Logic and **L**ogical **P**hilosophy
Volume 10 (2002), 45–58

Reinhard Kahle

STRUCTURED BELIEF BASES

1. Introduction

In this paper we discuss a formal approach to belief representation which stores *proof-theoretic* information together with formulae. It is illustrated how this additional information can be used in the context of *belief revision*. The general aims of this paper are the following three: First, we would like to give a *descriptive* approach to belief revision, in contrast to a normative one. Secondly, the given theory should avoid (the consequences of) *logical omniscience* of beliefs. Finally, from a broader point of view, the presented approach can be considered as a case study within the programme of *proof-theoretic semantics*. In this programme, the question is raised whether and how proof-theoretic information can be used as a basis for semantics.

For these purposes we will start with *belief bases*, i.e., sets of sentences, but provide the sentences with additional structure. Therefore, we would like to call the formal representation *structured belief bases*. The structure should collect information how a particular sentence finds its way into the belief. This information is provided by proof-theoretic considerations and will be used for the analysis of belief revision.

In this analysis we will introduce a crucial splitting: The operations of *perception* and *thought*. While the former one is modeled by a (trivial) addition of a new sentence to a structured belief base, together with the information that (the belief expressed by) the sentence was "perceived", the latter one tries to give a fine-grained account to the combinations of beliefs, especially the addition of new *derived* beliefs.

With respect to belief revision these distinction is essential for the case of *contradictory beliefs*. Since we assume that it is not possible to perceive a contradiction directly, the inconsistency of a belief has to be realized by thought. This trivial and seemingly minor observation leads us firstly to the conclusion, that a belief revision is usually caused by the realization of an inconsistency only. In our approach we will model such a realization by a *derivation*. Therefore, secondly, it seems to be quite natural that the sentence used for the belief revision has to be taken from those which are directly involved in such a derivation.

From a technical point of view, our approach has a particular advantage compared with most of the standard belief revision theories based on AGM and related frameworks: The operations defined here are *local* in the sense that the operations will depend on concrete objects, namely proofs, but not on abstract totalities, like all consistent subsets of a given set, which are, in practice, hard to grasp.

Using this locality, finally, the approach ensures a certain understanding of the paradigm of *minimal change*. This paradigm is one of the essential requirements demanded for every theory of belief revision.

The structure of the paper is as follows. In the next section we shortly address the formal background. Then we introduce structured beliefs. In section 4 we discuss how the realization of an inconsistency results in a belief revision. The following section is devoted to a description of the resulting contraction procedure. Sections 6–8 contain brief discussions of the relation to the AGM approaches for belief revision, the problem of logical omniscience and the truth maintenance systems. In the final section we discuss our approach with respect to the aims given at the beginning of this introduction and possible refinements and extensions.

2. The formal background

We will work in classical propositional logic. The terminological and notational conventions are as usual, cf. e.g. [Han98, Neb98]. However, for the proof-theoretic considerations our approach is based on the *derivability relation* instead of the (model-theoretic) consequence relation.

Starting with a set of atomic formulae, the language should be closed under the usual propositional connectives $\neg$ (negation), $\wedge$ (conjunction), $\vee$ (disjunction) and $\rightarrow$ (implication). As metavariables for arbitrary formulae we use Greek letters φ, ψ, $\ldots$. As a distinguished formula we need $\perp$ representing falsity, i.e., $\perp$ implies (semantically as well as syntactically) every formula. Assuming a standard axiomatization of propositional logic in a *Hilbert-style calculus* we have a derivability relation $\vdash \varphi$. For a set of formulae we use calligraphic capitals $\mathcal{A}$, $\mathcal{B}$, ...

A *belief base* is given by a set $\mathcal{A}$ of formulae. In particular, it does not have to be closed under logical consequences. Let us call a belief base $\mathcal{A}$ *inconsistent* if it contains $\perp$, we call it *contradictory* if $\perp$ is contained in the deductive closure of $\mathcal{A}$, i.e., $\perp$ can be derived by use of formulae from $\mathcal{A}$. For a descriptive approach to beliefs, one can hardly avoid to consider the case of contradictory beliefs which are not inconsistent in the defined sense. The easiest example for such a belief could be the set $\{\varphi, \neg\varphi\}$. Such a belief base can be considered as one where the inconsistency is not *realized*.

3. Structured beliefs

A key problem for our analysis of belief revision is the question, how or why a sentence was added to a belief base. Here, we would like to distinguish two reasons:

- Perception and
- Thought.

In the case of perception there could be different sources, like seeing or hearing something. Here, we will not discuss the difference between sources, but assume that all perceived beliefs are incorporated in the belief base in the same way.

Next to perception, there is the possibility that a belief was not directly perceived but *derived* from the given beliefs. In this case we would like to say that the new belief was caused by *thought*.

On the formal side, by a thought we mean a *derivation*. Therefore, given a belief base $\mathcal{A}$, φ can be added as a new belief to $\mathcal{A}$ if one finds a derivation $\mathcal{A} \vdash \varphi$ of φ from $\mathcal{A}$.

But, in contrast to the usual approaches to beliefs, we would like to build in such a "history" of a belief into the representation. Therefore, let us introduce the notion of *structured beliefs*. In the definition we will use B as a metavariable for proofs, in particular B_φ for a proof of the formula φ. Since we work in a Hilbert-style calculus, proofs can be represented as sequences of formulae, including the empty sequence which will be used for perceived formulae.

DEFINITION 1. A *structured belief base* $\mathfrak{A}$ is a finite set of tuples $\langle \varphi, \mathsf{B}_\varphi \rangle$, such that

1. $\mathcal{A}_{\mathfrak{A}} := \{\varphi \,|\, \exists \mathsf{B}. \langle \varphi, \mathsf{B} \rangle \in \mathfrak{A}\}$ is a belief base not containing $\bot$.
2. If $\langle \varphi, \mathsf{B}_\varphi \rangle \in \mathfrak{A}$, then B_φ is either the empty sequence — expressing that φ is a belief caused by perception — or a proof of φ from $\mathcal{A}'_{\mathfrak{A}'}$, the belief base associated with $\mathfrak{A}' := \mathfrak{A} \setminus \{\langle \varphi, \mathsf{B}_\varphi \rangle\}$ where $\mathfrak{A}'$ is a structured belief base.

A tuple $\langle \varphi, \mathsf{B}_\varphi \rangle$ is called *structured belief*.

The condition that $\mathfrak{A}'$ has to be a structured belief base rules out "non-wellfounded" beliefs. Otherwise, for example, $\{\varphi \wedge \psi, \varphi, \psi\}$ could be a belief base for a structured belief base in which every formula is derived by use of the others. Since we consider finite sets of beliefs only, the recursive definition of structured belief bases is obviously well-founded.

With respect to the operation of thought we have to stress the following points: It is not the idea to add to a belief $\mathfrak{A}$ *every* φ which is *derivable* from $\mathcal{A}_{\mathfrak{A}}$, but just those which are actually *derived*. In this view, our approach is *subjective*, i.e., depends on the person who is believing something. Thus, the problem what is or has to be derived is by no means deterministic. And we will not give "rationality criterions" from which one can determinate a particular φ which has to be derived. However, there could be some conditions which at least restrict the possible ones. One of them is, for instance, the condition that φ is indeed derivable from $\mathcal{A}_{\mathfrak{A}}$. Unfortunately, in reality, there will be more than enough examples where people getting by "thought" new "insights" which are not derivable from their beliefs. But let us leave these cases aside.

4. Revision caused by inconsistency

Given the notion of structured belief base and the operations of adding new beliefs by perception or thought we can model beliefs just by use of addition — up to the moment $\bot$ is derived. Thus, we claim that a *revision* of a structured belief base is caused only when $\bot$ is explicitly derived. It was build in in our definition that $\bot$ can not be an element of a belief base. Therefore, if one is deriving it from the given beliefs — *realizing the inconsistency* — this derivation forces a belief revision.

Of course, such a revision should convert the belief in a new, consistent one. Here, we have reached the point where a usual belief revision approach starts. However, in comparison with the standard approaches we have some kind of "history" available: We can ask for the *derivation* of $\perp$. And this derivation will be turned out as the key element in our analysis of belief revision.

Since we use a Hilbert-style calculus for the formal representation, such a derivation can be given as a sequence $\mathsf{B}_\perp = \langle\varphi_1, \varphi_2, \ldots, \varphi_n\rangle$ where φ_n is $\perp$ and for each $1 \leq i \leq n$, φ_i is either an element of $\mathcal{A}_{\mathfrak{A}}$ or derivable from $\varphi_1, \ldots, \varphi_{i-1}$. Since we have derived $\perp$ the set of formulae $\{\varphi_1, \varphi_2, \ldots, \varphi_{n-1}\}$ is a contradictory set. If we leave out the formulae derived within the proof, it is clear that also the set $\mathcal{C}_\perp = \{\varphi_i \,|\, \varphi_i \in \mathsf{B}_\perp \text{ and } \varphi_i \in \mathcal{A}_{\mathfrak{A}}\}$ is contradictory.

It is possible that there are other proofs of an inconsistency using other formulae of $\mathcal{A}_{\mathfrak{A}}$, but, in our terminology, such ones are not realized, at least not at the same moment. Also, there could be formulae in $\mathcal{C}_\perp$ which are not really involved in the proof of $\perp$. That means, $\perp$ could be provable from $\mathcal{C}_\perp \setminus \{\chi\}$ for a $\chi \in \mathcal{C}_\perp$. But, again, we would like to say that such a simplified proof of the inconsistency is not realized.

Now, the proposed belief revision has to choose an element ψ of $\mathcal{C}_\perp$ and to "remove" it (more exactly the tuple $\langle\psi, \mathsf{B}_\psi\rangle\rangle$) from $\mathfrak{A}$. If we define this removal by the simple contraction $\mathfrak{A} \setminus \{\langle\psi, \mathsf{B}_\psi\rangle\}$, we have at least destroyed the given proof of $\perp$. However, this is a quite weak result. Of course, there are at least two other desirable conditions:

($\star$) There should be no other proof of $\perp$,

($\star\star$) All consequences added to $\mathfrak{A}$ because of the presence of ψ should be removed, too.

Both conditions are quite problematic, and it turns out that we cannot really hope for a full realization of them. But before we turn to these considerations let us briefly discuss the significance of $\mathcal{C}_\perp$.

We have to emphasize that the restriction to the set $\mathcal{C}_\perp$ is indeed an important step. Compared with the structured belief base $\mathfrak{A}$, and the "unstructured" belief base $\mathcal{A}_{\mathfrak{A}}$, the set $\mathcal{C}_\perp$ could be much much smaller.

From a conceptional point of view, the biggest advantage is probably, that $\mathcal{C}_\perp$ can be obtained by a *local* operation, just considering a single, given derivation. This locality, which avoids totalities like the deductive closure or the intersection of all consistent subsets, is a major goal of our analysis. As far as possible, the defined operations should depend on existential conditions which can be exemplified and not on universal conditions which would require to grasp complex totalities. Moreover, $\mathcal{C}_\perp$ contains already some rather concrete information. We have $\mathcal{C}_\perp \vdash \perp$. Let $\mathcal{C}_\perp$ consist of the formulae $\psi_1, \ldots, \psi_m$. Thus, we have $\vdash \psi_1 \wedge \ldots \wedge \psi_m \to \perp$ which is equivalent with $\vdash \neg\psi_1 \vee \ldots \vee \neg\psi_m$. From this elementary equivalence, we get the justification to pick one of the ψ_i for the belief revision.

Note, that we only suggest to *remove* the chosen formula ψ (together with its proof) from $\mathfrak{A}$. Alternatively, we could also require, to add the negation of ψ to the new belief. Both variants could be maybe distinguished as a *weak* and a *strong* revision. But, as long as only ψ is removed, $\neg\psi$ is derivable from $\mathcal{C}_\perp \setminus \{\psi\}$, following trivially from the fact that $\mathcal{C}_\perp$ is inconsistent (at least if we work with classical logic). Thus, in accordance with our general analysis the addition of $\neg\psi$ has to be carried through an additional thought. A more substantial reservation against strong revision comes with the following considerations. The fulfilling of the two additional conditions ($\star$) and ($\star\star$) given above could result in further contractions, even with other elements of $\mathcal{C}_\perp$. In such a case, at the end of the revision procedure we could have the situation that ψ and χ have to be taken away from the original structured belief base. But in this case, we only get the information $\neg\psi \vee \neg\chi$ from the inconsistency of $\mathcal{C}_\perp$. From this point of view, a strong revision, realized by adding $\neg\psi$ *and* $\neg\chi$, could possibly too strong.

5. Let us contract

Now we would like to discuss a contraction operation on $\mathfrak{A}$ which fulfills, at least partially, the two conditions ($\star$) and ($\star\star$) given above.

For the contraction one should choose a formula ψ of $\mathcal{C}_\perp$ such that $\mathcal{D} := \mathcal{C}_\perp \setminus \{\psi\}$ is, hopefully, no longer contradictory. But even if $\mathcal{D}$ is still contradictory, there is the question whether the inconsistency is realized. If not, one would probably proceed by contracting ψ. But if an inconsistency of $\mathcal{D}$ is realized, i.e., if one finds by thought a proof $\mathcal{D} \vdash \perp$, the whole revision should start with $\mathcal{D}$ instead of $\mathcal{C}_\perp$. That means, when we have made a choice of ψ in $\mathcal{C}_\perp$, either an inconsistency of $\mathcal{C}_\perp \setminus \{\psi\}$ is realized and we go on with a revision based on this set, or it is not realized and we start a contraction with ψ. Therefore, the condition ($\star$) is only fulfilled in the way that no other *realized* inconsistencies will be left, but not that we end up with a consistent belief in any case. Additionally, we have to think of inconsistencies derivable with other formulae in $\mathcal{A}_{\mathfrak{A}}$. But, again, we will say that such inconsistencies are not realized.

Now, let us first discuss the case that the chosen formula ψ was not perceived but added to $\mathfrak{A}$ by thought. In this case, it seems to be quite inadequate to start the contraction with this formula. Instead, one has to choose another formula from the derivation of ψ. This can be done in the same way as before using the information attributed to the formulae. This procedure has to be iterated up to the moment, we have finally chosen a formula which was once added to $\mathfrak{A}$ by perception.

Thus, let a perceived formula ψ of $\mathfrak{A}$ be chosen which is considered as responsible for the realized inconsistency of a $\mathfrak{A}$. As discussed above, in the first step we contract ψ only. In particular, new beliefs should be added only afterwards by thought. But, secondly, we have to try to fulfill the second conditions ($\star\star$) demanding that formulae which "come from ψ" have to be removed form $\mathfrak{A}$ too. In principle, we could now go on with the standard contraction operations worked out in the AGM context, cf. e.g. [Fuh97]. However, then our analysis would just help to single out a particular formula ψ from a inconsistent belief. But, the concept of structured belief base even allows to give a more perspicuous procedure for the contraction. We have just to check those formulae which entered the belief by a derivation which used ψ. Thus, from the (intermediate) new belief $\mathfrak{A} \setminus \{\langle\psi, \mathsf{B}_\psi\rangle\}$ we have to contract the elements of the set $\{\chi \mid \langle\chi, \mathsf{B}_\chi\rangle \in \mathfrak{A} \wedge \psi \in \mathsf{B}_\chi\}$. Of course, this procedure has to be iterated. Working with finite beliefs this algorithm has to stop and we have a new belief in which the proof of the realized inconsistency does not any longer exists, and in which together with a former perceived formula ψ all its derived beliefs are removed.

In general, the procedure of contraction can be refined in several respects, and we will address some ideas in the discussion below. However, here it was our aim to illustrate how one can use the additional information provided by structured beliefs. In this view, the given analysis has to be understood as a qualitative example, only.

6. AGM and related approaches

For belief revision the framework of AGM named after the authors Alchourrón, Gärdenfors and Makinson of the seminal paper [AGM85] is nowadays chosen as a standard approach. The study of belief bases goes back to Alchourrón and Makinson [AM82] and was extensively discussed by Hansson in his dissertation [Han91], cf. also [Han98, Neb98, Han99].

In some sense, our operations are closely related to contraction operations known in the AGM context. In terms of Hansson our operation is an instance of *external revision* as it is described in [Han91, p. 5]: "The first of these is *external revision*, in which the incorporation of a new belief takes place as follows: (1) add the new belief to the belief base, and (2) if the belief base is now inconsistent, make it consistent by the rejection of some old belief(s)." Also, the procedure given here does not exclude *non-prioritized reception of epistemic input*, i.e., it allows that the sentence considered as responsible for an inconsistency could be the last perceived one.

The choice of a sentence used in a contraction of a contradictory belief is somehow related to the procedure of *safe contraction*, introduced by Alchourrón and Makinson [AM85, AM86], but operating on belief bases, cf. [Fuh91, Nay94]. However, we have to emphasize, that the formalizations of safe contraction require model-theoretic considerations. As noted, our approach is restricted to a purely proof-theoretic view of belief revision. In a certain sense, our approach can be considered as a special form of safe contraction (on belief bases), when the ordering relation used in safe contraction is built by use of the proof-theoretic information.

7. Logical omniscience

Up to now, we have not said anything about the problem of logical omniscience. However, it should become clear that the given approach does not respect logical equivalence in any form. If a formula φ is involved in a particular proof B this does not imply that an equivalent formula φ' is involved in the proof. Moreover, since we demand that all beliefs have to be derived from perceptions, it is clear that equivalent formulae are added to a belief only, if the equivalence is derived. And it is quite clear that one will not derive *all* logical consequences of a given set of formulae. From this point of view, one of the main objection against theories of beliefs which are deductively closed, namely the impossibility to deal with mathematical sentences, does not apply here.

The problem of logical omniscience in the usual approaches to belief revision and knowledge representation is well-known and well-discussed. For instance, chapter 9 of [FHMV95] serves as a good reference for a discussion of alternative approaches in knowledge representation. But all of them, *explicit representation of knowledge*, *nonstandard logic*, *impossible worlds*, *awareness*, and *local reasoning*, are heavily based on the perspective of *Kripke structures*.

The idea of explicit representation of knowledge is "instead of defining knowledge *in terms* of possible worlds, we let knowledge be defined directly. Intuitively, we think of each agent's knowledge as being explicitly stored in a database of formulas." [FHMV95, p. 313]. In its syntactic variant, it comes close to the idea of belief bases. Moreover, Fagin & al. even discuss the possibility of a proof-theoretic extension based on *limited deduction systems*: "As another example, a deduction system might be capable of certain limited reasoning about equality. For example, from $A = B$ and $B = C$, it might be able to deduce that $A = C$; however, given the information that $f(1) = 1$ and that $f(x) = x \cdot f(x-1)$, it might not be able to deduce that $f(4) = 24$. In both of these cases, agents have a base set of formulae and an incomplete set of inference rules" [FHMV95, p. 315f]. From this perspective, our approach can be seen as an extension which shifts the focus to the derivability relation and the information given by it.

From the other approaches, only *local reasoning* is related to our approach. Here, reasoning is divided in different *frames of mind* which are allowed to be incompatible. Nevertheless that beliefs *within a frame* are closed under logical consequence, in general they do not have be closed under conjunction. A very illustrative and adequate example is given by "the two great theories physicists reason with [...] the theory of quantum phenomena and the general theory of relativity" [FHMV95, p. 343]. In the proof-theoretic account this could be reflected by marking the basic beliefs with flags (e.g., a "quantum" and a "relativity" flag). Nevertheless, that the physicists are well aware of the inconsistency of both sets of axioms, they can work perfectly using only axioms from one or the other class and respect carefully that in no derivation axioms from both sets are used.

In general, compared with the several semantic approaches to overcome logical omniscience, our approach gives by the *proof(s)* of a formula an additional tool ad hand which is more fine-grained then the semantic interpretation of a formula, and which has more explanatory power.

8. Truth maintenance systems

There is another framework which is indeed very closely related to our approach: *Truth maintenance systems*. They were introduced by John Doyle [Doy79] in the area of *artificial intelligence* as implementations of problems solvers which can deal with derivations. In general, the idea is closely related to the approach presented here. Together with a formula, the *justification* for the formula is stored. The justification is a representation of its proof consisting of the used axioms and rules. The aim of truth maintenance systems is summarized in a text book on artificial intelligence [RN95, p. 326] as follows: "A *truth maintenance system* or *TMS* is a program that keeps track of dependencies between sentences so that retraction (and some other operations) will be more efficient. A TMS actually performed four important jobs. First, a TMS enables dependency-directed backtracking, to avoid the inefficiency of chronological backtracking. A second and equally important job is to provide *explanations* of propositions. A proof is one kind of explanation [...]. [... T]he third job of a TMS [is]: doing default reasoning. [...] Finally, TMSs help in dealing with inconsistencies. If adding P to the knowledge base results in a logical contradiction, a TMS can help pinpoint an explanation of what the contradiction is." In particular, the last point is directly related to our aims, cf. e.g., the description of the handling of inconsistencies in [McA90]. Although that the principle idea is quite similar to ours, the intention of the formalisms seems to be rather different. While truth maintenance systems are introduced and studied for implementation of problem solvers, we would like to give a *descriptive* approach to belief revision. While the development of truth maintenance systems was dominated by implementation and complexity issues, cf. e.g., the refined *assumption-based TMS* [dK86], we would like to focus on the role of proof-theoretic information given by a derivation which allows for explanations of various phenomena in "natural intelligence". However, it seems to be an omission that the results, discussions and experiences from the field of truth maintenance systems did not find their way into the philosophical discussion (at least to our knowledge) and it appears to be worth to study the relation more closely.

9. Discussion

In this section we will discuss various aspects of the proposed analysis. First of all, we have to stress that the given presentation shows a qualitative analysis, only. The main steps can be summarized as follows:

1. We consider *structured beliefs* by adding information about the *source* of a belief.
2. The two main sources are *perceptions* and *derivations*.
3. A *belief revision* is only enforced when an inconsistency of the beliefs is realized by an actual derivation of $\perp$.
4. A responsible formula ψ has to be chosen from this derivation for the revision procedure.
5. Was ψ a derived formula, one has to choose a responsible formula from its proof, and iterate this procedure up to the case that the chosen formula was perceived.

6. This formula has to be *contracted* from the structured belief base $\mathfrak{A}$. Therefore, we remove ψ and contract hereditary all formulae which were derived by use of ψ.

Each of these steps can probably be refined by its own. But we would like to emphasize on two points: First, the access to an additional information, in particular *proofs* referring to concrete formulae. The second point is the *locality* of the operation: A concrete proof is a finite object and easy to handle.

In addition, it is a standard belief revision paradigm that revisions should be *minimal*. But, both, the set $\mathcal{C}_\perp$ and the contraction operation with a chosen formula, can be considered as minimal *with respect to the given derivation of* $\perp$. Nevertheless that the derivation $\perp$ does not have to be minimal, the set $\mathcal{C}_\perp$ (and therefore the formula for the contraction operation) is obviously not arbitrarily chosen. Finally, the contraction operation removes only formulae from the belief base which are involved in a *realized* contradiction and its consequences (as along as they are realized as consequences, see below).

The proofs involved in our analysis are not only useful with respect to the operation of belief revision. They give the approach the descriptive character, which was one of our initial aims. We claim that indeed the proofs stored in the structured belief bases control the belief revision operation. In addition, they can serve as an appropriate account to *justification*. They *are* the answer, if one is asked "Why do you believe φ?". This fact was also one of the motivations for the set up of truth maintenance systems. So, we can distinguish *knowledge*, classically defined as *true and justified belief*, by judging the given proof of the belief. In particular, we can reject an incidentally true belief φ as knowledge, if the given "proof" is defective. But such an approach to knowledge has to be worked out in more detail elsewhere.

The crucial role of proofs in our account makes clear in which form the approach can be considered as an example within the programme of *proof-theoretic semantics*. There exist a closely related approach to *necessity*, outlined in [Kah99]. For the general programme of proof-theoretic semantics we refer to the forthcoming volume [KSH0x].

In the following we will briefly address some potential refinements and extension of our approach.

First, we have to discuss a seemingly defect of the approach, at least with respect to the desired locality. In the last step of the procedure described above we require that *all* formulae which were derived by use of ψ have to be removed. In principle, this universal quantifier would require to check all elements of $\mathfrak{A}$ with respect to the use of ψ. In accordance with the general approach, we could improve this step by removing only those formulae for which it is *realized* that ψ is involved in the derivation. However, even this is not a local operations since we have to consider the totality of *realized derivations*. The problem could be resolved by extending the structured beliefs and adding a third component to a formula ψ containing a list of all formulae which have used this formula in its derivation (and this is probably the way to deal with it in a refined version of the approach). However, in the given approach one can even use this problem to describe another effect which obviously happens in the procedure of belief revision: One does not really remove all formulae which he has derived from ψ, but only those for which he *realizes* or *remembers* that it was derived by use of ψ. Thus, beliefs contain often remnants of "incomplete" revisions what, obviously, causes troubles when they are used later on in thoughts.

There are some operations we left out of consideration here. One is *forgetting*. The interesting case is that in structured belief base, one can forget not a formula but, maybe, the attached information only. This is probably the case when someone claims: "I know φ, but I don't remember why." Another operation can be described (slightly disrespectfully) by an expression used in computer memory management: *Garbage collection*. In this case, one is just reflecting his beliefs (not necessarily caused by a realization of an inconsistency). During such a reflecting process one can refine the derivations of formulae or derive more abstract one which allows to produce a more "efficient" representation of belief, etc. On that occasion one can even realize inconsistencies or find overlooked formulae from a former belief revision which should be removed.

Finally, a refined approach should probably attribute the perceived beliefs by weights measuring the reliability of the perception. Such weights, which are obviously completely subjective, could play an important role in the choice of the responsible formula for an inconsistency.

In general, the question which information should be collected and how it has to be stored is a topic for itself.

We will close this discussion by addressing two points with respect to the formal representation. First, we have used a Hilbert-style calculus for the representation. But the analysis should not depend on the chosen calculus. The important notion we need is the notion of *use*. It has to be clear what it means that a formula is used in a proof. A discussion of this notion can be found in the proof-theoretic account to *necessity* as it is given in [Kah99]. Probably, Gabbay's framework of *Labelled Deductive Systems* [Gab96] could support the representation of the formal concepts we need in our approach.

The second point concerns the underlying logic. We already addressed the possibility of limited reasoning as mentioned in [FHMV95]. But moreover, when modeling beliefs, there is no hope to assume that people really use logic to derive theirs beliefs. Therefore, on can think of replacing the logical derivability relation $\vdash$ by (probably rather strange) *subjective derivablity* relations $\mid\!\sim$. These could involve the well-known amateur mistakes, e.g., allowing to derive $\neg\psi$ from $\varphi \rightarrow \psi$ and $\neg\varphi$. But this, of course, is a wide sphere outside the scope of this paper.

Acknowledgment

This paper goes back to the workshop *Logic and Logical Philosophy* organized by Heinrich Wansing in spring 2001 at the TU Dresden. I thank the organizers and participants for this stimulating and fruitful meeting. In addition, I have to thank an anonymous referee whose comments on this paper gave rise to substantial improvements.

References

[AGM85] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[AM82] Carlos Alchourrón and David Makinson. On the logic of theory change: Contraction functions and their associated revision functions. *Theoria*, 48:14–37, 1982.

[AM85] Carlos Alchourrón and David Makinson. On the logic of theory change: Safe contraction. *Studia Logica*, 44:405–422, 1985.

[AM86] Carlos Alchourrón and David Makinson. Maps between some different kinds of contraction function: the finite case. *Studia Logica*, 45:187–198, 1986.

[dK86] J. de Kleer. An assumption-based TMS. *Artificial Intelligence*, 28:127–162, 1986.

[Doy79] J. Doyle. A truth maintenance system. *Artificial Intelligence*, 12:231–272, 1979.

[FHMV95] Ronald Fagin, Joseph Halpern, Yoram Moses, and Moshe Vardi. *Reasoning about Knowledge*. MIT Press, 1995.

[Fuh91] André Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20:175–203, 1991.

[Fuh97] André Fuhrmann. *An Essay on Contraction.* Studies in Logic, Language and Information. CSLI Publications, 1997.

[Gab96] Dov Gabbay. *Labelled Deductive Systems, Volume 1*, volume 33 of *Oxford Logic Guides.* Oxford University Press, 1996.

[Han91] Sven Ove Hansson. *Belief Base Dynamics.* PhD thesis, Uppsala University, 1991.

[Han98] Sven Ove Hansson. Revision of belief sets and belief bases. In Didier Dubois and Henri Prade, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Volume 3: Belief Change*, pages 16–75. Kluwer Academic Publishers, Dordrecht, 1998.

[Han99] Sven Ove Hansson. *A Textbook of Belief Dynamics. Theory Change and Database Updating.* Kluwer, 1999.

[Kah99] Reinhard Kahle. A proof-theoretic view of intensionality. In Paul Dekker, editor, *Proceedings of the 12th Amsterdam Colloquium*, pages 163–168. Amsterdam University, 1999.

[KSH0x] Reinhard Kahle and Peter Schroeder-Heister, editors. *Proof-theoretic semantics.* In preparation, 200x.

[McA90] David McAllester. Truth maintenance. In Thomas Dietterich and William Swartout, editors, *Proceedings of the Eight National Conference on Artificial Intelligence*, pages 1109–1116. AAAI Press, 1990.

[Nay94] Abhaya Nayak. Foundational belief change. *Journal of Philosophical Logic*, 23:495–533, 1994.

[Neb98] Bernhard Nebel. How hard is it to revise a belief base? In Didier Dubois and Henri Prade, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Volume 3: Belief Change*, pages 77–145. Kluwer Academic Publishers, Dordrecht, 1998.

[RN95] Stuart Russell and Peter Norvig. *Artificial Intelligence — A Modern Approach.* Prentice Hall, 1995.

REINHARD KAHLE
WSI, University of Tübingen
IFI, LMU München
Centria, DI, Universidade Nova de Lisboa
Quinta da Torre
P-2825-114 Caparica, Portugal
`kahle@di.fct.unl.pt`