



LUND UNIVERSITY

On the Role of the Research Agenda in Epistemic Change

Olsson, Erik J; Westlund, David

Published in:
Erkenntnis

DOI:
[10.1007/s10670-006-9001-6](https://doi.org/10.1007/s10670-006-9001-6)

2006

[Link to publication](#)

Citation for published version (APA):

Olsson, E. J., & Westlund, D. (2006). On the Role of the Research Agenda in Epistemic Change. *Erkenntnis*, 65(2), 165-183. <https://doi.org/10.1007/s10670-006-9001-6>

Total number of authors:

2

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

On the Role of the Research Agenda in Epistemic Change

Erik J. Olsson and David Westlund
Lund University

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

Abstract

The standard way of representing an epistemic state in formal philosophy is in terms of a set of sentences, corresponding to the agent's beliefs, and an ordering of those sentences, reflecting how well entrenched they are in the agent's epistemic state. We argue that this wide-spread representational view – a view that we identify as a “Quinean dogma” – is incapable of making certain crucial distinctions. We propose, as a remedy, that any adequate representation of epistemic states must also include the agent's research agenda, i.e., the list of question that are open or closed at any given point in time. If the argument of the paper is sound, a person's questions and practical interests, on the one hand, and her beliefs and theoretical values, on the other, are more tightly interwoven than has previously been assumed to be the case in formal epistemology.

1. Introduction

In the AGM model of belief revision, a state of full belief is represented syntactically as a set of sentences of some formal language. The set is assumed to be closed under logical consequence. Three main ways of changing one's epistemic state are identified: expansion, revision, and contraction. Expansion is the process of accepting an epistemic input α , represented as a sentence. The result of expansion is the logical closure of the set-theoretical union of the original state of belief and the new belief. In formulas, $K+\alpha = \text{Cn}(K \cup \{\alpha\})$. Revising K by α , in the case where K is inconsistent with α , is the result of replacing $\neg\alpha$ by α . Finally, the contraction of K by α is the result of deleting α from K . It is generally thought that revision can be defined in terms of expansion and contraction. Revising K by α amounts to first contracting K by $\neg\alpha$ so as to make room for α and then expanding K by α . In a formula:

$$K * \alpha = (K \div \neg\alpha) + \alpha$$

This formula is called the Levi identity.¹ Given the Levi identity, the remaining problem of belief revision is to define a contraction operation. For once we have a contraction operation we can define revision with the help of the Levi identity. The problem of defining contraction is, however, a non-trivial one. First of all, little can be said about contraction in purely logical terms. For instance, if we want to contract $\alpha \& \beta$ from K , there is, as far as logic is concerned, a number of ways in which this can be achieved. We can delete α but retain β , delete β but retain α , or delete both α and β . The solution to this problem is generally taken to be to add more structure to the representation of an epistemic state. The main idea is to add an ordering of the sentences in the belief state that reflects the beliefs' relative degree of entrenchment. Let us call such an ordering an entrenchment ordering. Suppose for example that α is more entrenched in K than β is. Then contracting $\alpha \& \beta$ from K would involve deleting β but retaining α . Entrenchment-based contraction can now be defined along these lines. A more general approach is to add instead a choice function that selects the sentences to be removed (Alchourron, Gärdenfors, and Makinson 1985, Rott 2001). This approach is more general since the choices may but need not be based on considerations of entrenchment.

Since the AGM model was presented, a large number of variations have been proposed and investigated, including belief base revision (Hansson 1994), relational belief revision (Lindström and Rabinowicz 1991), ranking theory (Spohn 1988), and dynamic doxastic logic (Segerberg 1995, 1998). They share the basic structure of the AGM model in that epistemic states are thought of as belief states together with some ordering or choice mechanism defined on the beliefs. Certainly, these models differ in a number of interesting respects, but these differences will not be important for the purposes of this paper.

As we just said, Segerberg's dynamic doxastic logic is a case in point. Segerberg (1995, 536) introduces the operator \mathbf{B} from doxastic logic, read as "the agent believes that". He also introduces the operators $[\mathbf{+B}\alpha]\beta$ read as "after expanding the set of beliefs by $\mathbf{B}\alpha$, it is the case that β " and $[\mathbf{-B}\alpha]\beta$ read as "after contracting the set of beliefs by $\mathbf{B}\alpha$, it is the case that β ". The truth conditions of these operators is given by a system of spheres roughly in the sense of Lewis's semantics for counterfactuals, but where the innermost sphere is a set of worlds and not (as in Lewis's original theory) a particular world. The innermost sphere represents what the agent believes and the other spheres are used in contraction, the idea being that the agent should

¹ We would like to thank Bengt Hansson and the participants of his weakly research seminar in philosophy of science for their lively discussions and helpful advice on earlier drafts. Thanks also to Sven Ove Hansson who commented on an earlier version. Finally, we are obliged to Ulrich Gähde and the members of his research group in Hamburg for their input.

contract by moving to the smallest set of spheres where he or she no longer believes the proposition to be contracted.

As the reader may have anticipated, modeling epistemic states as a belief state plus an entrenchment orderings of its elements corresponds closely to Quine's view of how our belief systems are structured:

[...] total science is like a field of force whose boundary conditions are experience. A conflict with experience at the periphery occasions readjustment in the interior of the field. Truth values have to be redistributed over some of our statements. Reevaluation of some statements entails reevaluation of others, because of their logical interconnections [...] (Quine, 1953, 42)

Later Quine writes:

Certain statements, though *about* physical objects and not sense experience, seem peculiarly germane to sense experience - and in a selective way: some statements to some experiences, others to others. Such statements, especially germane to particular experiences, I picture as near the periphery. But in this relation of 'germaneness' I envisage nothing more than a loose association reflecting the relative likelihood, in practice, of our choosing one statement rather than another for revision in the event of recalcitrant experience. (Quine, 1953, 43)

It is natural to think of "total science" as an idealized belief state and of the "loose association reflecting the relative likelihood, in practice, of our choosing one statement rather than another for revision" as corresponding to an entrenchment ordering of the beliefs that form part of that state.

Our first task in this paper will be to argue, in section 2, that the Quinean representation of epistemic states has some inherent limitations making it incapable of modeling certain types of belief change convincingly. The basic assumption behind Quinean belief change – the assumption that we find objectionable – says that whether or not a change from one epistemic state E_0 to another E_1 is rational depends only on the beliefs and entrenchment orderings of those beliefs in E_0 and E_1 . We will refer to it, somewhat provocatively, as a Quinean dogma of epistemic change. According to it, we could in principle decide whether a given change from E_0 to E_1 was rational by inspecting the changes that the beliefs and orderings of those beliefs underwent in the transition from E_0 to E_1 . In section 3, we make some suggestions as to what belief change without the Quinean dogma might look like. Section 4 is devoted to a comparison with Isaac Levi's theory.

2. Counterexamples to the Quinean dogma

In this section we argue that the following types of change cannot be properly handled on the Quinean dogma.

- (1) Falling into inconsistency
- (2) Ceasing to believe something in the absence of inconsistency
- (3) Accepting an ad hoc hypothesis in order to save a given theory from some sort of trouble

In order to model these types of change more structure is needed, or so we will argue.

Let us start with inconsistency. Suppose you are convinced that Peter is on the other side of the globe. Suddenly, you see Peter just entering his office. You have fallen into inconsistency, for you believe both that Peter is over there and that he is on the other side of the globe, and so something must be done. One may, to be sure, quarrel about how exactly to analyze this type of example. Does it really involve inconsistency or merely some sort of anomaly? What one cannot quarrel about, however, is that it is urgent to resolve the conflict, whatever its exact nature. Our claim is that the urgency aspect of this situation cannot naturally be modeled in purely Quinean terms, i.e., in terms of beliefs and entrenchment only. To clarify the point, suppose you are interested in Peter's whereabouts so that you would really like to know where he is. Then it would be irrational of you not to inquire further into his whereabouts given the inconsistency (anomaly) just described. You are, in some sense, committed to conducting further inquiries, given, of course, that you take interest in the matter in the first place. You cannot rest content with the current state of affairs. And yet there seems to be no natural way to represent this element of dissatisfaction along Quinean lines.

We now turn to the second type of change: ceasing to believe something in the absence of inconsistency. Suppose you believe that the National Lottery is fair. Now you are informed that the wife of the organizer has won the highest prize. This is not inconsistent with your belief that the lottery is fair, but you decide nonetheless that you cannot take it for granted anymore that the lottery is fair. Perhaps the lottery has been rigged in favor of the organizer's wife. It seems that, given that you care about the fairness of the lottery in the first place, your retracting your belief that it is fair should be accompanied by a reopening of the question as to its fairness. That is, you are, in some sense, obliged to take continued interest in the matter. Again, it is far from clear how this new obligation fits into the Quinean picture. It seems that, if all we have are beliefs and orderings of beliefs, we cannot distinguish the case just described from one in which you simply drop your belief that the lottery is fair without undertaking any commitments of the sort described. That, however, would seem irrational, unless you lost interest in the matter in the

process of learning about the winner. The point is that this sort of contraction is (normally) rational just in case it is made for the purpose of opening a neutral investigation into what the best answer is. In the case of the lottery, it would be rational in the circumstances as described to give up the view about the fairness if, and only if, that means opening up for a neutral examination of “fair” contra “rigged”.

A similar point can be made in connection with the acceptance of auxiliary hypotheses. Suppose you accept planetary theory T . Now you make some telescope observations incompatible with T . Suppose now that you come to believe that T would be compatible with the data provided one more epicycle is added.² Call the resulting theory T' . You decide to accept T' , at least for the time being. Under what conditions is such addition of auxiliary hypotheses in order?³ The idea is that this may well be legitimate if it is combined with a commitment to search for independent evidence for the auxiliary hypothesis with the aim of relieving it of its auxiliary character. By contrast, it is irrational, or at least not scientifically virtuous, to introduce an auxiliary hypothesis without undertaking such a commitment to further investigation. Once more, it is difficult to see how these observations could be accommodated by the Quinean representational paradigm.

3. Belief revision without the dogma

The commitments that we are interested in and that, in our view, play a crucial role in belief change are typically what we would find on the inquirer’s research agenda. Our aim is to make the research agenda and its role in rational belief change explicit. We are not, at this point, in a position to offer a full formal theory of how this should work. What we can do is to make some suggestions in that direction. The main purpose of this section is to substantiate the modest claim that it is possible to devise a formal model of belief change that takes seriously the research agenda and its role in epistemic change.

In the Quinean model, epistemic change is seen essentially as a function from a pair $\langle K, E \rangle$ to another pair $\langle K', E' \rangle$ where K (K') is a belief state and E (E') an ordering of epistemic entrenchment. What we propose is that it should be viewed as a function from a triple $\langle K, E, A \rangle$ to another triple $\langle K', E', A' \rangle$ where A (A') is the inquirer’s research agenda. Following AGM, a belief state is represented by a logically closed set of sentences. We will leave the entrenchment

² In the example, the motivation for introducing the auxiliary hypothesis is one of removing inconsistency which makes it tempting to subsume this strategy under the heading of inconsistency-removal. In general, however, auxiliary hypotheses may be introduced for other reasons as well, including reasons of coherence.

³ As Frank Zenker has pointed out, what is commonly referred to as addition of auxiliary hypotheses is actually not to be seen as mere addition but rather as a kind of revision (personal communication).

ordering aside for the purposes of this paper. There is no shortage of suggestions in the literature as to what properties the entrenchment ordering should be assumed to satisfy and how it should behave in epistemic change (e.g. Gärdenfors 1988). Nor will we have much to say about how to contract a set of sentences with a given proposition. A more urgent question for our purposes is how to represent the research agenda in formal terms.

We may think of the research agenda as a set of questions which the researcher wants to have answered.⁴ How then should we represent a question? A simple idea is to identify a question with the set of all its potential answers. If we take an answer to be a sentence in the underlying formal language L , this amounts to identifying a question with a set of sentences. Not all sets of sentences correspond, however, intuitively to questions. The answers to a genuine question are mutually exclusive and exhaustive relative to what is known or believed to be the case. This means that, for all we know, precisely one of the answers is true. Moreover, all answers to a genuine question are serious possibilities, meaning that they are not excluded by what is known or believed to be the case. We define a *K-question* to be a set Q of sentences that is exclusive and exhaustive relative to K and that is such that K does not entail that the true answer is in a proper subset of Q . Thus $Q = \{\alpha_1, \dots, \alpha_n\}$ is a *K-question* if and only if (1) K entails $\alpha_1 \underline{\vee} \dots \underline{\vee} \alpha_n$, and (2) there is no non-empty $Q' \subset Q$ such that K entails $\vee Q'$. The symbol $\underline{\vee}$ stands for exclusive disjunction. For the sake of simplicity, we consider only questions with finitely many potential answers. $\vee Q$ ($\underline{\vee} Q$) denote arbitrary inclusive (exclusive) disjunction of all the elements of Q . If $Q = \{\alpha\}$, then $\vee Q = \underline{\vee} Q = \alpha$. We note that, if $\alpha \in K$, then $\{\alpha\}$ is a *K-question*. We will refer to (1) and (2) as the *presuppositions* of the question. We let Q_K denote the set of all *K-questions*. We can now say what a *K-relative research agenda*, or *K-agenda* for short, is: A is a *K-agenda* if and only if $A \subseteq Q_K$. If $Q \in A$, we sometimes say that Q is *on* the agenda A . We note that a *K-agenda* cannot contain two questions where one is a proper subset of the other. For if Q is on the *K-agenda* A , then Q is a *K-question*. Hence, by condition (2) for *K-questions*, there is no $Q' \subset Q$ such that K entails $\vee Q'$. Hence, by condition (1) for *K-questions*, there is no Q' on A such that $Q' \subset Q$.⁵

We require the following of an epistemic state:

⁴ This assumption presupposes that the commitments on the research agenda are always commitments to answer questions. The reasonableness of this assumption would have to be substantiated in a more extensive treatment of the subject.

⁵ We are, of course, aware of the fact that identifying a question with its set of potential answers is not always a reasonable policy. There are certainly cases where we entertain questions without having any particular idea of how they might be answered. A more realistic model would have to be able to accommodate such situations as well.

Postulate 1: If $S = \langle K, E, A \rangle$ is an epistemic state, then A is a K -agenda.

We now turn to the dynamic aspects of the research agenda. We will focus on (consistent) expansion and contraction in the absence of inconsistency.

3.1. Non-Quinean expansion

In this section we consider the problem of how to define expansion in a non-Quinean setting. $S+\alpha$ denotes the result of adding α to the epistemic state S . We assume that α is consistent with the belief state component of S . We want to define expansion in a way that makes it a function from epistemic states to epistemic states:

Postulate 2: $S+\alpha$ is an epistemic state.

The effect of expansion on the belief state component of an epistemic state is simply that the new belief is added set-theoretically to the set representing the old belief whereupon the result is closed under logical consequence. We leave the entrenchment part open, as explained above. The interesting issue, from our perspective, is what should happen to the research agenda in expansion. The first idea that comes to mind is that expansion by α should have the sole effect of deleting from the research agenda all questions that have α as a potential answer.⁶ Given an epistemic state $S = \langle K, E, A \rangle$, we propose the following definition of expansion for a start:

(D1) $S+\alpha = \langle \text{Cn}(K \cup \{\alpha\}), E', A' \rangle$, where $A' = A / \{Q \in A \mid \alpha \in Q\}$.

According to this proposal, the new research agenda after expansion is the same as the old minus all questions for which α is one of the potential answers.

There are, however, at least two problems with this proposal. The most acute problem is that it fails to guarantee that the result after expansion is a new epistemic state, that is to say, it may lead to violations of Postulate 2. To see the problem, suppose that $S = \langle K, E, A \rangle$ is an epistemic state such that $Q \in A$, where $Q = \{\alpha, \beta, \gamma, \}$. Since S is an epistemic state, A is a K -agenda (by Postulate 1). Hence K entails the presuppositions of Q , i.e., it entails the truth of exactly one of α , β and γ ,

⁶ We allow, for the sake of generality, that one and the same belief can answer more than one question on the current research agenda. There may be philosophical reasons for thinking that this cannot be the case. We will leave this matter aside for the purposes of this paper. The matter is related to the lottery paradox. For relevant discussions, see Levi (1967) and Olsson (2005).

and none of them can be excluded on the basis of K . Suppose now that we expand by $\alpha \vee \beta$. By D1, this will give rise to a new belief state, K' , in which $\alpha \vee \beta$ is believed to be the case and a new agenda $A' = A / \{Q \in A \mid \alpha \in Q\}$. Since $\alpha \vee \beta$ is not an answer to Q , $Q \in A'$. But K' entails $\alpha \vee \beta$, and so Q is not a K' -question. We recall that for Q to be a K' -question, its presuppositions must most hold relative to K' . In particular, it must not follow from K' that the true answer is in a proper subset of Q . That presupposition is violated here. Hence, by Postulate 1 and in violation of Postulate 2, the new state after expansion is not an epistemic state.

Second, D1 does not allow us to distinguish in the model between beliefs that are held because they were adopted in response to some question and those that do not. The reason is that as soon as a belief becomes incorporated via expansion, the question it answers is removed without a trace from the research agenda. From that point onward, the belief is on a par with beliefs that are held habitually or for no particular reason at all. This suggests that a more adequate model should keep track not only of questions in need of answers but also of questions that have already been answered. Beliefs that answer questions have a special status. It is natural to think of them as having a higher degree of informational value than other beliefs. Perhaps they are also in general more entrenched than other beliefs are. In any case, the special status of question-answering beliefs should arguable be reflected in a formal model.⁷

We will now give an alternative definition of expansion that solves both these problems in one swoop. As a preliminary, we define $Q/_K \alpha = \{\beta \in Q \mid K \cup \{\alpha\} \text{ does not entail } \neg\beta\}$. We call $Q/_K \alpha$ the *K-truncation of Q by α* . The K -truncation of Q by α is the result of deleting all answers from Q that are incompatible with K plus α . Here are some elementary properties of truncation:

(T1) If Q is a K -question and K entails α , then $Q/_K \alpha = Q$.

(T2) If Q is a K -question, then $Q/_K \alpha$ is a $K \cup \{\alpha\}$ -question.⁸

⁷ It is tempting to identify the set of beliefs which answer some question with the agent's belief base. This would be an alternative way to think about the nature of belief bases (cf. Hansson, 1989). There is also connection here to Levi's postulate of weak monotonicity which states that a logically stronger set of beliefs is at least as valuable as a weaker set (1991, 82). This would follow directly if we take informational value as an indication of belief state's capacity to answer the inquirer's questions. The question-answering power of a stronger set will be no less than that of a weaker set. However, more information does not necessarily mean greater power to answer questions. For the residual information of the stronger set may not correspond to any questions entertained by the inquirer. This, in turn, is related to the status of the Recovery postulate in the AGM theory. The Recovery postulate says that a contraction by α followed by an expansion by α should result in status quo: no beliefs should be lost in the process. But presumably beliefs can be lost in the contraction step, beliefs that are not worth reinstating in the expansion step that follows. We are thinking of beliefs that fail to answer any interesting questions.

⁸ For $Q/_K \alpha$ to be a $K \cup \{\alpha\}$ -question, the presuppositions must be fulfilled. The first presupposition states that the exclusive disjunction of the elements of $Q/_K \alpha$ must be entailed by $K \cup \{\alpha\}$. We know that the exclusive disjunction of Q is entailed by K , and because our logic is monotonic, it is also entailed by

(T3) If Q is a K -question, $\beta \in Q$ and $\beta \in \text{Cn}(K \cup \{\alpha\})$, then $Q/K\alpha = \{\beta\}$.

We can now define expansion, more plausibly, as follows:

(D2) $S+\alpha = \langle \text{Cn}(K \cup \{\alpha\}), E', A' \rangle$, where $A' = \{M \mid M = Q/K\alpha \text{ for some } Q \in A\}$.

According to this definition, the result of expanding epistemic state S by α , as far as the research agenda is concerned, is that all questions on S 's agenda are modified so that all answers that are incompatible with the new belief state are deleted.

Let us see how this solves the problems. First of all, the new state after expansion is guaranteed to be a new epistemic state since, by T2, all questions on the new agenda will be of the appropriate kind, i.e., $K \cup \{\alpha\}$ -questions or, equivalently, $\text{Cn}(K \cup \{\alpha\})$ -questions. Second, the new definition makes it possible to distinguish between beliefs that are accepted as answers to some question and beliefs that are not. Suppose that, while in epistemic state $S = \langle K, E, A \rangle$, we expand by α where $\alpha \in Q \in A$, that is to say, α is an answer to the K -question Q on the agent's agenda. This will result in a new epistemic state $S' = \langle K', E', A' \rangle$ where $K' = \text{Cn}(K \cup \{\alpha\})$. By D2, $\{\alpha\} = Q/K\alpha \in A'$. In other words, a belief α that was accepted as the answer to some question is distinguished by the fact that its singleton set, $\{\alpha\}$, is subsequently on the research agenda. Let us call such singleton questions *settled*. Hence, a belief α that was accepted as the answer to some question is distinguished by the fact that it is now the answer to a settled question.

3.2. Non-Quinean contraction

We will not venture a full-blown definition of contraction. Rather, we will be content with proposing some reasonable postulates. The first is trivial:

(CP1) $S \div \alpha$ is an epistemic state.

$K \cup \{\alpha\}$. From that exclusive disjunction and the negation of the element removed from Q when creating $Q/K\alpha$, we can infer an exclusive disjunction of the elements of $Q/K\alpha$. No proper subset of $Q/K\alpha$ can be entailed by $K \cup \{\alpha\}$, because then some element of $Q/K\alpha$ would have to be false relative to $K \cup \{\alpha\}$, but all such elements were removed in the K -truncation of Q by α . The second presupposition, which says that no proper subset of a question can be a question, is thus also satisfied.

The effect of contraction on the research agenda is quite the opposite to that of expansion. Where expansion closes questions, contraction opens them up. This is reflected in the following proposal for a contraction postulate:

(CP2) $S \dot{-} \alpha = \langle K \dot{-} \alpha, E', A' \rangle$, where $Q \in A'$, for some $Q \in Q_{K \dot{-} \alpha}$ such that $\alpha \in Q$.

The effect of contracting an epistemic state by a sentence α is that α is deleted from the belief state component. Here, we rely on the reader to pick his or her favorite contraction recipe. More interesting in this context is the fact that contraction affects the research agenda in certain systematic ways. As CP2 postulates, contraction by α has the effect of introducing a new question on the agenda with α as one of its potential answers.

However, if, as was indeed suggested in the previous section, opening up inquiry into the truth of some alternative hypothesis is the only compelling motivation for contraction, what we need is an operation of contraction by α *with respect to some alternative hypothesis* β . Let us denote the result of this operation by $S \dot{-}_{\beta} \alpha$. The following postulate reflects the actual motivation behind contraction even better than does CP2:

(CP3) $S \dot{-}_{\beta} \alpha = \langle K \dot{-} \alpha, E', A' \rangle$, where $Q \in A'$, for some $Q \in Q_{K \dot{-} \alpha}$ such that $\alpha, \beta \in Q$.

According to CP3, the effect on the research agenda of contracting S by α with respect to β is that a new question is introduced that has both α and β as potential answers.

Obviously, this simple model also allows for simple expansions and contractions of questions from the research agenda. Adding a question corresponds to a new interest in an issue. Removing a question corresponds to losing one's interest. There seems to be no reason to impose constraints on what (unsettled) K-question the inquirer may add to his or her research agenda.⁹ Removing questions, by contrast, is a more sensitive matter. If this paper is on the right track, an inquirer cannot simply remove commitments to conduct further inquiry if those commitments arose in the process of contracting or in that of adding an auxiliary hypothesis.

In addition to giving rise to new questions, a contraction can render an old question vacuous by weakening the belief state so that it no longer implies that exactly one of its potential answers is true. For instance, the question whether god is evil or benevolent presupposes that there either exist a god that is evil or one that is benevolent. Take away the belief that god exists and the

⁹ An agent must not be allowed to add settled questions to the agenda on pain of making impossible a formal distinction between those beliefs that have been added in response to unsettled questions and those that have not.

presupposition of the question will be removed, leaving the question itself hanging in the air. One belief may contribute essentially to the derivability of the presuppositions of many questions. Contracting by such a belief pulls away the carpet, as it were, under many questions and has repercussions throughout the research agenda. There is, however, no need to add an extra postulate guaranteeing that questions for which the presuppositions are no longer derivable must be removed from the agenda. This is already guaranteed by CP1 and Postulate 1 in combination. CP1 states that the effect of contraction is a new epistemic state, and Postulate 1 dictates that the new epistemic state must have a research agenda that is appropriate for that state. Appropriateness implies, among other things, that it must follow from the current state of belief that exactly one of the answers is true.

Here is a final point. We have seen that in the case of inconsistency, epistemic change is not only obligatory; it is even urgent. Something must be done in order to get rid of contradiction and it should be done as soon as possible. This means that introducing inconsistency should be accompanied by updating the research agenda with a commitment to further inquiry where that commitment should be assigned highest possible priority. In order to capture this formally, we would have to assign different priorities to different questions on the research agenda. This is missing from present sketch of a formal model for non-Quinean belief change. In an improved model, one might try to define degree of entrenchment/informational value in terms of the priority of questions. The guiding idea is that a belief α is more entrenched than a belief β just in case the question answered by α is more important (has higher priority) than the question answered by β . There is an obvious complication, though: one single belief may conceivably answer many different questions. The degree of entrenchment/informational value should presumably be a function not only of the priorities but also of the sheer number of questions answered by the belief.

4. Comparison with Levi's theory

We are, to be sure, not the first to argue along these lines against what we have identified as a Quinean dogma. A similar criticism is implicit in Isaac Levi's writings on belief change. In Levi's view, falling into inconsistency as the result of making a belief-contravening observation necessitates "coerced contraction". This type of change he describes as follows:

When the inquirer expands inadvertently into inconsistency via routine expansion (normally because observation, experimental data, or reliable testimony contradicts background knowledge), *the inquirer should contract to extricate himself from inconsistency*. Contraction is coerced by the inadvertent stumbling into inconsistency. (Levi, 1991, 152, our italics)

Thus, allowing oneself to fall into inconsistency involves, in Levi's view, a commitment to contract. An inconsistency is not something we can rest content with. From what Levi writes on the matter, it is clear that he takes contraction from inconsistency to be not only obligatory but also urgent.

Removing some belief in the absence of inconsistency is called "uncoerced contraction" by Levi. He makes the following remark on the rationale for engaging in such contraction:

... contraction does not import error from the inquirer's point of view. It incurs a loss of information. But if this loss is accompanied by the opportunity to undertake an inquiry to settle the question as to the truth of the rival theories without begging the question, there may be a benefit in doing so to compensate for the initial loss of information. If this is right, a good reason for implementing an uncoerced contraction would be that it allows a promising theory incompatible with current doctrine to be examined without prejudice. (Levi, 1991, 153)

As this paragraph suggests, removing a belief in the absence of inconsistency is legitimate, in Levi's view, to the extent that it involves taking on a commitment to examine a promising theory incompatible with the current doctrine; it is not OK, otherwise.¹⁰

Interestingly, however, Levi thinks that there is a further type of situations that tends to generate commitments to carry out further investigation. Consider the following passage from Levi and Morgenbesser (1978):

Dispositions predicates, like *ceteris paribus* clauses, function as place-holders for predicates specifying conditions in generalized statements. But they are not simply place-holders and differ from *ceteris paribus* clauses in a number of respects: *Ceteris paribus* clauses entail no commitment as to the kinds of predicates to be employed in replacing them; disposition predicates do. Hence, a given disposition predicate may be a surrogate for standing conditions in many different laws; and unlike *ceteris paribus* clauses, *this involves a commitment (which may prove erroneous but which may have at the moment some inductive backing) to replacing the disposition predicate by the same specification of standing conditions in each of the laws in which the disposition predicate appears*. Thus, if magnetism is a disposition, it is a disposition to attract and repel iron filings and to induce electricity in electric coils. The reference to magnetism rather than the use of a *ceteris paribus* clause in explaining these two phenomena implies that a certain standing condition is the same in explanations of both phenomena. (Levi and Morgenbesser, 1978, 401, our italics)

¹⁰ We have not found any textual evidence suggesting that Levi would endorse our view on the legitimacy of introducing auxiliary hypotheses. Our view, again, is that such introduction is legitimate to the extent that a commitment is undertaken to search for independent evidence for the hypothesis. There is a potential tension between this proposal and Levi's general epistemological outlook. Levi subscribes to the Peircean belief-doubt model according to which a state of full belief is an idle state entirely free of doubt. An inquirer who accepts an auxiliary hypothesis as true has, on this model, no reason to inquire further into the matter. Hence, he or she has, it seems, no reason to search for independent evidence.

According to Levi and Morgenbesser, accepting a law containing a *ceteris paribus* clause involves a commitment to engage in further inquiry with the goal of replacing the *ceteris paribus* clause. Consider for instance the law “Objects break when tapped lightly *ceteris paribus*”. On the present view, an agent can rationally accept this law as true only if he or she promises to make further inquiries aiming at eliminating the *ceteris paribus* part. The same goes for laws that make use of disposition predicates. The introduction of a law containing a disposition predicate involves a commitment to engage in further inquiry. Consider for example the law “Fragile objects break when tapped lightly”. Acceptance of this law entails a commitment to investigate the finer structure of materials. There are some differences as to what commitments are involved in the two cases (*ceteris paribus* clauses and disposition predicates), but for the purposes of the present discussion we set these difference aside.

While Levi is admirably sensitive to the limitations of the Quinean style of belief representation in his informal motivations, this awareness is not reflected in his (semi-) formal model, which is essentially of the Quinean sort. Thus, epistemic states are taken to consist essentially of a state of full belief and a measure of informational value that plays roughly the same role as the entrenchment ordering in the AGM model, i.e. it is used for evaluating, among other things, the different contraction alternatives open to the agent at a given point in his or her investigations.¹¹ In Levi’s official theory, commitments to inquire further into certain matters play no identifiable role. In this paper, we have taken the first steps toward assigning the research agenda a formally precise role in epistemic change.

5. Conclusion

In this paper we argued, in opposition to Quinean tradition, that the role of the research agenda cannot be neglected in an adequate theory of epistemic change. There are several types of change whose rationality depends crucially on how the research agenda is managed. Adding an auxiliary hypothesis, for instance, is arguable legitimate only if a commitment is undertaken to search for independent evidence for the hypothesis. By the same token, contraction, if it is to be a virtuous type of epistemic change, involves updating the research agenda in certain ways as part of the contraction process. If the argument of the paper is sound, a person’s questions and practical

¹¹ Levi’s formal or semi-formal model includes, besides the belief state and a measure of informational value, several other interesting components. Yet none of these further components plays the role of keeping track of the inquirer’s commitments to engage in further inquiry. Levi may contest this description by pointing out that what he calls the ultimate partition plays exactly this role. The ultimate partition is the set of potential answers to the inquirer’s question at hand. Even if this is true, Levi’s theory is silent regarding the dynamic aspects of the research agenda, i.e., regarding how new ultimate partitions are created or removed in the course of the inquirer’s investigations.

interests, on the one hand, and her beliefs and theoretical values, on the other, are more tightly interwoven than has previously been assumed to be the case in formal epistemology.

We also claimed that it is possible to model formally such non-Quinean belief change. In order to substantiate this contention of ours, we constructed a simple formal model which contains as a part a formal representation of the research agenda. We also made some suggestions as for how the agenda is affected in expansion and contraction.

We have already seen that the main theme of this paper – the role of the research agenda in epistemic change – is related to important issues in philosophy of science concerning the status of disposition predicates and *ceteris paribus* laws. We would, finally, like to point out connections to two other areas of philosophy. First, part of the purpose of the present paper is to codify in a precise manner legitimate reasons for contraction or “doubt”. In this regard, our problem is related to the long-standing problem of skepticism and, in particular, to the issue of how far legitimate doubt extends. Second, the argument of this paper has bearing on the problem of demarcating science from pseudo-science. The practice of introducing auxiliary hypotheses is, as we have argued, not objectionable *per se*. It becomes objectionable only if it is not taken to involve commitments to further inquiry. A reasonable conjecture is that one factor that distinguishes science from pseudo-science is how the research agenda is managed. Recognizing commitments to further inquiry, a scientist assigns these commitments high priority in his or her scientific conduct. The pseudo-scientist, by contrast, either does not recognize such commitments at all or, alternatively, he or she does recognize them but assigns them such low priority that they *de facto* play no role in his or her conduct.

References

Alchourrón, C., Gärdenfors, P, and Makinson, D. (1985), “On the Logic of Theory Change: Partial Meet Contraction and Revision Functions”, *The Journal of Symbolic Logic*, Vol. 50, No. 2, (June, 1985), 510-530.

Gärdenfors, P. (1988), *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, The MIT Press.

Hansson, S. O. (1994), “Taking Belief Bases Seriously”, in *Logic and Philosophy of Science in Uppsala*, D. Prawitz and D. Westerståhl (eds.), 13-28.

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

Levi, I. (1991), *Fixation of Belief and Its Undoing: Changing Beliefs Through Inquiry*, Cambridge University Press.

- Levi, I., and Morgenbesser, S. (1978), "Belief and Disposition", *Dispositions*, ed: R. Tuomela, Dordrecht: Reidel, 389-410. Also in (1964) *American Philosophical Quarterly*, Vol. 1, 221-232.
- Lindström, S., and Rabinowicz, W. (1991), "Epistemic entrenchment with incomparabilities and relational belief revision", *The Logic of Theory Change*, A. Fuhrmann and M. Morreau (eds.), Springer Verlag, 93-126.
- Olsson, E. J. (2005), "Levi and the Lottery", in *Knowledge and Inquiry: Essays on the Pragmatism of Isaac Levi*, Cambridge University Press.
- Quine, W. V. O. (1953), "Two Dogmas of Empiricism", in *From a Logical Point of View*, Cambridge, Mass.: Cambridge University Press.
- Rott, H. (2001), *Change, Choice, and Inference: A study of Belief Revision and Nonmonotonic Reasoning*, Oxford University Press.
- Seegerberg, K. (1995), "Belief Revision from the Point of View of Doxastic Logic", *Bulletin of the IGPL*, Vol. 3, 534-53.
- Seegerberg, K. (1998), "Irrevocable Belief Revision in Dynamic Doxastic Logic", *Notre Dame Journal of Formal Logic*, Vol. 39, No. 3, 287-306
- Spohn, W. (1988), "Ordinal Conditional Functions. A Dynamic Theory of Epistemic States", *Causation in Decision, Belief Change, and Statistics*, Vol. 2, ed: W.L. Harper, B. Skyrms, Kluwer, Dordrecht, 105-134.