Published in R. Brachman, H. J. Levesque, and R. Reiter (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the 1st International Conference (KR'89)*, p. 301–311, Toronto, ON, May 1989. Morgan Kaufmann.

A Knowledge Level Analysis of Belief Revision

Bernhard Nebel*
IBM Germany, Scientific Center
Stuttgart, West Germany
e-mail: nebel@ds0lilog.bitnet

Abstract

Revising beliefs is a task any intelligent agent has to perform. For this reason, belief revision has received much interest in Artificial Intelligence. However, there are serious problems when trying to analyze belief revision techniques developed in the field of Artificial Intelligence on the *knowledge level*. The symbolic representation of beliefs seems to be crucial. The *theory of epistemic change* shows that a partial knowledge-level analysis of belief revision is possible, but leaves open the question of how this theory is related to belief revision approaches in Artificial Intelligence. In particular, it remains an open question whether the results achieved in the knowledge-level analysis are valid. Furthermore, the idea of *reason maintenance*, which is considered to be essential in AI, has no counter-part in the theory of epistemic change. Addressing these problems, it is shown how to reconstruct symbol-level belief revision on the knowledge level.

1 Introduction

Any intelligent agent has to account for a changing environment and the fact that its own beliefs might be inaccurate. For this reason, *belief revision* is a task central for any kind of intelligent behavior. For instance, learning [Diettrich, 1986], diagnosis from first principles [Reiter, 1987], and interpretation of counterfactuals [Ginsberg, 1986] are all activities requiring the revision of beliefs.

In Artificial Intelligence, a number of so-called *truth-maintenance systems* [Doyle, 1979; McAllester, 1982; de Kleer, 1986] were developed which support belief revision. However, the question remains how belief revision can be described on an abstract level, independent of how beliefs are represented and manipulated inside a machine. In particular, it is unclear how to describe belief revision on the *knowledge level* as introduced by Newell [1981]. Levesque and Brachman [1986] demanded that every information system should be describable on the knowledge level without any reference to how information is represented or manipulated by the system. However, this seems to be difficult for belief revision. A large number of authors seem to believe that a knowledge-level analysis of belief revision is impossible [Diettrich, 1986; Fagin *et al.*, 1983; Fagin *et al.*, 1986; Ginsberg, 1986]. Considerations of how beliefs are represented on the *symbol level* seem inevitable for belief revision. Reconsidering Newell's original intentions when he introduced the notion of the knowledge level, we note that the main idea was describing the potential for generating actions by knowledge and not providing a theory of how knowledge or beliefs are manipulated [Newell, 1981]: "...there are not well-defined structural properties associated with access and augmentation." Hence, we may conclude that belief revision is a phenomenon not analyzable on the knowledge level.

However, the *theory of epistemic change* and the *logic of theory change* developed by Alchourrón, Gärdenfors, and Makinson [Alchourrón and Makinson, 1982; Alchourrón *et al.*, 1985; Makinson, 1985; Makinson, 1987; Gärdenfors, 1988; Gärdenfors, 1989], which will be described briefly in Section 2, show that at least some aspects of belief revision can be subject to a knowledge level analysis. Based on some *rationality postulates* any *epistemic change operation* should satisfy, various epistemic change operations on *deductively closed theories* are analyzed—some results of this investigation will be presented in detail in Section 3. This approach, which recently received a lot of interest in the AI community (e.g. [Gärdenfors and

*This paper describes research done at the Technical University of Berlin as part of ESPRIT project 311 and at IBM Germany as part of the LILOG project.

Makinson, 1988; Dalal, 1988]), suffers from some deficiencies, though. It is not clear how the results of the theory of epistemic change relate to belief revision as done in AI. Second, the theory of epistemic change does ignore what is usually called *reason maintenance*. These problems will be discussed in Section 4.

In Section 5, some approaches to belief revision in AI and database theory are presented. These will be analyzed by adapting the rationality predicates of the theory of epistemic change, which leads to the conclusion that these approaches satisfy most of the basic rationality postulates.

Based on that, in Section 6, an explicit reconstruction of *symbol-level* belief revision in terms of the theory of epistemic change is given—showing that a *knowledge level* analysis of belief revision techniques as developed in AI is indeed possible. It also shows that reason maintenance needs not to be integrated as a primitive notion in any theory of belief revision, but that it results as a side-effect of the reconstruction, contrary to the opinion most authors seem to have (cf. [Ginsberg, 1986; Gärdenfors, 1988]).

Finally, in Section 7, we will refine the reconstruction in order to satisfy all rationality postulates—leading to a belief revision strategy similar to the one used in the RUP system [McAllester, 1982].

2 The Theory of Epistemic Change

For the following discussion, we will assume a propositional language $\mathcal{L}$ containing propositions x, y, z and the standard sentential connectives ($\neg$, $\vee$, $\wedge$, $\rightarrow$, $\leftrightarrow$). Sets of propositions will be denoted by A, B, C. Furthermore, $\vdash$ shall denote classical propositional derivability[1] and *Cn* should be a function mapping sets of propositions to sets of propositions by applying $\vdash$, i.e.

$$Cn(A) \stackrel{\text{def}}{=} \{x \in \mathcal{L} |\; A \vdash x\} \qquad (1)$$

Formalizing Newell's notions of the knowledge level in this setting, sets of propositions A closed with respect to *Cn* (i.e. $A = Cn(A)$)—technically speaking *propositional theories*—can be identified with *knowledge-level knowledge bases* as argued in [Diettrich, 1986; Levesque and Brachman, 1986]. Arbitrary, finite set of propositions $B \subseteq \mathcal{L}$ can be identified with *symbol-level knowledge bases*.

In the theory of epistemic change [Gärdenfors, 1988], only knowledge-level knowledge bases are considered, which are called *belief sets*. *Epistemic change operations* on such belief sets are

Expansion: the monotonic addition of a belief with the requirement that the result is again a belief set (written $A + x$),

Contraction: the removal of a proposition from a belief set resulting in a new belief set (written $A \dot{-} x$),

Revision: incorporation of a new proposition into a belief set under the requirement that the result is a consistent belief set (written $A \dot{+} x$).

While expansion is a well-defined, unique operation, namely:

$$A + x \stackrel{\text{def}}{=} Cn(A \cup \{x\}) \qquad (2)$$

the other two operations are problematical. An immediate criterion for them is that a belief set shall be changed minimally by an epistemic change operation (but cf. [Winslett, 1986]). However considering contraction, given a belief set B and a proposition x, in general there is no *unique greatest* belief set $C \subseteq B$ such that $C \not\vdash x$.

The problem of finding intuitively plausible change operations is approached by formulating sets of *rationality postulates* any epistemic change operation should satisfy. A set of such postulates for contraction can be given as follows (A a belief set, x, y propositions):

($\dot{-}$1) *$A \dot{-} x$ is a belief set (closure);*

($\dot{-}$2) *$A \dot{-} x \subseteq A$ (inclusion);*

($\dot{-}$3) *If $x \notin A$ then $A \dot{-} x = A$ (vacuity);*

($\dot{-}$4) *If $\not\vdash x$, then $x \notin (A \dot{-} x)$ (success);*

($\dot{-}$5) *If $Cn(\{x\}) = Cn(\{y\})$ then $A \dot{-} x = A \dot{-} y$ (preservation);*

($\dot{-}$6) *$A \subseteq (A \dot{-} x) + x$ (recovery);*

($\dot{-}$7) *$(A \dot{-} x) \cap (A \dot{-} y) \subseteq A \dot{-} (x \wedge y)$;*

($\dot{-}$8) *If $x \notin A \dot{-} (x \wedge y)$, then $A \dot{-} (x \wedge y) \subseteq A \dot{-} x$.*

Most of these postulates are straightforward. The *closure* postulate ($\dot{-}$1) tells us that we always get a belief set when applying $\dot{-}$ to a belief set and a proposition. The *inclusion* postulate ($\dot{-}$2) assures that when a proposition is removed, nothing previously unknown can enter into the belief set, setting an upper bound for any possible contraction operation. Postulate ($\dot{-}$3) takes care of one of the limiting cases, namely, that the proposition to be removed is not part of the belief set, while the next postulate ($\dot{-}$4) describes the effect of

[1] The results presented in the following can be generalized to conservative extensions of propositional logics, provided they are compact and monotonic (cf. [Gärdenfors, 1988, p. 21–26]).

the other case. If the proposition to be removed is not a logically valid one, then the contraction operation will effectively remove it. The *preservation* postulate ($\dot{-}5$) assures that the syntactical form of the proposition to be removed will not effect the resulting belief set. Any two propositions which are logically equivalent shall lead to the same result. Finally, the *recovery* postulate ($\dot{-}6$) describes the lower bound of any contraction operation. The contracted belief set should contain enough information to recover all propositions deleted. Note that ($\dot{-}6$) together with ($\dot{-}1$)–($\dot{-}5$) entails the following conditional equation:

$$\textit{If } x \in A \textit{ then } A = (A \dot{-} x) + x \tag{3}$$

The two postulates ($\dot{-}7$) and ($\dot{-}8$) are less obvious and not as basic as the former ones—a reason for calling them "supplementary postulates." ($\dot{-}7$) states that retracting a conjunction should remove less information than retracting both conjuncts individually in parallel, with ($\dot{-}8$) its conditional converse. Although this does not sound like a strong restriction, not all conceivable contraction operations satisfy it. In order to shed some more light on these supplementary postulates, it might be worthwhile to present some principles derivable from ($\dot{-}7$) and ($\dot{-}8$).[2] First, there is the following "factoring" condition:

$$A \dot{-} (x \wedge y) = \begin{cases} (A \dot{-} x) \cap (A \dot{-} y) & \textit{or} \\ A \dot{-} x & \textit{or} \\ A \dot{-} y & \end{cases} \tag{4}$$

This condition is actually equivalent to ($\dot{-}7$) and ($\dot{-}8$) if a contraction operation already satisfies ($\dot{-}1$)–($\dot{-}6$).

Another interesting property derivable from the supplementary postulates is an identity criterion for contracted belief sets:

$$\textit{If } (x \rightarrow y) \in A \dot{-} y \textit{ and } (y \rightarrow x) \in A \dot{-} x \textit{ then } A \dot{-} y = A \dot{-} x \tag{5}$$

Turning now to revision, we note that there are two independent ways to characterize this operation. First, a set of rationality postulates for revision could be specified capturing the idea that a revised belief set should minimally differ from the original belief set, as done in [Alchourrón *et al.*, 1985]. Second, one could define the revision operation $A \dot{+} x$ by first contracting A with respect to $\neg x$ in order to avoid inconsistencies, and then expanding the result by x:

$$A \dot{+} x \stackrel{\text{def}}{=} (A \dot{-} \neg x) + x \tag{6}$$

[2]Proofs for the principles (4) and (5) can be found in [Gärdenfors, 1988].

This way of defining a revision operation was proposed by Levi [1977]. As it turns out, both ways of characterizing revision are equivalent, as shown in [Alchourrón *et al.*, 1985]. Any revision operation satisfying the rationality postulates for revision could be generated by (6) and a contraction operation satisfying the contraction postulates and *vice versa*. What should be noted at this point is that in the case when $\vdash \neg x$ (and only in this case), the revised belief set will be inconsistent—which cannot be avoided, however.

Parallel to defining revision by contraction, we could try it the other way around—defining contraction in terms of revision:[3]

$$A \dot{-} x \stackrel{\text{def}}{=} (A \dot{+} \neg x) \cap A \tag{7}$$

This means, revision and contraction are interdefinable and it suffices to analyze one of these operations. Whether contraction or revision is taken as the basic one is mostly a matter of taste and philosophy (cf. [Makinson, 1985; Dalal, 1988]).

3 Constructing Contraction Functions

Using the rationality postulates, a number of possible contraction functions (and the associated revision functions) are studied and evaluated in [Alchourrón *et al.*, 1985] and [Gärdenfors, 1988].

All of these operations are defined using the family of maximal subsets not implying a given proposition, denoted by $A{\downarrow}x$ (pronounced "A less x"):

$$A{\downarrow}x \stackrel{\text{def}}{=} \{B \subseteq A \mid B \not\vdash x \textit{ and } \textit{if } B \subset C \subseteq A \textit{ then } C \vdash x\} \tag{8}$$

Note that all elements of $A{\downarrow}x$ are again belief sets because of the maximality condition. Trying to construct contraction functions based on $A{\downarrow}x$, a first idea could be to take into account all possible outcomes of removing a proposition, and, since we do not have a measure of what is a better solution, to choose the intersection of the outcomes as the result of the contraction operation. If $A{\downarrow}x$ is empty—which can only happen if x is a logically valid proposition—A itself will be taken as the solution.

$$A \stackrel{\text{f}}{-} x \stackrel{\text{def}}{=} \begin{cases} \bigcap(A{\downarrow}x) & \textit{if } \not\vdash x \\ A & \textit{otherwise} \end{cases} \tag{9}$$

This operation satisfies obviously most of the rationality postulates.[4]

[3]Note that the intersection of two belief sets is again a belief set.

[4]In order to make the paper self-contained, I included proofs for all lemmas in this section in Appendix A.

Lemma 1 *Full meet contraction satisfies ($\dot{-}$1)–($\dot{-}$5).*

In order to see that full meet contraction also satisfies ($\dot{-}$6), the following lemma is helpful.

Lemma 2 *Let A be a belief set, and let x be a proposition such that $x \in A$ and $\not\vdash x$. Then*

$$A \stackrel{\mathrm{f}}{-} x \quad = \quad A \cap Cn(\{\neg x\}) \tag{10}$$

Applying this result to revision by using (6), it becomes obvious that full meet contraction is most probably not an operation one wants to use. Full meet contraction removes too much information.

Corollary 3 *For a revision operation defined by (6) and (9), for any x such that $\neg x \in A$ and $\not\vdash \neg x$ it holds that*

$$A \dot{+} x \quad = \quad Cn(\{x\}) \tag{11}$$

Nevertheless, full meet contraction satisfies all the rationality postulates for contraction.

Lemma 4 *Full meet contraction as defined by (9) satisfies ($\dot{-}$1)–($\dot{-}$8).*

Looking for a more reasonable contraction function, another way to contract a belief set could be to choose one of the elements in $A{\downarrow}x$—employing a *choice function* $\mathcal{C}$—instead of using the intersection over all elements:

$$A \stackrel{\mathrm{m}}{-} x \quad \stackrel{\mathrm{def}}{=} \quad \begin{cases} \mathcal{C}(A{\downarrow}x) & \textit{if } \not\vdash x \\ A & \textit{otherwise} \end{cases} \tag{12}$$

It is easy to see that this operation satisfies ($\dot{-}$1)–($\dot{-}$6), but the supplementary postulates are not satisfied unconditionally [Alchourrón and Makinson, 1982]. Ignoring this fact for the moment, let us try to characterize the result of such contraction operations. As it turns out, the contraction function defined by (12) generates belief sets which are far too large to be plausible.

Lemma 5 *Let A be a belief set with $x \in A$. Then for any proposition y:*

$$(x \vee y) \in A \stackrel{\mathrm{m}}{-} x \quad \textit{or} \quad (x \vee \neg y) \in A \stackrel{\mathrm{m}}{-} x \tag{13}$$

This property has a rather counter-intuitive consequence for revision. Applying again (6), we get the following result.

Corollary 6 *Let $\dot{+}$ be a revision operation defined by using (12) and (6). Then, for any proposition x and belief set A with $\neg x \in A$:*

$$y \in A \dot{+} x \quad \textit{or} \quad \neg y \in A \dot{+} x \tag{14}$$

This means that by applying maxichoice revision to an arbitrary belief set, we get all of the sudden a *complete* belief set, provided that $\neg x \in A$. However, starting with an arbitrary belief set in which there may be no belief in some proposition z or its negation $\neg z$ and ending up with a belief set in which for all propositions z, either z or $\neg z$ is believed, is clearly something not desirable.

Viewing full meet contraction and maxichoice contraction as two extreme points, it might be worthwhile to explore the "middle ground" between them. Instead of chosing one element from $A{\downarrow}x$ or the entire family of belief sets, a subfamily of $A{\downarrow}x$ is used to generate the contracted belief set. For this purpose, let us assume a *selection function* $\mathcal{S}$ which selects a subset of $A{\downarrow}x$:

$$A \stackrel{\mathrm{p}}{-} x \quad \stackrel{\mathrm{def}}{=} \quad \begin{cases} \bigcap \mathcal{S}(A{\downarrow}x) & \textit{if } \not\vdash x \\ A & \textit{otherwise} \end{cases} \tag{15}$$

This contraction function, called *partial meet contraction*, unconditionally satisfies the basic postulates, which can be easily verified.

Lemma 7 *Any partial meet contraction operation $\stackrel{\mathrm{p}}{-}$ satisfies ($\dot{-}$1)–($\dot{-}$6).*

What is more interesting is that the converse holds as well. Any operation satisfying ($\dot{-}$1)–($\dot{-}$6) is a partial meet contraction [Gärdenfors, 1988, Theorem 4.13]. In order to satisfy the supplementary postulates, some restrictions on $\mathcal{S}$ must be imposed. Let us assume a "preference relation" $\sqsubseteq$ over all subsets of a belief set A independent of x such that for all $B, C \subseteq A$:

$$\textit{if } B \in \mathcal{S}(A{\downarrow}x) \textit{ and } C \in A{\downarrow}x \textit{ then } C \sqsubseteq B \tag{16}$$

in which case the contraction function is called *relational contraction*.

Lemma 8 *Any relational contraction function satisfies ($\dot{-}$1)–($\dot{-}$7)*

If the relation $\sqsubseteq$ is a transitive relation, the corresponding contraction function—called *transitively relational contraction*—satisfies all postulates.

Lemma 9 *Any transitively relational contraction function satisfies ($\dot{-}$1)–($\dot{-}$8).*

Furthermore, it is possible to show that any contraction function satisfying ($\dot{-}$1)–($\dot{-}$8) is a transitively relational contraction [Gärdenfors, 1988, Theorem 4.16].

4 Problems with the Approach

Although the results presented in the previous sections sound interesting and provide some insights into the problem of belief revision, it seems arguable whether the approach could be used in a computational context, as in AI or in the database field. The theory of epistemic change seems to capture only the *idealized* process of belief revision—ignoring some important problems of belief revision appearing in the "real world."

First of all, closed theories cannot be dealt with directly in a computational context because they are too large. At least, if we deal with them, we would like to have a finite representation (i.e. a finite axiomatization), and there seems to be no obvious way to derive a finite representation from a revised or contracted belief set in the general case.

Second, it seems to be preferable for pragmatic reasons to modify *belief bases*, i.e. finite sets of propositions, instead of belief sets, i.e. deductively closed theories. Propositions in belief bases usually represent something like facts, observations, rules, laws, etc., and when we are forced to change the belief set we would like to stay as close as possible to the original formulation of the finite base. In particular, when it becomes necessary to give up a proposition in the belief base, we would like to throw away the consequences of the retracted proposition if they are not supported otherwise, i.e. to perform *reason maintenance*. More formally, given a belief base B and propositions $x, y \in Cn(B)$ and assuming that a proposition y is removed from B to accomplish a contraction $Cn(B) \dot{-} x$, then we expect that $z \notin (Cn(B) \dot{-} x)$ if y was responsible for z, i.e. $z \in Cn(B)$ but $z \notin Cn(B \setminus \{y\})$.

For instance, let a be the proposition "the device is ok," let b be the proposition "the output voltage is 5V" and let us assume we have the base $B = \{a, (a \rightarrow b)\}$, i.e. "the device is ok" and "if the device is ok then the output voltage is 5V." That means that from B we can infer that "the output voltage is 5V." Now, when we learn that the device is defect, then together with a we would like to get rid of b because the reason for the belief that the output voltage is 5V has vanished. This, however, cannot be easily accomplished by the approach described above. On the contrary, since the theory of the epistemic change formalizes the idea of keeping as much of the old propositions (in the belief set) as possible, it seems likely that b will be among the propositions in the contracted belief set since it does not contradict $\neg a$. Gärdenfors puts it in the following way [Gärdenfors, 1988, p. 67]:

> However, belief sets cannot be used to express that some beliefs may be *reasons* for other beliefs. (This deficiency was one of the motivations behind Doyle's TMS ...). And intuitively, when we compare degrees of similarity between different epistemic states, we want the structure of reasons or justifications to count as well.

Actually, viewing this property from a cognitive angle, it could be defended by the argument that the theory of epistemic change models what is called the *coherence theory* of belief revision [Gärdenfors, 1989]. This means that in the course of revising beliefs the main emphasis is to arrive at a new *coherent* set of beliefs, which may be interpreted as a *logically consistent* set of beliefs. Identifying and discarding *derived* beliefs when their justifications are undermined, on the other hand, is not viewed as essential in the *coherence theory*. It is argued that it is intellectually much too expensive to keep track of all justifications—a fact supported by empirical evidence.

Nevertheless, although this theory may be right in the general case, in a problem-solving context, as modeled in typical AI applications, we usually want *reason maintenance*—as e.g. in the toy example given above.

Summarizing, we see that *belief revision* and *reason maintenance* are not genuinely connected with each other, as it sometimes seems to be perceived in AI (cf. [Martins and Shapiro, 1988]). However, as will be shown, it is not necessary to add *reason maintenance* as a primitive notion to a theory of belief revision. Reason maintenance will result as a side-effect when we choose the "right" contraction operation.

5 Contracting Finite Bases

As spelled out in the previous section, there are good reasons to perform belief revision on belief bases—considering the propositions in the base as the basic beliefs. As a matter of fact, such operations were adopted in an analysis of update semantics for logical databases [Fagin *et al.*, 1983; Fagin *et al.*, 1986] and in modelling counterfactual reasoning [Ginsberg, 1986].

Basically, revision ($\tilde{+}$) and contraction ($\sim$) on a belief base B is defined in the following way:

$$B \sim x \quad \stackrel{\text{def}}{=} \quad \begin{cases} \bigvee_{C \in (B \downarrow x)} C & \text{If } \not\vdash x \\ B & \text{otherwise} \end{cases} \tag{17}$$

$$B \;\tilde{+}\; x \quad \stackrel{\text{def}}{=} \quad (B \sim \neg x) \wedge x \tag{18}$$

with $B \downarrow x$ being the same operation as defined by equation (8) without requiring that B is deductively closed. The guiding idea behind (17) and (18) is that we want to retain as many old propositions as possible, and if there are ties, we take the disjunction. Moreover, (18) is logically very similar to (6). Obviously, such change operations realize some form of reason maintenance, as one can see in (19).

$$\{a, a \to b\} \sim a \quad = \quad \{a \to b\} \not\vdash b \tag{19}$$

Based on (17) and (18), both Ginsberg [1986] and Fagin *et al.* [1983] consider more elaborated versions of contraction and revision, which distinguish between different kinds propositions in the belief base. For instance, Fagin *et al.* distinguish between *facts* and *integrity rules* in the belief base, and Ginsberg proposes to protect some propositions against retraction. We will ignore this issue here. However, one should note that such a construction is not qualitatively different from $\sim$ and $\tilde{+}$ [Nebel, 1989]. In particular, the results in this and the next section are valid for such operations.

In trying to relate $\sim$ to $\dot{-}$, we see that the rationality postulates presented in Section 2 cannot be applied immediately to $\sim$ since it does not operate on belief sets. However, it seems possible to adapt the postulates to belief bases by setting

$$A \stackrel{\text{def}}{=} Cn(B) \tag{20}$$

$$A \dot{-} x \stackrel{\text{def}}{=} Cn(B \sim x) \tag{21}$$

Thus, in a sense, we view $\sim$ as an implementation of $\dot{-}$.

Lemma 10 *Under the assumption of (20) and (21), $\sim$ satisfies ($\dot{-}$1)–($\dot{-}$5).*

Proof: ($\dot{-}$1) holds trivially because of (21). If $\not\vdash x$, ($\dot{-}$2) is satisfied because

$$Cn(\bigvee_{C \in (B \downarrow x)} C) \quad = \quad \bigcap_{C \in (B \downarrow x)} Cn(C) \tag{22}$$

and for all $C : Cn(C) \subseteq Cn(B)$. If $\vdash x$, ($\dot{-}$2) is satisfied as well since the belief base is not changed. ($\dot{-}$3) is satisfied because if $B \not\vdash x$, then $B \downarrow x = \{B\}$. ($\dot{-}$4) holds because for all $C \in (B \downarrow x)$ we have $C \not\vdash x$ and, hence, $\bigvee C \not\vdash x$. Finally, ($\dot{-}$5) holds since for the determination of $B \downarrow x$ the syntactic form of x does not matter. ■

Unfortunately, however, the recovery postulate is not satisfied. For instance, we have

$$Cn((\{a, a \to b\} \sim b) \cup \{b\}) \not\supseteq Cn(\{a, a \to b\}) \tag{23}$$

Trying to find the reason for this unsatisfying behavior, one notes that even the weakest possible contraction function on belief sets—full meet contraction—generates belief sets such that (Lemma 2)

$$A \dot{-} x = A \cap Cn(\{\neg x\})$$

which is sufficient for restoring the original belief set as we have seen in Lemma 4.

Adding a finite conjunct, logically equivalent to $A \cap Cn(\{\neg x\})$, to the outcome of $\sim$ leads to a new contraction function which has the desired property:

$$B \mathrel{\dot{\sim}} x \stackrel{\text{def}}{=} \begin{cases} (\bigvee_{C \in (B \downarrow x)} C) \wedge (B \vee \neg x) & \textit{if } \not\vdash x \\ B & \textit{otherwise} \end{cases} \tag{24}$$

Lemma 11 *Under the assumption of (20) and (21), $\dot{\sim}$ satisfies ($\dot{-}$1)–($\dot{-}$6).*

Proof: The satisfaction of ($\dot{-}$1), ($\dot{-}$3), and ($\dot{-}$5) can be shown with the arguments used in the proof of Lemma 10. That ($\dot{-}$2) holds becomes obvious by observing that $Cn(B \vee \neg x)$ and $Cn(\bigvee C)$ are both subsets of $Cn(B)$, and, hence, the set of consequences of their unions can only be a subset of $Cn(B)$. ($\dot{-}$4) holds because for the added conjunct we have $(B \vee \neg x) \not\vdash x$ and no $C \in B \downarrow x$ implies x. Finally, ($\dot{-}$6) holds since $Cn(B \mathrel{\dot{\sim}} x)$ contains $Cn(B) \cap Cn(\{\neg x\})$, which is sufficient for recovery as shown in the proof of Lemma 4. ■

Actually, if revision is the only operation of interest, it does not make a difference whether we employ $\dot{\sim}$ or $\sim$. The revision operation $\tilde{+}$ gives identical results (*wrt Cn*) regardless of whether $\dot{\sim}$ or $\sim$ is used.

Theorem 12 *The operations $\dot{\sim}$ and $\sim$ are revision-equivalent wrt $\tilde{+}$ as defined by (18), i.e.*

$$Cn((B \mathrel{\dot{\sim}} \neg x) \wedge x) \quad = \quad Cn((B \sim \neg x) \wedge x) \tag{25}$$

Proof: In the limiting cases when $\vdash \neg x$ or $\neg x \notin Cn(B)$, $\sim$ and $\dot{\sim}$ give the same results trivially. For the principal case, $\not\vdash \neg x$ and $\neg x \in Cn(B)$, we have:

$$\begin{aligned} Cn((B \mathrel{\dot{\sim}} \neg x) \wedge x) &= \\ &= Cn((\bigvee_{C \in (B \downarrow \neg x)} C) \wedge (B \vee x) \wedge x) \\ &= Cn((\bigvee_{C \in (B \downarrow \neg x)} C) \wedge x) \\ &= Cn(B \sim \neg x) + x \end{aligned}$$

■

This result might raise the question of the value of the recovery postulate—a problem discussed in [Makinson, 1987]. Despite the fact that it is necessary to establish some of the theoretical results cited in Section 3, it has some practical value, I believe. In a setting where revision and contraction are equally important, as in the case of database updates (cf. [Fagin *et al.*, 1986]), we had better use $\stackrel{\sim}{\sim}$ instead of $\sim$. Otherwise, more information is lost than intended. In particular, we might be unable to *undo* a contraction operation.

Before we now go on trying to verify that the supplementary postulates are satisfied by $\stackrel{\sim}{\sim}$, we will try to establish a connection between belief-base and belief-set contraction. In [Ginsberg, 1986], as well as in [Fagin *et al.*, 1983], some thoughts are devoted to the issue of modifying closed theories, i.e. belief sets. However, by considering (22) and permitting infinite disjunctions, it becomes quickly obvious that in the limiting case when the belief base is a belief set, $\sim$ is identical to full meet contraction. This kind of contraction is rather useless, however, as demonstrated by Lemma 2 and Corollary 3. Thus, Fagin *et al.* [1983] and Ginsberg [1986] conclude that belief revision is a phenomenon which can only be modeled if the *syntactic representation* of a belief set (its belief base) is taken into account. A knowledge-level analysis of belief revision seems to be impossible.

Addressing this problem, Ginsberg [1986, Section 8.1] proposes to incorporate *reason-maintenance* information into the logic by using derivations as truth-values. This proposal leads to the desired results, i.e. changes of belief sets (in the reason-maintenance style logic) are identical to changes of a finite belief bases. However, this approach does not shed too much light onto the relation between modifications of belief sets and modifications of finite belief bases.

6 Base Contraction Viewed as Partial Meet Contraction

Reconsidering the arguments from above, we note that we are not interested in the particular syntactical form of a belief base, but we regard the propositions in the belief base as somehow more important than any derived propositions. In particular, given two belief bases B and B' such that for all $x \in B$ there exists $x' \in B'$ with $Cn(\{x\}) = Cn(\{x'\})$ and *vice versa*, it is immediate that $Cn(B \stackrel{\sim}{\sim} y) = Cn(B' \stackrel{\sim}{\sim} y)$. This means that it is not the syntactical form of a belief base we are interested in, but the "logical force" of the propositions in the base. Using this idea, it is tempting to define a selection function which selects elements from $(A \downarrow x)$ which contain maximal subsets of B, a set of propositions considered as "interesting":

$$\mathcal{S}_B(A \downarrow x) \stackrel{\text{def}}{=} \{C \in (A \downarrow x) | \forall C' \in A \downarrow x : C' \cap B \not\supset C \cap B\}$$

Based on this selection function, we can define a partial meet contraction on belief sets which has the property of being identical to $\stackrel{\sim}{\sim}$ (*wrt* Cn). In order to show this, the next lemma proves to be helpful.

Lemma 13 *Let A be a belief set and S be any subset of A such that $S \not\vdash x$. Then*

$$\bigcap \{C \in (A \downarrow x) |\, S \subseteq C\} = Cn(S \cup (A \cap Cn(\{\neg x\})))$$

Proof: "$\supseteq$": First, by Lemma 2 we know that any intersection over a subfamily of $A \downarrow x$ must contain $(A \cap Cn(\{\neg x\}))$. Second, since all C in the subfamily chosen contain S, the LHS contains S. Finally, because the intersection over any subfamily of $A \downarrow x$ is a belief set, the LHS is a belief set containing S and $(A \cap Cn(\{\neg x\}))$ and, hence, the right hand side.

"$\subseteq$": Assume the contrary, i.e. there is a y such that $y \in$ LHS and $y \notin$ RHS. Using set theory and the properties of Cn, we can transform the RHS in the following way:

$$\begin{aligned} Cn(S \cup (A \cap Cn(\{\neg x\}))) = \\ &= Cn((S \cup A) \cap (S \cup Cn(\{\neg x\}))) \\ &= Cn(A \cap Cn(S \cup Cn(\{\neg x\}))) \\ &= A \cap Cn(S \cup \{\neg x\}) \end{aligned}$$

Since $y \in A$ because of our assumption, we must have $y \notin Cn(S \cup \{\neg x\})$ and, in particular, $\neg x \not\vdash y$. Using this, we can derive $(x \vee \neg y) \not\vdash x$, following the same line of arguments as in the proof of Lemma 2. By that and the observation that $y \notin Cn(S)$, we can conclude that $x \notin Cn(S \cup \{(x \vee \neg y)\})$. Since adding y to this set would lead to the derivation of x, there must be a set $E \supseteq S \cup \{(x \vee \neg y)\}$ with $y \notin E$ and $E \in (A \downarrow x)$, which means that y cannot be a member of all sets in $A \downarrow x$ which contain S, and we have a contradiction. ■

Based on this lemma, we can establish the connection between contractions on belief bases and contractions on belief sets.

Theorem 14 *Contraction of finite premise sets B using $\stackrel{\sim}{\sim}$ is identical (wrt $\vdash$) to a partial meet contraction $\stackrel{\text{p}}{-}$ defined by the selection function $\mathcal{S}_B$, i.e.*

$$Cn(B \stackrel{\sim}{\sim} x) \quad = \quad Cn(B) \stackrel{\text{p}}{-} x \tag{26}$$

Proof: In the limiting cases when $\vdash x$ or $x \notin Cn(B)$ the result is immediate. For the principal case, we note that

$$\begin{aligned} &\mathcal{S}_B(Cn(B)\downarrow x) = \\ &= \bigcup_{D \in (B \downarrow x)} \{C \in (Cn(B)\downarrow x) | D \subseteq C\} \end{aligned} \quad (27)$$

Applying Lemma 13, it follows that

$$\begin{aligned} Cn(B) \stackrel{p}{-} x &= \\ &= \bigcap \mathcal{S}_B(Cn(B)\downarrow x) \\ &= \bigcap_{D \in (B\downarrow x)} Cn((D) \cup (Cn(B) \cap Cn(\{\neg x\}))) \\ &= Cn((\bigcap_{D \in (B\downarrow x)} Cn(D)) \cup (Cn(B) \cap Cn(\{\neg x\}))) \\ &= Cn((\bigvee_{D \in (B\downarrow x)} Cn(D)) \wedge (B \vee \neg x)) \\ &= Cn(B \dot{\sim} x) \end{aligned}$$

■

Thus, contrary to the assumptions spelled out in [Diettrich, 1986; Fagin *et al.*, 1983; Fagin *et al.*, 1986; Ginsberg, 1986], revision and contraction on finite belief bases are not qualitatively different from epistemic change operations on deductively closed belief sets. The finite case can be modeled without any problem by a particular selection function. Viewed from a knowledge-level perspective, the only additional information needed for belief revision is a preference relation on sets of propositions. It should be noted, however, that the construction did not lead to an epistemic change function which satisfies all rationality postulates.

Theorem 15 *Any partial meet contraction using* $\mathcal{S}_B$ *satisfies the rationality postulates* ($\dot{-}$ 1)–($\dot{-}$ 7).

Proof: Because of Lemma 7 and the fact that $\stackrel{p}{-}$ constructed by $\mathcal{S}_B$ is a partial meet contraction, ($\dot{-}$1)–($\dot{-}$6) are satisfied. That ($\dot{-}$7) is satisfied follows from Lemma 8 and the fact that $\mathcal{S}_B$ satisfies (16), if $\sqsubseteq$ is defined as:

$$X \sqsubseteq Y \quad \textit{iff} \quad X \cap B \not\supset Y \cap B \quad (28)$$

■

Since $\mathcal{S}_B$ is not a transitively relational selection function, ($\dot{-}$8) is not satisfied in general. In order to give an example where ($\dot{-}$8) is violated, let us assume that $\mathcal{S}_B$ is used to define the partial meet contraction $\stackrel{p}{-}$ and that

$$B \quad = \quad \{a, b \wedge c, a \wedge b \wedge d, a \wedge d\}$$

Setting, $x = (a \wedge c)$ and $y = (b \wedge d)$, we see that

$$x \notin Cn(B) \stackrel{p}{-} (x \wedge y)$$

but

$$Cn(B) \stackrel{p}{-} (x \wedge y) \not\subseteq Cn(B) \stackrel{p}{-} x$$

because we have the following relationships:

$$\begin{aligned} (a \wedge c) \quad &\notin \quad Cn(B) \stackrel{p}{-} ((a \wedge c) \wedge (b \wedge d)) \\ a \quad &\in \quad Cn(B) \stackrel{p}{-} ((a \wedge c) \wedge (b \wedge d)) \\ a \quad &\notin \quad Cn(B) \stackrel{p}{-} (a \wedge c) \end{aligned}$$

As can be verified, the factoring condition (4) is violated, as well. So what? Does this result imply that $\stackrel{p}{-}$ defined by using $\mathcal{S}_B$ and, hence, $\dot{\sim}$ is not a reasonable contraction operation? Actually, it seems to make a lot more sense than $\stackrel{f}{-}$.

The disadvantage of not having ($\dot{-}$8) is actually very subtle. One consequence is that revision operations based on $\dot{\sim}$ violate the respective postulate for revision, which is needed to derive an identity criterion for revised belief sets similar to (5):

$$\begin{aligned} &\textit{If } x \in A \dot{+} y \textit{ and } y \in A \dot{+} x \\ &\textit{then } A \dot{+} y = A \dot{+} x \end{aligned} \quad (29)$$

This criterion in turn is similar to a principle Stalnaker [1968] postulated for the interpretation of counterfactual conditionals on "neighboring possible worlds." However, it seems to be difficult to come up with an example demonstrating that the violation of this principle leads to obviously counter-intuitive results.

7 Maxichoice Contraction on Belief Bases

As we noted above, we could achieve satisfaction of ($\dot{-}$8) if the ordering $\sqsubseteq$ is transitive. Embedding the partial order defined by set-inclusion in a total ordering would certainly help. This can be achieved by starting from an arbitrary total ordering on the propositions in a belief base, for instance.[5] Thus, let us assume such an ordering $\leq$ on the propositions of a belief base B. Furthermore, define $\sqsubseteq'$ on 2^B:

$$X \sqsubseteq' Y \;\; \textit{iff} \; \forall x \in (X \setminus Y)\, \exists y \in (Y \setminus X) : \; x \leq y$$

[5]Note that such a construction is fundamentally different from the notion of epistemic entrenchment as introduced in [Gärdenfors, 1988; Gärdenfors and Makinson, 1988].

This means, in case when X and Y are incomparable by set-inclusion, we assume Y to be larger than X by $\sqsubseteq'$ iff there is an element in Y which is larger by $\leq$ than any element in X which is not in Y. Based on that, let us define a selection function as follows:

$$\mathcal{S}_{B,\leq}(A{\downarrow}x) \stackrel{\text{def}}{=} \{C \in (A{\downarrow}x) \mid \forall C' \in (A{\downarrow}x) : C' \cap B \sqsubseteq' C \cap B\}$$

Theorem 16 *A partial meet contraction defined by using $\mathcal{S}_{B,\leq}$ satisfies all rationality postulates.*

Proof: Obviously, $\sqsubseteq'$ is a transitive relation, and, thus, the partial meet contraction defined by $\mathcal{S}_{B,\leq}$ is a transitively relational contraction. ■

Another interesting point about $\mathcal{S}_{B,\leq}$ is that it is similar to a maxichoice contraction on belief bases.

Lemma 17 *For partial meet contractions defined by using $\mathcal{S}_{B,\leq}$, for all x with $\not\vdash x$, there is an $E \in (Cn(B){\downarrow}x)$ such that*

$$Cn(B) \stackrel{\text{p}}{-} x = Cn(E \cup (Cn(B) \cap Cn(\{\neg x\}))) \quad (30)$$

Proof: The main point is that there is exactly one element in $E \in (B{\downarrow}x)$ such that for all $C \in (B{\downarrow}x)$ we have $C \sqsubseteq' E$. Applying Lemma 13, we get the desired result. ■

From this we can conclude two things. First, maxichoice contraction on belief bases does not have the undesirable result as the same operation applied to belief sets. Second, considering the rationality postulates, it is more "rational" to chose one alternative of $(B{\downarrow}x)$ than to take the disjunction of all alternatives. Furthermore, from a practical point of view, this strategy has the advantage that the belief base is not cluttered with disjunctions, but it simply shrinks.

Summarizing, although it might seem arbitrary to choose one maximal consistent set during a contraction, it is despite its practical value justified because it is at least as "rational" as using the disjunctions. Thus, for example, the *truth-maintenance system* RUP [McAllester, 1982], which implements this strategy for belief revision, could be characterized as a "fully rational belief revision system" (modulo its inferential incompleteness).

8 Conclusion and Outlook

We have shown that belief revision, as exercised in many applications in AI, is not an activity which can only be analyzed on the symbol level. Employing the theory of epistemic change we demonstrated how to reconstruct symbol-level belief revision in the theory of epistemic change—resulting in a knowledge-level analysis of some aspects of belief revision. In particular, we have shown that *reason maintenance* is a symbol-level notion, which although not present in the theory of epistemic change, appears as a side-effect. Furthermore, analyzing the approaches, we noted that chosing one maximal consistent subset of a belief base seems to be more rational than taking disjunctions of all maximal consistent sets—considering the rationality postulates.

However, a number of issues remain unresolved—of course. For instance, iterated contractions were ignored because they present serious problems. One has to provide update operations for the preference relation among propositions. Furthermore, implausible contraction operations, such as $\{a \wedge b\} \sim a = \emptyset$, were ignored. Choosing the "right" form of the premises seems to be one of the central tasks before any kind of belief revision can be applied.

Acknowledgments

I am grateful to Alex Borgida, Mukesh Dalal, and Karl Schlechta who provided me with a number of hints and ideas concerning the issues discussed in this paper. I would also like to express my thanks to Peter Gärdenfors, who sent me copies of his papers, made available the manuscript of his book, and provided me with some helpful comments an earlier version of this paper.

A Proofs of Lemmas of Section 3

Lemma 1 *Full meet contraction satsfies ($\dot{-}$1)–($\dot{-}$5).*

Proof: ($\dot{-}$1) is satisfied because the intersection of belief sets is a belief set. ($\dot{-}$2) is satisfied because for all $E \in (A{\downarrow}x)$, $E \subseteq A$. ($\dot{-}$3) holds because $A{\downarrow}x = \{A\}$, if $A \not\vdash x$. ($\dot{-}$4) holds since for all $E \in (A{\downarrow}x)$, $E \not\vdash x$, if $\not\vdash x$. ($\dot{-}$5) is satisfied since the syntactical form of x in $A{\downarrow}x$ is irrelevant. ■

Lemma 2 *Let A be a belief set, and let x be a proposition such that $x \in A$ and $\not\vdash x$, then*

$$A \stackrel{\text{f}}{-} x \quad = \quad A \cap Cn(\{\neg x\})$$

Proof: First, consider the case when $y \in A$ and $\neg x \vdash y$. Now assume that $y \notin A \stackrel{\text{f}}{-} x$. That means that there is a set $E \in A{\downarrow}x$ such that $y \notin E$. Because of the maximality condition on all such sets, we know that $x \in Cn(E \cup \{y\}))$. Using contraposition on our premise $\neg x \vdash y$, we get $\neg y \vdash x$ and hence $x \in Cn(E \cup \{\neg y\})$. Together with the previous result, we have $x \in Cn(E \cup \{(y \vee \neg y)\}) = Cn(E)$ and a contradiction.

For the case $y \in A$ and $\neg x \not\vdash y$, we know by contraposition that $\neg y \not\vdash x$ and hence $x \vee \neg y \not\vdash x$. Because of the maximality of the sets in $A \downarrow x$, there are two sets E', E'' with $y \in E'$ and $(x \vee \neg y) \in E''$, but there can be no set which includes both because $x \in Cn(\{y, (x \vee \neg y)\})$, and thus $y \notin \bigcap(A{\downarrow}x)$. ■

Corollary 3 *For a revision operation defined by (6) and (9), for any x such that $\neg x \in A$ and $\not\vdash \neg x$ it holds that*

$$A \dot{+} x \quad = \quad Cn(\{x\})$$

Proof: Using (6) and (10) leads to:

$$\begin{aligned} A \dot{+} x &= Cn((A \cap Cn(\{x\})) \cup \{x\}) \\ &= Cn((A \cup \{x\}) \cap (Cn(\{x\}) \cup \{x\})) \\ &= Cn(Cn(A \cup \{x\}) \cap Cn(\{x\})) \end{aligned}$$

Since we assumed $A \vdash \neg x$, we know $Cn(A \cup \{x\}) = \mathcal{L}$, and, thus, the result is $Cn(\{x\})$. ■

Lemma 4 *Full meet contraction as defined by (9) satisfies ($\dot{-}$1)–($\dot{-}$8).*

Proof Sketch: Since ($\dot{-}$1)–($\dot{-}$5) are obvious, we will focus on ($\dot{-}$6)–($\dot{-}$8). These are satisfied trivially for the two limiting cases $\vdash x$ and $x \notin A$. That ($\dot{-}$6) holds in the principal case, $x \in A$ and $\not\vdash x$, becomes obvious when substituting the right hand side of equation (10) for $A \dot{-} x$ in ($\dot{-}$6), which leads to:

$$A \quad \subseteq \quad Cn((A \cap Cn(\{\neg x\})) \cup \{x\}) \tag{31}$$

Now, because for any $y \in A$ we know that $(y \vee \neg x) \in Cn(\{\neg x\})$ and that this together with x implies y, the right hand side of (31) is clearly a superset of the left hand side. Furthermore, using Lemma 2, it can be easily derived that ($\dot{-}$7) (it can even be strengthened to equality) and ($\dot{-}$8) hold as well. ■

Lemma 5 *Let A be a belief set with $x \in A$. Then for any proposition y:*

$$(x \vee y) \in A \overset{\text{m}}{-} x \quad \textit{or} \quad (x \vee \neg y) \in A \overset{\text{m}}{-} x$$

Proof: In the limiting cases when $\vdash x$, the lemma holds trivially. In the other case, we know that $y \not\vdash x$ or $\neg y \not\vdash x$. Thus, because of the maximality of the elements of $A{\downarrow}x$, either y or $\neg y$ is in $A \overset{\text{m}}{-} x$. Since for any z: $y \vdash (y \vee z)$, the lemma holds. ■

Corollary 6 *Let $\dot{+}$ be a revision operation defined by using (12) and (6). Then, for any proposition x and belief set A with*

$$y \in A \dot{+} x \quad \textit{or} \quad \neg y \in A \dot{+} x$$

Proof: Expanding $A \overset{\text{m}}{-} \neg x$ by x leads by Lemma 5 for all propositions y to

$$\begin{aligned} ((\neg x \vee y) \wedge x) &\in (A \overset{\text{m}}{-} \neg x) + x \quad \textit{or} \\ ((\neg x \vee \neg y) \wedge x) &\in (A \overset{\text{m}}{-} \neg x) + x \end{aligned}$$

from which the desired result is immediate. ■

Lemma 7 *Any partial meet contraction operation $\overset{\text{p}}{-}$ satisfies ($\dot{-}$1)–($\dot{-}$6).*

Proof Sketch: ($\dot{-}$1)–($\dot{-}$5) can be easily verified. ($\dot{-}$6) holds because $A \dot{-} x$ always contains $A \cap Cn(\{\neg x\})$, which is sufficient for recovery. ■

Lemma 8 *Any relational contraction function satisfies ($\dot{-}$1)–($\dot{-}$7)*

Proof: ($\dot{-}$1)–($\dot{-}$6) follow from Lemma 7. If $\vdash x$, $\vdash y$, $x \notin A$, or $y \notin A$, the proof is immediate. For the other cases, we have to show that

$$\bigcap \mathcal{S}(A{\downarrow}x) \cap \bigcap \mathcal{S}(A{\downarrow}y) \quad \subseteq \quad \bigcap \mathcal{S}(A{\downarrow}(x \wedge y))$$

This could be done by showing that

$$(\mathcal{S}(A{\downarrow}x) \cup \mathcal{S}(A{\downarrow}y)) \quad \supseteq \quad \mathcal{S}(A{\downarrow}(x \wedge y))$$

Assume the contrary, i.e. there is an $E \in$ RHS, but $E \notin$ LHS. Since we cannot have $x, y \in E$, assume *wlg* $x \notin E$. Now it could be the case that $E \notin$ LHS because $E \notin (A \downarrow x)$. However, since $x \notin E$, there must be $D \in (A \downarrow x)$ with $E \subset D$. Furthermore, because $(x \wedge y) \notin D$, there is an $F \in (A{\downarrow}(x \wedge y))$ with $E \subset D \subseteq F$, which is a contradiction. This means $E \in (A{\downarrow}x)$. Accepting this, we must have $D \in (A{\downarrow}x)$ with $D \not\sqsubseteq E$. Because D is maximal, either $y \in D$ or $x \in Cn(D \cup \{y\})$. Thus, $D \in (A{\downarrow}(x \wedge y))$, and because of relationality, the premise $E \in$ RHS is contradicted. ■

Lemma 9 *Any transitively relational contraction function satisfies ($\dot{-}$1)–($\dot{-}$8).*

Proof: ($\dot{-}$1)–($\dot{-}$7) follows from Lemma 8. Similar to the proof above, for the principal case, we will show that when $x \notin \bigcap \mathcal{S}(A{\downarrow}(x \wedge y))$, then

$$\mathcal{S}(A{\downarrow}(x \wedge y)) \quad \supseteq \quad \mathcal{S}(A{\downarrow}x)$$

Assume the contrary, i.e. $E \in$ RHS but $E \notin$ LHS. Because E is maximal, either $y \in E$ or $x \in Cn(E \cup \{y\})$, and hence $E \in (A \downarrow (x \wedge y))$. Since $E \notin$ LHS, there is $D \in (A \downarrow (x \wedge y))$ with $D \not\sqsubseteq E$. Because of our premise $x \notin \bigcap \mathcal{S}(A \downarrow (x \wedge y))$, there must be at least one $F \in \mathcal{S}(A \downarrow (x \wedge y))$ with $x \notin F$, which is also an element of $(A{\downarrow}x)$. Thus $D \sqsubseteq F$ and $F \sqsubseteq E$, and because of transitivity $D \sqsubseteq E$, which is a contradiction. ■

References

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

[Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *The Journal of Symbolic Logic* **50**(2): 510–530, June 1985.

[Dalal, 1988] Mukesh Dalal. Investigations Into a Theory of Knowledge Base Revision: Preliminary Report. *Proc. AAAI-88*, , pp. 475–479. Saint Paul, Minn., 1988.

[de Kleer, 1986] Johan de Kleer. An Assumption-based TMS. *Artificial Intelligence* **28**(2): 127–162, March 1986.

[Diettrich, 1986] Thomas G. Diettrich. Learning at the Knowledge Level. *Machine Learning* **1**: 287–316, 1986.

[Doyle, 1979] Jon Doyle. A Truth Maintenance System. *Artificial Intelligence* **12**(3): 231–272, 1979.

[Fagin *et al.*, 1983] Ronald Fagin, Jeffrey D. Ullman, and Moshe Y. Vardi. On the Semantics of Updates in Databases. *Proc. Second ACM SIGACT-SIGMOD Symposium on Principles of Database Systems*, pp. 352–365. Atlanta, 1983.

[Fagin *et al.*, 1986] Ronald Fagin, Gabriel M. Kuper, Jeffrey D. Ullman, and Moshe Y. Vardi. Updating Logical Databases. *Advances in Computing Research* **3**: 1–18, 1986.

[Gärdenfors, 1988] Peter Gärdenfors. *Knowledge in Flux—Modeling the Dynamics of Epistemic States*. Cambridge, Mass.: MIT Press, 1988.

[Gärdenfors, 1989] Peter Gärdenfors. The Dynamics of Belief Systems: Foundations vs. Coherence Theories. Lund, Sweden: Department of Philosophy, Lund University, in preparation, January 1989.

[Gärdenfors and Makinson, 1988] Peter Gärdenfors and David Makinson. Revision of Knowledge Systems Using Epistemic Entrenchment. *Proceedings of the 2nd Workshop on Theoretical Aspects of Reasoning about Knowledge*, Los Altos, Cal.: Morgan Kaufmann, 1988.

[Ginsberg, 1986] Matthew L. Ginsberg. Counterfactuals. *Artificial Intelligence* **30**(1): 35–79, 1986.

[Levesque and Brachman, 1986] Hector J. Levesque and Ronald J. Brachman. Knowledge Level Interfaces to Information Systems. In M. L. Brodie and J. Mylopoulos (eds.), *On Knowledge Base Mangement Systems*, pp. 13–34. Berlin: Springer, 1986.

[Levi, 1977] Isaac Levi. Subjunctives, dispositions and chances. Synthese **34**: 423–455, 1977.

[Lewis, 1973] David K. Lewis. *Counterfactuals*. Cambridge, Mass.: Harvard University Press, 1973.

[Makinson, 1985] David Makinson. How to give it up: A Survey of Some Formal Aspects of Theory Change. *Synthese* **62**: 347–363, 1985.

[Makinson, 1987] David Makinson. On the Status of the Postulate of Recovery in the Logic of Theory Change. *Journal of Philosophical Logic* **16**: 383–394, 1987.

[Martins and Shapiro, 1988] João P. Martins and Stuart C. Shapiro. A Model for Belief Revision. *Artificial Intelligence* **35**(1): 25–79, May 1988.

[McAllester, 1982] David A. McAllester. Reasoning Utility Package User's Manual. AI Memo 667, Cambridge, Mass.: AI Laboratory, MIT, 1982.

[Nebel, 1989] Bernhard Nebel. *Reasoning and Revision in Hybrid Representation Systems*. Ph.D. thesis, in preparation, 1989.

[Newell, 1981] Allen Newell. The Knowledge Level. The Presidential Address, AAAI-80, Stanford, Cal. *The AI Magazine* **2**(2): 1–20, Summer 1981.

[Reiter, 1987] Raymond Reiter. A Theory of Diagnosis from First Principles. *Artificial Intelligence* **32**(1): 57–95, April 1987.

[Stalnaker, 1968] Robert C. Stalnaker. A Theory of Conditionals. In N. Rescher (ed.), *Studies in Logical Theory*. Oxford: Oxford University Press, 1968.

[Winslett, 1986] Marianne S. Winslett. Is Belief Revision Harder than You Thought?. *Proc. AAAI-86*, pp. 421–427. Philadelphia, Pa., August 1986.