, achieves the satisfaction of these three principles *ad hoc* by introducing a measure of damped informational value that significantly lacks independent support.

Now, it would be interesting to find out whether our new, more complex framework, involving research agendas, could come to Levi’s rescue here. In other words, does the shift to the agenda-based framework allow the simultaneous satisfaction of Levi’s three principles without there being any need for substantial extra assumptions? This, too, is a question I must leave uninvestigated.

Summing up, the present way out of the dilemma between the “bold” and “skeptical” approaches is based on the observation that the skeptical approach has not been advocated in its most plausible form. AGM style theorizing on the matter conveys the impression that the skeptical position resulting from judgment-suspension is fine as it is, i.e., that there is no need to inquire any further. However, there is a more plausible form of the skeptical approach that combines taking what the theories or belief states have in common as the new belief content with updating the research agenda in the manner suggested. What I am getting at can be described as the difference between a “static” and a “dynamic” caution; between being satisfied with having suffered informational loss and accepting such loss only as part of a temporary retreat. The kind of caution that I favor is dynamic.

5. Conclusion

In an attempt to contribute conceptual clarity I have distinguished three ways of drawing the line between functionalism and relationalism: weak functionalism vs. strong relationalism, strong functionalism vs. weak relationalism and finally function-basedness vs. relation-basedness. These distinctions were then employed to shed light on the functionalist-relationalist debate in philosophical logic. In the final section, I argued for a weak functionalism according to which the belief state resulting from revision is rationally determined given the new datum, the prior belief state and other relevant features of the

agent's cognitive makeup. My brand of such functionalism was seen to be one with a "relationalist touch": it concedes to the relationalist that the skeptical position resulting from suspension of judgment is unreasonable if interpreted as a static end-point of inquiry. Rather, such suspension must be interpreted dynamically as an intermediate position in an on-going investigation aiming at the eventual acceptance of one of the optimal theories among which judgment was originally suspended. Making this precise is a task left for future research. It requires, among other things, a formal framework that includes a representation of the agent's research agenda.

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.

Doyle, J. (1991), "Rational Belief Revision: Preliminary Report", pp. 163-174 *Principles of Knowledge Representation and Reasoning*, Allen (Ed.), Morgan Kaufmann, Los Altos, CA.

Enqvist, S. (2009), "Interrogative Belief Revision in Modal Logic", *Journal of Philosophical Logic* 38 (5): 527-548.

Galliers, J. R. (1992), "Autonomous Belief Revision and Communication", pp. 220-246 in *Belief Revision*, Gärdenfors, P. (Ed.), Cambridge Tracts in Theoretical Computer Science 29, Cambridge University Press.

Genot, E. J. (2009), "The Game of Inquiry: The Interrogative Approach to Inquiry and Belief Revision Theory", *Synthese*, published online.

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

Gärdenfors, P. (1992), "Belief Revision: An Introduction", in Gärdenfors, P. (Ed.), *Belief Revision*, Cambridge Tracts in Theoretical Computer Science 29, Cambridge University Press.

Hansson, S. O. (1991), *Belief Base Dynamics*, Uppsala: Acta Universitatis, Upsaliensis.

Hansson, S. O. (1998), "Revision of Belief Sets and Belief Bases", pp. 17-75 in *Belief Change*, Gabbay, D. M. and Smets, P. (Eds.), Handbook of Defeasible Reasoning and Uncertainty Management Systems, vol. 3, Kluwer.

Levi, I. (1967), *Gambling with Truth*, MIT Press.

Levi, I. (1991), *The Fixation of Belief and Its Undoing*, Cambridge University Press.

Levi, I. (2004), *Mild Contraction: Evaluating Loss of Information due to Loss of Belief*, Clarendon Press: Oxford.

Levi, I. (2006), "Answering back: Reply to Olsson", in *Knowledge and Inquiry: Essays on the Pragmatism of Isaac Levi*, Olsson, Erik J. (ed), Cambridge University Press: New York, pp. 153-159.

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.

Neurath, O. (1981), "Die Verirrten des Cartesius und das Auxiliarmotiv (Zur Psychologie des Entschlusses)", in *Gesammelte philosophische und methodologische Schriften*, vol. 1, Haller, R. and Rutte, H. (eds.), Wien: Hölder-Psichler-Tempsky, 57-67. The paper is dated 1913.

Olsson, E. J. (2006), "Potential Answers – To What Question?", in *Knowledge and Inquiry: Essays on the Pragmatism of Isaac Levi*, Olsson, Erik J. (ed), Cambridge University Press: New York, pp. 157-166.

Olsson, E. J., and Westlund, D. (2006), "On the Role of the Research Agenda in Epistemic Change", *Erkenntnis* 65 (2): 165-183.

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.

Stölzner, M. (2000), “An Auxiliary Motive for Buridan’s Ass: Otto Neurath on Choice without Preference in Science and Society”, *Conceptus* 33 (82): 23-44.

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

Tennant, N. (2006b), “On the Degeneracy of the Full AGM-theory of Belief Revision”, *Journal of Symbolic Logic* 71(2), pp. 661-676.