Ten philosophical problems in belief revision

Sven Ove Hansson

March 2, 2002

Abstract

The paper introduces ten open problems in belief revision theory, related to the representation of the belief state, to different notions of degrees of belief, and to the nature of change operations. It is argued that these problems are all issues in philosopical logic, in the strong sense of requiring inputs from both logic and philosophy for their solution.

1 Introduction

The AGM model of belief change [1] is an unusually simple and elegant theory. Not surprisingly for such a highly simplified theory, there are many features of actual belief systems that it does not capture. Unfortunately, much of its mathematical beauty is lost when it is subjected to various amendments and extensions in order to make it more realistic. The theory is in a sense aesthetically closed.

This aesthetic closure may be a major reason why several able researchers have, in private conversations, confided to me that they believe the subject may be more or less finished since there is nothing more important to add to it. In my view, this is very far from the truth. On the contrary, there are a large number of important open problems in this field of research. Some of these may require the introduction of some measure of inelegance into the formal framework, but mathematical elegance should not be an absolute priority.

The purpose of this paper is to introduce ten philosophical problems in belief revision that I believe to be worth further study.[1] My general approach to belief revision theory is that it should aim at providing us with models of changes in the beliefs of actual agents, that may be human beings or computers.

[1] The word "revision" is commonly used in two senses in this context: (1) change, in general, and (2) consistency-preserving acceptance of new information. The whole field of research is mostly called "theory change" or "belief revision".

2 Choosing a level of idealization

Actual processes of belief change are extremely complex. In order to obtain a model that is at all manageable, substantial simplifications are necessary. In other words, a useful formal model of belief change has to be idealized. But what exactly does it mean to say that these models are idealized?

The words "ideal" and "idealize" have two senses. First, to idealize may mean to perform a simplification for the sake of clarity. In this sense, to say that a model of the belief state is an ideal means that it is "[s]omething existing only as a mental conception". (OED) To idealize in this sense means to perform a "deliberate simplifying of something complicated (a situation, a concept, etc.) with a view to achieving at least a partial understanding of that thing. It may involve a distortion of the original or it can simply mean a leaving aside of some components in a complex in order to focus the better on the remaining ones."[34]

Secondly, by idealization we can mean the act of entertaining or expressing a (too) high opinion of something. Formal models may or may not represent something that is "perfect or supremely excellent in its kind". (OED)

Formal models of belief revision, such as that used in AGM, have been obtained through both types of idealizations. These models are both (1) idealizing-simplifying, i.e. they leave out many of the complexities of real life, and (2) idealizing-perfecting, i.e. they represent patterns that satisfy standards of rationality that are higher than those that actual (doxastic) agents usually live up to. In many cases, idealizing-perfecting also leads to idealizing-simplifying, so that the two forms of idealization may coincide. Nevertheless, they are conceptually distinct. In order to better understand the models we are working with, we need to distinguish between the two forms of idealization.

The level of idealization-perfection is an important characteristic of a system of belief revision. Some researchers seem to conceive the ideal rational agents of belief change theory as having unlimited cognitive capacity. It is then fairly unproblematic to construct formal models in which these agents have to process infinite entities (such as infinite sets of sentences). This is the most common approach, and mathematically the least cumbersome one. But is it the approach from which we can learn the most about real-world rationality? An alternative view is that the ideal agents of belief change theory should have limited cognitive capacity, of which they make rational use. On that view, finiteness and computability are important *desiderata* of formal models. This is an argument for the use of models with finite representation of the belief state, such as finite belief bases. However, finiteness is only a very weak restriction. It takes us in the direction of a realistic approach to cognitive capacity, but it does not take us far enough. This leads to the first of my ten questions:

1. Can stricter cognitive limitations than finiteness be represented in an interesting way?

3 The object language

In a model of change there must be something that changes. In belief change theory, that object of change is called the belief state. However, the very idea of a belief state is in itself an idealization. It artificially isolates beliefs from other constituents of a state of mind such as emotions and preferences. The idea of an isolable belief state requires, among other conditions, that the fact-value distinction can be drawn with perfect precision.

The available belief change models are sentential: beliefs are represented by sentences. This, too, is clearly an idealization. Actual beliefs do not necessarily have the structure of sentences in a language. However, although sentences do not capture all aspects of beliefs, they provide the best available general-purpose representation of beliefs.

In most belief change models, the belief-representing sentences are assumed to be elements of a simple, truth-functional propositional language. There seems to be a consensus that the addition of quantifiers to the belief-representing language would not provide new insights in proportion to the complications that would ensue. The inclusion of modal or conditional sentences in the language gives rise to interesting puzzles, but also seems to make further formal developments difficult. [9, 14, 18, 32] If we include modal sentences in the belief set, standard conditions on operations of change will no longer hold. This applies, for instance, to the inclusion postulate for contraction ($K \div a \subseteq K$). The sentence "it is possible that a is not the case" may be an element of $K \div a$ but not of K.[2] We seem to "lose the logic" when the language is enriched. This problem is well-known, but certainly nevertheless worth a place on the list of open questions.

2. How can modal and conditional sentences be represented?

4 Degrees of belief

The notion of belief can be conceived as an all-or-nothing concept: either you believe something, or you do not. Alternatively, it may be thought of as admitting of degrees: you may believe something to various degrees. Mainstream belief change models are *dichotomous*: they divide the sentences

[2]There is a sense of possibility according to which one takes something for possible (exactly) when one does not believe in its negation. [6, 29, 30] There is another sense of possibility according to which one believes something to be possible if and only if there is some input (or series of inputs) that would make one believe it. [16]

of the language into two distinct categories, those representing beliefs and those not doing so.[3]

There are two notions of degree of belief. One of these is the static concept of *degree of confidence*. In this sense, your degree of belief in a sentence is higher, the more confidently you entertain that belief. The other notion is the dynamic concept of *degree of resistance to change*. In that sense, your degree of belief in a sentence is higher, the more difficult it is to change that belief.

Standard probability theory makes no difference beween these two notions, so in a sense we are brought up not to distinguish between them in formal representations. However, some hints of this distinction can be found in the belief change literature. Isaac Levi distinguishes carefully between certainty and unchangeability (incorrigibility). [30] This distinction cannot be made in probability theory, since there a proposition is certain if and only if its probability is one. The distinction that I propose, between confidence and resistance to change, can be seen as a generalization of Levi's distinction. Furthermore, the notion of epistemic entrenchment [10] can be seen as an attempt to capture resistance to change. "It is the epistemic entrenchment of a sentence in an epistemic state that determines the sentence's fate when the state is contracted or revised." [10]

In the light of this, the above statement that standard models of belief change are dichotomous can be made more specific: They are dichotomous with respect to degrees of confidence, but not with respect to degrees of resistance to change.

The dichotomous picture is far from uncontroversial. According to the Bayesian ideal of rationality, a rational subject should not have many full beliefs. Only logically true sentences are irretractable and can be assigned probability one. The resulting belief system is a complex web of interconnected probability statements. [24] In practice, however, such a belief system would be unmanageable for human subjects. [33] Our cognitive limitations are so severe that massive reductions from high probability to full belief (certainty) are inevitable in order to make us capable of reaching conclusions and making decisions. In other words, we treat things as certain although they are not. This reduction to full belief, or "fixation of belief" [42], helps us to achieve a cognitively manageable representation of the world.

The prevalence of this reduction (fixation) process is one of the reasons why dichotomous belief models represent some features of doxastic behaviour (notably those related to logic) more realistically than probabilistic models. Clearly, there are other features that can be more realistically represented in the latter models. Note, however, that this argument for dichotomous models refers to their relevance for agents with limited cognitive capacities. The use of dichotomous models in a discussion of (ideal)

[3] On non-dichotomous models see [4, 45, 46].

agents with unlimited cognitive capacities does not seem to be equally well motivated.

What we have, then, are two types of models: (1) Probabilistic models in which no distinction is made between confidence and resistance to change, and both are represented by one and the same probability function, and (2) entrenchment-based and related models in which degrees of resistance to change are represented by the entrenchment relation, whereas there is no distinction made with respect to the degrees of confidence with which beliefs are held. We need to know what is the actual relationship between the two notions of degrees of belief. Intuitively, they seem to be neither completely independent nor reducible into one and the same concept.

3. What is the (formal and informal) relationship between the two notions of degree of belief, confidence and resistance to change?

In what follows we will be concerned with dichotomous, sentential models of belief change. In such models, the relation between belief states and believed sentences can be expressed with a support function s that sorts out the sentences that are supported by the belief state. [14, 20] Let $\mathcal{K}$ be a belief state. Then $s(\mathcal{K})$ is the set of all sentences that are beliefs in $\mathcal{K}$.

5 Vulnerability and justification

It is commonly assumed in belief change theory that logical consequences of beliefs are themselves beliefs, i.e. that $s(\mathcal{K})$ is closed under logical consequence, $s(\mathcal{K}) = \mathrm{Cn}(s(\mathcal{K}))$, where Cn is an operation of logical consequence.[4] This is not a realistic assumption, but it has turned out to be extremely helpful as a means to obtain a manageable formal structure. An interesting argument in its favour was put forward by Isaac Levi [27, 30]; according to him, $s(\mathcal{K})$ should be interpreted as consisting of the sentences that someone is committed to believe, not those that she actually believes in.

The simplest and most obvious representation of belief states is to identify each belief state with its respective belief set, so that $s(\mathcal{K}) = \mathcal{K}$. Then operations of change are performed on the belief set, rather than on some underlying belief state from which it can be derived. This is how the AGM model is often described; the belief set is said to represent or even be equal to its belief state. However, this description of AGM is misleading. In a model of belief change, a static construct such as a belief set is not a sufficient description of the belief state. We need to know not only the beliefs presently endorsed, but also what will be the fate of these beliefs after var-

[4] In logical parlance a closed set is called a theory. In belief change theory, it is called a belief set.

ious operations of change have been performed.[5]

Two major ways to represent this dynamic information have been explored in the literature. One of these is to introduce a direct representation of resistance to change, or conversely, of the vulnerability of the elements of the original belief set, indicating how easily different beliefs are given up. This is the dominant approach in AGM-style belief revision. The idea is, of course, that when choosing which previous beliefs to give up, less vulnerable ones are retained as far as possible. Vulnerability is independent of the particular operations to be performed (and hence of the input sentence). Suppose that there are two operations $\circ_1$ and $\circ_2$, and two sentences a and b that are candidates for being retracted in both of these operations. Then, according to a vulnerability approach, a is more vulnerable than b in operation $\circ_1$ if and only if it is so in operation $\circ_2$. (For concreteness, we may think of $\circ_1$ as the retraction of the belief $a\&b$ and $\circ_2$ as the retraction of the belief $a\&b\&c$.)

The other type of dynamic information relates to the justificatory structure of the belief set. Some beliefs have no independent standing, but are held only because they are justified by some other belief(s). When the justification of a belief has been lost, that belief should arguably also be deleted. There are two major ways to express information about justificatory structure. The simplest of these employs belief bases. A belief base is a set of sentences that is not (except as a limiting case) closed under logical consequence. Its elements represent beliefs that are held independently of any other belief or set of beliefs. The logical closure of a belief base is a belief set. Those elements of the belief set that are not in the belief base are "merely derived", i.e., they have no independent standing. [16] Changes are performed on the belief base, and derived beliefs are changed only as a result of changes of the base.

It must be emphasized that belief bases provide only a very rough approximation to justificatory structure. They capture deductive justifications, which comprise only a fraction of the actual justificatory relationships. A more realistic justificatory structure is contained in what may be called track-keeping representations. [13] Here, to each sentence is appended a list of its justifications or origins. This approach has been much explored by computer scientists, beginning with the truth maintenance systems (reason maintenance systems) developed by Jon Doyle [3]. In correspondence, David

[5] Alvaro del Val has argued that the belief state should not be allowed to contain dynamic information, since there is a "need to separate the specifications of the agent's beliefs from the specification of the agent's revision policy, which are fully orthogonal, independent issues".[47] In my view, these issues are far from orthogonal. There are two major reasons for this. First, some forms of modal and conditional beliefs seem to be strongly connected with how one would change one's (non-modal, non-conditional) beliefs upon receipt of certain inputs. See above, Section 3. Secondly, there are obvious connections between justificatory structure and propensities to change. If p represents my only justification for believing that q, then we should expect q to be lost when p is lost.

Makinson pointed out to me that a distinction can be drawn between actual and potential justificatory structure. The former recalls how a belief actually came to be held, whereas the latter states some possible way(s) of coming to hold the belief. Track-keeping representations seem to be particularly well-suited to represent an actual justificatory structure, but a premise set (such as a belief base) can also be used. To express a potential justificatory structure, a set of possible derivations, or a set of possible premise sets, may be suitable.

The relation between vulnerability and justificatory structure remains an open issue. It is not clear, either on a conceptual or a technical level, to what degree the justificatory structure can be expressed in terms of vulnerability, or vice versa.

4. What is the relation between vulnerability/resistance and justificatory structure?

6 Representing change

Given a formal representation of the belief state, let us now consider how changes in that state can be expressed in the formal framework. In what may be called time-indexed models, a (discrete or continuous) variable is employed to represent time. The object of change (such as a state of affairs, state of the world, or belief state) can then be represented as a function of this time variable. (This framework can also be made indeterministic by allowing for a bundle of functions, typically structured as a branching tree.) This approach has been useful in many areas of research, but it has hardly been used at all in studies of belief change.

Instead, most models of belief change are constructed as *input-assimilating* models. In such models, the object of change is exposed to an input, and is changed as a result of this. No explicit representation of time is included. Instead, the characteristic mathematical constituent is a function that, to each pair of a state and an input, assigns a new state. (There are interesting parallels between belief change theory and automata theory.)

In a well-constructed input-assimilating model of belief change, the representation of a belief state after a change has taken place should have the same format as the representation of the belief state before the change. This has been called the principle of *categorial matching*. [11] As an example, if we begin with a belief base, then the outcome of a change should be a new belief base, not a belief set. Similarly, if the original belief state is a belief set combined with information about the vulnerability of its elements, the new belief state after change should contain the same two constituents (and not, e.g., be a belief set with no accompanying vulnerability information).

Categorial matching may seem to be a very elementary criterion. In actual fact, however, it has turned out to be very difficult to achieve. The original AGM model does not satisfy it (a fact that has often been hidden by the practice of treating the belief set alone as the belief state, instead of the pair consisting of the belief set and the selection function or entrenchment relation). Subsequent research has shown that it is indeed very difficult to achieve categorial matching in a credible way.

> 5. Which is the best way to change the AGM model to achieve categorial matching?

This is, of course, a well-known problem in belief change theory. It is more often referred to as the problem of iterated belief change. Probably, changes in the input representation are necessary for a satisfactory solution.[6]

Input-assimilating models have the advantage of focusing on the causes and mechanisms of change. They exhibit the effects of external causes on systems that change only in response to such external influences (inputs) and are otherwise stable. This makes them tolerably well suited to represent important aspects of changes in human states of mind, and of compartments of mind such as states of belief. At least for some purposes, it is a reasonable idealization to disregard such changes in a person's beliefs that have no direct external causes, in order to focus better on the mechanisms of externally caused changes.

In the presence of conflicting information, selections are necessary. We have a choice between (1) making these selections as part of the operations of change when new information is received, and (2) letting operations of change leave conflicts unresolved, and instead make the necessary selections when information is recovered from the system. [43] There is a trade-off in simplicity between retrieval and change. In the AGM model, the retrieval operation is as simple as possible — it is just the identity operation. The change operations of AGM are more complex. In other models, with a more complex retrieval operation, simpler operations of change may be sufficient. The relations between retrieval and change remain to be investigated, both from a formal and a more philosophical point of view. We do not know whether these approaches are fundamentally different or one of them can in some way be reduced to the other.

> 6. To what extent are retrieval and change operations interchangeable?

[6] Abhaya Nayak [36] has proposed a variant of revision in which the input contains not only a sentence to be incorporated into the belief set, but also its degree of priority in the resulting new belief set. Hence, in his model there are different ways to revise one and the same belief state by one and the same sentence. See also [46] and Hans Rott's contribution to the present issue.

7 The nature of contraction

In the AGM framework, there are three types of belief change. (The basic ideas derive from earlier work by Isaac Levi, [27] and [28].) In contraction, a specified (belief-representing) sentence is removed from the belief set, without anything else being added to it. Hence, contraction of a belief set by a non-tautologous sentence a results in a new belief set, which does not contain a. By expansion is meant that a specified sentence is set-theoretically added to the belief set (without anything else being excluded), and this expanded set is then closed under logical consequence. By revision is meant that a specified sentence is added to the belief set, under the condition that the new belief set be consistent and closed under logical consequence. [1] Since revision is definable in terms of contraction and expansion, $K * a = (K \div \neg a) + a$ (the Levi identity), all belief changes in the AGM model are constructible from two primitive types of operation.

It should be noted that the AGM operations are all sentential[7]: The inputs of contraction, expansion, and revision are taken to be sentences in the belief-representing formal language. This is by no means unproblematic. Actual epistemic agents are moved to change their beliefs largely by non-linguistic inputs, such as sensory impressions. Sentential models of belief change (tacitly) assume that all inputs can, in terms of their effects on belief states, be adequately represented by sentences. When I see a hen on the roof (a sensory input), I am assumed to adjust my belief state *as if* I modified it to include the sentence "There is a hen on the roof" (a linguistic input). [17])

Hence, in AGM and related systems, each operation of change is related

[7]Mukesh Dalal [2] has proposed that systems of belief change should satisfy the principle of irrelevance of syntax. This means the outcome of any operation of change is independent of the syntactic form of either the old or the new information.

In actual systems of belief change, a single sentence may or may not carry information in addition to its propositional contents. (In AGM and related systems, it does not.) Similarly, a set of sentences may or may not represent something more than, or other than, the combined propositional contents of its elements. Hence, the belief bases $\{p, q\}$ and $\{p, p \leftrightarrow q\}$ are expected to have different dynamic properties, although they are statically equivalent. [16] (Some authors, including Schlechta [44], Nebel [37], and del Val [47] seem reluctant to allow for more than one belief base with the same logical closure. However, it is only when this is allowed that the enhanced expressive power of belief bases over belief sets is made use of.) In the same way, in an extension of AGM that allows for contraction by sets of sentences rather than only by single sentences ("multiple contraction"), contraction by different sets with the same logical closure may yield different outcomes. [8]

To programmatically refrain from using sentences or sets of sentences to represent more than their respective propositional contents would mean that one gives up much of the potential expressive power of language-based models. Syntax-independence should therefore not be seen as a criterion of adequacy for language-based models of information-processing, but rather as a property that emerges from some but not all of the idealization processes through which such models are constructed. For a more detailed discussion, see [21].

to a sentence in one of three ways (contraction, expansion, or revision by that sentence). Furthermore, uniqueness holds in the sense that for each of these types of operation, there is exactly one operation for each sentence (indeed, for each class of logically equivalent sentences); hence there is exactly one contraction by each sentence, etc.

Even if we accept the basic underlying idealizations, this typology of change operations is open to criticism of at least two kinds. First, the realism and relevance of the three proposed types of operations can be questioned. Secondly, it may be argued that these three types of operations (or their combinations in sequences) do not cover all the types of belief changes that there are.

The first type of criticism has been directed primarily at the operation of contraction. (The realism of expansion or revision does not seem to have been seriously contested.) In contraction, as conventionally defined, the outcome is a subset of the original belief set, that does not contain the input sentence. Hence, this is an operation in which old beliefs are deleted but no new beliefs are added. It is difficult, however, to find examples of such *pure* contraction, in which no new belief is added. When we give up a belief, this is typically because we have learnt something new that forces the old belief out. For concreteness, suppose that I previously believed that the dinosaurs died out due to sudden climatic change (a). Then a geologist told me that this is only one out of several competing hypotheses. This makes me give up my belief in a (without starting to believe in its negation). Strictly speaking, this is not a case of (pure) contraction, since a new belief was acquired to the effect that there are several competing scientific hypotheses on the extinction of the dinosaurs. In the literature on belief dynamics, examples such as this are often interpreted as referring to (pure) contraction. The new belief that gave rise to contraction is neglected, and is not included in the new belief set. This is an imprecise but convenient convention, that makes it much easier to find examples of contraction.

We sometimes hypothetically give up a belief in order to give a contradictory belief a hearing. Such hypothetical contractions, or contractions for the sake of argument, have sometimes been taken to be pure contractions. [7, 8, 30] However, this interpretation is questionable since these contractions are not seriously undertaken by the agent.

The elusiveness of pure contraction should not lead us to believe that contraction is unimportant. Contraction is an essential element of rational belief change. It typically occurs as a part of more complex changes that involve both losses and acquisitions of information. For the formal analysis, it is useful to develop models of pure contraction. However, this still leaves us with an open question, namely how to deal adequately with the operations that are intuitively taken to be contractions although they are not pure contractions since subsidiary beliefs are added.[8]

7. How should ordinary, non-pure, contraction be represented?

8 Decomposition problems

The other line of attack is that other types of change than those of AGM should be allowed for.

It is a fundamental assumption in belief dynamics — introduced by Isaac Levi [27] — that complex changes can be analyzed as sequences of changes of these simple types:

> *Decomposition principle* (Fuhrmann [6])
>
> Every legitimate belief change is decomposable into a sequence of contractions and expansions.

The decomposition principle need not be read as a requirement that you actually change your beliefs in this stepwise fashion: one expansion, or contraction at a time. All that is required is that the outcomes of complex changes are the same *as if* you had performed them in this way.

Several additional categories of operations have been proposed. In *multiple contraction* and *multiple revision*, the input consists of sets of sentences rather than single sentences. [5, 8, 12, 31, 38] In *updating*, the change takes part in the real world, rather than in the agent's beliefs about an unchanging

[8] Pure contraction can be conceived as taking place on an important fraction of the belief set, such as the fraction that does not contain any beliefs that refer to other beliefs or their justifiability. Let $\mathcal{L}_R$ be a fraction of the language $\mathcal{L}$ such that $\mathcal{L}_R$ consists of those sentences that contain only certain atoms, or are logically equivalent with sentences only containing these atoms. For each set A, let $|A| = A \cap \mathcal{L}_R$.

Let K be a logically closed, consistent belief set in $\mathcal{L}$ and $*$ an AGM revision operator, based on an entrenchment relation $\leq$, whose symmetric part is denoted $\equiv$.) Let f be a function from (a part of) $\mathcal{L}_R$ to $\mathcal{L} \setminus \mathcal{L}_R$ such that

(i) If $a \leftrightarrow a' \in \mathrm{Cn}(\emptyset)$, then $f(a) \leftrightarrow f(a') \in \mathrm{Cn}(\emptyset)$
(ii) If $a \in \mathcal{L}_R \setminus K$, or $a \in \mathrm{Cn}(\emptyset)$, then $f(a) \in K$.
(iii) If $a \in |K| \setminus \mathrm{Cn}(\emptyset)$, then
(iii.a) $f(a) \to \neg a \equiv f(a) \to a$, and
(iii.b) $f(a) \to d \leq f(a) \to \neg d$ for all $d \in \mathcal{L}_R \setminus K$.

Intuitively speaking, $f(a)$ is the sentence that we revise by in order to contract by a. We can then define the operation $|K| \div a = |K * f(a)|$. For $a \in \mathcal{L}_R$ it satisfies the following conditions:

Success: If $a \notin \mathrm{Cn}(\emptyset)$ then $a \notin |K| \div a$.
Inclusion: $|K| \div a \subseteq |K|$.
Closure: $|\mathrm{Cn}(|K| \div a)| \subseteq |K| \div a$.
Extensionality: If $a \leftrightarrow a' \in \mathrm{Cn}(\emptyset)$, then $|K| \div a = |K| \div a'$.
Vacuity: If $a \notin |K|$ then $|K| \subseteq |K| \div a$.
Failure: If $a \in \mathrm{Cn}(\emptyset)$, then $|K| \div a = K$.

world. [25, 26] Operations of *non-prioritized revision* differ from those of conventional revision in that the input sentence is not always accepted. Operations of *consolidation* enhance the integrity of the belief state by making it consistent [16, 19] or coherent [39, 40]. It has been suggested that *abduction*, or inference to the best explanation, should be representable in a system of belief change. [41] *Impure contraction* was just mentioned, and still other operations are also possible. In view of this multiplicity, the issue of decomposition into basic or atomic operation-types is worth a thorough investigation.

8. Are there atomic operations in terms of which all belief changes can be represented?

The last two open issues that I wish to raise are both special cases of the decomposition issue. The first of them is concerned with the decomposition of revision.

The two major tasks of a revision operator $*$ are to add the new belief a to the belief set K, and to ensure that the resulting belief set $K * a$ is consistent (unless a is inconsistent). The first task can be accomplished by expansion by a. The second task can be accomplished by prior contraction by its negation $\neg a$. If a belief set does not imply $\neg a$, then a can be added to it without loss of consistency. An operator of revision can therefore be constructed out of two suboperations. The recipe is as follows:

(1) Contract by $\neg a$.

(2) Expand by a.

or in more concise notation: $K * a = (K \div \neg a) + a$ (the Levi identity).

It is not possible to perform the two operations in reverse order since if $K \cup \{a\}$ is inconsistent, then $K + a$ is always the same (namely identical to the whole language), so that all distinctions between initial values of K are lost.[9]

Just like the corresponding operators for belief sets, revision operators for belief bases can be constructed out of two suboperations: in order to revise by a we expand by a and contract by $\neg a$. In this case, however, it is possible to perform the two operations either in the order of the Levi identity or in the reverse order:

(1) Expand by a

(2) Contract by $\neg a$

[9]For a belief set K, the outcome of expansion is logically closed, $K + a = \text{Cn}(K \cup \{a\})$. For a belief base B, we have instead $B + a = B \cup \{a\}$. The difference is due to the principle of categorial matching.

More compactly, this is expressed by the reversed Levi identity, $B * a = (B + a) \div \neg a$.

The two revision operators for belief bases have been shown to differ in their formal properties [15].[10] They correspond to different intuitions about how belief-contravening information should be accommodated by a rational doxastic agent. Consistency is preserved in every step of contraction+expansion, but there is an intermediate non-committed state in which neither the input sentence a nor its negation $\neg a$ is believed. In (belief-contravening) expansion+contraction, there is instead an intermediate inconsistent state in which both a and $\neg a$ are believed, Which of the two operations is the more plausible? Our intuitions about this seem to differ between different cases:

> (1) Anthony and Beatrice are a married couple. I used to think that they were both Roman Catholics. Then I heard Beatrice say: "In our marriage, it was never a problem that we belong to different denominations." When I heard this, I gave up my belief that Beatrice was a Roman Catholic, but I retained my belief that Anthony was so (since I have seen him enter the local Catholic Church several times).
>
> (2) When Joseph Black learned of the results of Lavoisier's new experiments, he gave up his previous belief in the phlogiston theory of combustion, and accepted Lavoisier's oxygen theory.
>
> (3) I believed that John was dead. Then I met him in the street.

In case 1, expansion+contraction seems to be the most plausible account. More generally: if it is obvious that the new information must be accepted, but less obvious which previous beliefs it should push out, then expansion+contraction seems to be closest to the actual psychological process. In case 2, there was a phase of hesitation in which neither the new belief nor its negation was accepted. Contraction+expansion is closer than expansion+contraction to this kind of process. In case 3, it is difficult to determine which of the two is the most adequate model. Intuitively, the two operations seem to be simultaneous — a feature that is not easy to capture in logical representation.

9. What are the roles of intermediate non-committed and intermediate inconsistent belief states?

[10] For a simple example, let $B = \{\neg p \vee q, \neg p \vee \neg q\}$. It is easy to construct a partial meet operation such that revision via the reversed Levi identity yields $q \in \mathrm{Cn}(B * (p\&r))$ and $q \notin (B * (p\&s))$. If we instead use the original Levi identity, then this is not possible since we have $B \perp \neg(p\&r) = B \perp \neg(p\&s)$, hence for any partial meet operation $\div$ we have $B \div \neg(p\&r) = B \div \neg(p\&s)$, so that it holds for every operation $*$ of partial meet revision that $B * (p\&r)) = B * (p\&s)$.

Finally, let us have a look at the decomposition problem for non-prioritized belief revision, i.e. revision in which the input sentence is not necessarily accepted. One way to construct this operation is to base it on the following two-step process: First we decide whether to accept or reject the input. After that, if the input was accepted, then it is incorporated into the belief state through (conventional) revision.

Decision+revision:

(1) Decide whether the input a should be accepted or rejected.

(2) If a was accepted, revise by a.

The decision+revision model is foreshadowed in some of Isaac Levi's work, but the first fully formalized model of it seems to be David Makinson's [35] screened revision. (See [22] for a detailed study of the formal properties of this construction.)

Another approach to non-prioritized revision is to provisionally accept the new information and, if this led to inconsistency, afterwards regain consistency by throwing out either the input or some of the previous beliefs.

Expansion+consolidation [16, 19]:

(1) Expand by a.

(2) Consolidate the belief state.

where consolidation is a procedure that makes the belief state consistent. This approach has been developed only for belief bases. Consolidation can be defined as contraction by a contradictory sentence. Erik Olsson [39, 40] has developed another variant, in which the consistency requirement is replaced by a requirement of coherence. Another interesting development is to use a "localized" consolidation operator that consolidates only a compartment of the belief base. Contrary to full consolidation, this process will not eradicate all inconsistencies. [23] This is a realistic feature, since in real life inconsistencies are often tolerated, and do not propagate to make the whole belief state degenerate.

10. What is the relation between decision+revision and expansion+consolidation?

These ten questions, and many other open questions in belief change theory, are all issues of philosophical logic, in the strong sense of requiring inputs from both logic and philosophy for their solution.

Acknowledgement I would like to thank Dov Gabbay, David Makinson, and Hans Rott for valuable comments on earlier versions of this paper.

References

[1] Alchourrón CE, Gärdenfors P, Makinson D. "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions". *Journal of Symbolic Logic* 50:510-530, 1985.

[2] Dalal, M. "Investigations Into a Theory of Knowledge Base Revision: Preliminary Report", *Seventh National Conference on Artificial Intelligence* (AAAI-88), 475-479, 1988.

[3] Doyle J. "A Truth Maintenance System". *Artificial Intelligence* 12:231-272, 1979.

[4] Dubois D, Moral S, Prade H. "Belief change rules in ordinal and numerical uncertainty theories". Pp. 311-392 in DM Gabbay and P Smets (eds), *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, vol 3 D Dubois and H Prade (eds.) *Belief Change*, Kluwer 1998.

[5] Fuhrmann A. Relevant Logics, *Modal Logics, and Theory Change.* Doctoral Thesis, Australian National University, Canberra, 1988.

[6] Fuhrmann A. "Reflective Modalities and Theory Change". *Synthese* 81:115-134, 1989.

[7] Fuhrmann A. "Theory Contraction Through Base Contraction". *Journal of Philosophical Logic* 20:175-203, 1991.

[8] Fuhrmann, A. and Hansson, S.O. "A Survey of Multiple Contraction", *Journal of Logic, Language and Information* 3:39-76, 1994.

[9] Gärdenfors P. "Belief Revision and the Ramsey Test for Conditionals", *Philosophical Review* 95:81-93, 1986.

[10] Gärdenfors P. *Knowledge in Flux. Modeling the Dynamics of Epistemic States.* Cambridge, Massachusetts: The MIT Press, 1988.

[11] Gärdenfors, P. and Rott, H. "Belief Revision", pp. 35-132 in Gabbay, Hogger, and Robinson (eds.) *Handbook of Logic in AI and Logic Programming*, volume 4, Oxford University Press, 1995.

[12] Hansson S.O. "New Operators for Theory Change". *Theoria* 55:114-132, 1989.

[13] Hansson S.O. *Belief Base Dynamics.* Doctoral Thesis, Uppsala University, 1991.

[14] Hansson S.O. "In Defense of the Ramsey Test". *Journal of Philosophy* 89(10):522-540, 1992.

[15] Hansson S.O. "Reversing the Levi Identity". *Journal of Philosophical Logic* 22:637-669, 1993.

[16] Hansson S.O. "Taking Belief Bases Seriously". In: Prawitz D, and Westersthl, Dag ed. *Logic and Philosophy of Science in Uppsala*. Uppsala: Kluwer Academic Publishers, 1994: 13-28.

[17] Hansson S.O. "Changes in Preference". Theory and Decision 38:1-28,1995.

[18] Hansson S.O. "The Emperor's New Clothes. Some recurring problems in the formal analysis of counterfactuals." Pp. 13-31 in: Crocco G, Fariñas del Cerro L, Herzig A ed. *Conditionals: from Philosophy to Computer Science*. Oxford: Clarendon Press, 1995..

[19] Hansson S.O. "Semi-Revision". *Journal of Applied Non-Classical Logic*, 7:151-175, 1997.

[20] Hansson, S.O. "Coherentist Contraction", *Journal of Philosophical Logic* 29:315-330, 2000.

[21] Hansson, S.O. "The role of language in belief revision", manuscript.

[22] Hansson S.O. Ferm E. Cantwell J. Falappa M. "Credibility-limited Revision". *Journal of Symbolic Logic*, in press.

[23] Hansson S.O, Wassermann R. "Local change", *Studia Logica*, in press.

[24] Jeffrey R.C. "Valuation and Acceptance of Scientific Hypotheses". *Philosophy of Science* 23:237249, 1956.

[25] Katsuno H., Mendelzon A.O. "On the Difference between Updating a Knowledge Base and Revising it". Pp. 183-203 in Gärdenfors P. ed. *Belief Revision*. Cambridge: Cambridge University Press, 1992.

[26] Keller A., Winslett M. "On the use of an extended relational model to handle changing incomplete information". *IEEE Transactions on Software Engineering* SE-11:7:620-633, 1985.

[27] Levi I. "Subjunctives, Dispositions and Chances". *Synthese* 34:423-455, 1977.

[28] Levi I. *The Enterprise of Knowledge*. Cambridge, Massachusetts: The MIT Press, 1980.

[29] Levi, I. "Iteration of Conditionals and the Ramsey Test", *Synthese* 76:49-81, 1988.

[30] Levi, I. *The Fixation of Belief and Its Undoing*, Cambridge University Press,. Cambridge, Mass. 1991.

[31] Li J. A "Note on Partial Meet Package Contraction". *Journal of Logic, Language and Information* 7:139-142, 1998.

[32] Lindström S., Rabinowicz W. "Conditionals and the Ramsey test". Pp. 147-188 in DM Gabbay and P Smets (eds), *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, vol 3 Didier Dubois and Henri Prade (eds.) *Belief Change*, Kluwer 1998.

[33] McLaughlin A. "Science, Reason and Value". *Theory and Decision* 1:121-137, 1970.

[34] McMullin, Ernan "Galilean Idealization", *Studies in History and Philosophy of Science* 16:247-273, 1985.

[35] Makinson D. "Screened Revision". *Theoria* 63:14-23, 1997.

[36] Nayak, A. C. "Iterated Belief Change Based on Epistemic Entrenchment", *Erkenntnis* 41:353-390, 1994.

[37] Nebel, B. "Syntax based approaches to belief revision", pp. 52-88 in Gärdenfors P. (ed.), *Belief Revision*, Cambridge University Press, Cambridge, 1992.

[38] Niederée R. "Multiple contraction. A further case against Gärdenfors's principle of recovery". Pp. 322-334 in Fuhrmann A, Morreau M ed. *The Logic of Theory Change*. Berlin: Springer-Verlag, 1991.

[39] Olsson EJ. "A Coherence Interpretation of Semi-Revision". *Theoria* 63:105-134, 1997.

[40] Olsson EJ. Coherence. *Studies in Epistemology and Belief Revision*. PhD Thesis. Uppsala: Uppsala University, 1997.

[41] Pagnucco M. *The Role of Abductive Reasoning within the Process of Belief Revision*, PhD Thesis. University of Sydney, Australia, 1996.

[42] Peirce C. "The fixation of belief". PP. 223-247 in Hartshorne C, Weiss P ed. *Collected Papers of Charles Sanders Peirce*. Harvard University Press 1934.

[43] Rott H. *Change, Choice and Inference*, Oxford University Pres, 2001.

[44] Schlechta, K. "Some results on theory revision", pp. 72-92 in Fuhrmann (ed.), *The Logic of Theory Change*, Springer-Verlag, Berlin 1991.

[45] Smets P. "Numerical representations of uncertainty". Pp. 265-309 in D. Gabbay and P Smets (eds), *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, vol 3 D Dubois and H Prade (eds.) *Belief Change*, Kluwer 1998.

[46] Spohn, W. "Ordinal conditional functions: A dynamic theory of epistemic states", in W.L. Harper and B. Skyrms, eds., *Causation in Decision, Belief Change and Statistics*, vol 2, pp. 105-134, 1988.

[47] del Val, A. "Non monotonic reasoning and belief revision: syntactic, semantic, foundational and coherence approaches", *Journal of Applied Non-Classical Logics* 7:213-240, 1997.