REVISIONS OF KNOWLEDGE SYSTEMS USING EPISTEMIC ENTRENCHMENT

Peter Gärdenfors
Department of Philosophy
Lund University
S-223 50 Lund
SWEDEN

David Makinson
Les Etangs B2, La Ronce
F-92410, Ville d'Avray
FRANCE

ABSTRACT

A major problem for knowledge representation is how to revise a knowledge system in the light of new information that is inconsistent with what is already in the system. Another related problem is that of contractions, where some of the information in the knowledge system is taken away.

Here, the problems of modelling revisions and contractions are attacked in two ways. First, two sets of rationality postulates or integrity constraints are presented, one for revisions and one for contractions. On the basis of these postulates it is shown that there is a natural correspondence between revisions and contractions.

Second, a more constructive approach is adopted based on the "epistemic entrenchment" of the facts in a knowledge system which determines their priority in revisions and contractions. We introduce a set of computationally tractable constraints for an ordering of epistemic entrenchments.

The key result is a representation theorem which says that a revision method for a knowledge system satisfies the set of rationality postulates, if and only if, there exists an ordering of epistemic entrenchment satisfying the appropriate constraints such that this ordering determines the retraction priority of the facts of the knowledge system. We also prove that the amount of information needed to uniquely determine the required ordering is linear in the number of atomic facts of the knowledge system.

1. PROGRAM

One of the main problems concerning knowledge representation is how to *revise* a knowledge system in the light of new information -- information that may be inconsistent with what is already in the system. When a system is revised some of the old information has to be retracted. The main problems are to determine which information should be given up and how this should be handled computationally (cf. *updates of databases* as studied by Fagin, Ullman and Vardi (1983) and Fagin, Kuper, Ullman and Vardi (1986). Related projects are pursued by Ginsberg (1986), Foo and Rao (1986), and Martins and Shapiro (1986)). Apart from revisions, there is a closely related type of change of a knowledge system which we call *contractions*. Such a change occurs when some of the information in the knowledge system is retracted without adding any new items, for example because it is discovered that the information derives from a faulty source.

We will attack the problems of revisions and contractions of knowledge systems by two methods. In Section 2 we will present two sets of *rationality postulates* for these processes and outline the connections between the postulates. These postulates should be viewed as *dynamic intregrity constraints* or *transition laws*. The postulates will depend on viewing a knowledge set as a *set* of (logically related) facts. This part is mainly a summary of earlier work (Gärdenfors (1984,1988), Alchourrón, Gärdenfors and Makinson (1985), Makinson (1985)).

In Section 3 we adopt a more *constructive* approach. It will be assumed that apart from the logical relations, a knowledge set has some additional structure which makes it possible to determine the *epistemic entrenchment* of the facts in the system. The epistemic entrenchment of a fact represents how important it is for problem solving or planning on the basis of the knowledge system and in this way determines the database *priority* of the fact. We introduce a set of logical constraints for an ordering of epistemic entrenchment.

The key result of the paper is a representation theorem which says, roughly, that a revision method for a knowledge set satisfies the set of rationality postulates presented in Section 2, if and only if, there exists an ordering of epistemic entrenchment satisfying the logical constraints such that this ordering determines the retraction priority of the facts. We also prove that, due to the logical constraints on the ordering of epistemic entrenchment, the amount of information needed to uniquely determine the required ordering (and thereby also to determine the revision method) is *linear* in the number of atomic facts of the knowledge set. We conclude by some comments on implementations of revision (and contraction) methods.

The proof of theorems are grouped together in an appendix.

2. POSTULATES FOR REVISIONS AND CONTRACTIONS

We assume that the items of a knowledge system are expressed in some language L which is closed under applications of the boolean operators - (negation), & (conjunction), v (disjunction), and -> (implication). Further details of the language will be left unspecified here. We will use A, B, C etc. as variables over sentences in L. A *knowledge set* is a set K of sentences in L which satisfies the integrity constraint:

(I) If K logically entails B, then $B \in K$.

In logical parlance, this means that a knowledge set is a *theory* which can be seen as a partial description of the world (knowledge sets are called "belief sets" in Gärdenfors (1988)). "Partial" because in general there are sentences A such that neither A nor -A are in K. The sentences in a knowledge set will also be called *facts*. Here we assume that the underlying logic includes classical propositional logic and that it is compact. If K logically entails A we will write this as K

|- A. We also assume that |- satisfies "disjunction in the premises", i.e. that K U {B v C} |- A whenever both K U {B} |- A and K U {C} |- A.

By classical logic, whenever K is inconsistent, then K |- A for every sentence A of the language L. This means that there is exactly one inconsistent knowledge set under our definition, namely the set of all sentences of L. We introduce the notation $K_{\perp}$ for this knowledge set. Clearly, $K_{\perp}$ is useless for information handling purposes, but for technical reasons we have included it as a knowledge set. When we want to exclude it from consideration, we may simply speak of *consistent* knowledge sets.

The integrity constraint (I) that a knowledge set is supposed to be closed under logical consequence will cause problems when it comes to implementing a system, since there are in general infinitely many logical consequences to take care of. We will return to implementation problems at the end of Section 3.

We believe that this formal framework is appropriate for representing knowledge systems for computational approaches. This approach is propounded in Fagin, Ullman and Vardi (1983), Reiter (1984), Gärdenfors (1984, 1988) among others. It has the advantages of handling facts, logical integrity constraints and derivation rules in a uniform way and it is a convenient way of modelling partial information.

The logic determined by |- and the integrity constraint (I) specify the *statics* of knowledge sets. We now turn to their *dynamics*. What we need are methods for updating knowledge sets. Three kinds of updates will be discussed here:

(i) *Expansion*: A new sentence together with its logical consequences is *added* to a knowledge set K regardless of the consequences of the larger set so formed. The knowledge set that results from expanding K by a sentence A will be denoted K^+A.

(ii) *Revision*: A new sentence that is *inconsistent* with a knowledge set K is added, but in order that the resulting knowledge set be consistent some of the old sentences in K are deleted. The result of revising K by a sentence A will be denoted K*A.

(iii) *Contraction*: Some sentence in K is retracted without adding any new facts. In order that the resulting system satisfies (I) some other sentences from K must be given up. The result of contracting K with respect to the sentence A will be denoted K^-A.

Expansions of knowledge sets can be handled comparatively easy. K^+A can simply be defined as the logical closure of K together with A:

(Def +) $K^+A = \{B: K \cup \{A\} \vdash B\}$

As is easily shown, K^+A defined in this way will satisfy (I) and will be consistent when A is consistent with K.

It is not possible to give a similar explicit definition of revisions and contractions in logical and set-theoretical notions only. To see the problems for revisions, consider a knowledge set K which contains the sentences A, B, A&B -> C and their logical consequences (among which is C). Suppose that we want to revise K by adding -C. Of course, C must be deleted from K when forming K*-C, but also at least one of the sentences A, B, or A&B -> C must be given up in order to maintain consistency. There is no purely *logical* reason for making one choice rather than the other, but we have to rely on additional information about these sentences. Thus, from a logical point of view, there are several ways of specifying the revision K*A. What is needed here is a (computationally well defined) method of determining the revision. We will handle this technically by using the notion of a *revision function* * which has two arguments, a knowledge set K and a sentence A, and which has as value the knowledge system K^*A.

The contraction process faces parallel problems. A concrete example is provided by Fagin, Ullman and Vardi (1983, p. 353):

> "Consider for example a relational database with a ternary relation SUPPLIES, where a tuple *<a,b,c>* means that supplier *a* supplies part *b* to project *c*. Suppose now that the relation contains the tuple *<Hughes, tiles, Space Shuttle>*, and that the user asks to delete this tuple. A simpleminded approach wouuld be to just go ahead and delete the tuple from the relation. However, while it is true that Hughes does not supply tiles to the Space Shuttle project anymore, it is not clear what to do about three other facts that were implied by the above tuple, i.e. that Hughes supplies tiles, that Hughes supplies parts to the Space Shuttle project, and that the Space Shuttle project uses tiles. In some circumstances it might not be a bad idea to replace the deleted tuple by three tuples with null values:
>
> *<Hughes,tiles,NULL>*
> *<Hughes,NULL,Space Shuttle>*
>
> and
>
> *<NULL,tiles,Space Shuttle>.*
>
> The common denominator to both examples is that the database is not viewed merely as a collection of atomic facts, but rather as a collection of facts from which other facts can be derived. It is the interaction between the updated facts and the derived facts that is the source of the problem."

Also here we introduce the concept of a *contraction function* - which has the same two arguments as before, i.e. a knowledge set K and a sentence A, and which produces as value the knowledge set $K^{-}A$. Later in this section we will show that the problems of revision and contraction are closely related -- being two sides of the same coin.

From a computational point of view, the ultimate goal is to develop algorithms for computing appropriate revision and contraction functions for an arbitrary knowledge set. However, in order to know whether an algorithm is sucessful or not it is necessary to determine what an "appropriate" function is. Our standards for revision and contraction functions will be two sets of *rationality postulates.* One guiding idea is that the revision K^*A of K with respect to A represents the *minimal change* of K needed to accomodate A consistently. Here, we will just list the postulates, with only a few comments. The postulates are defended and further investigated in Gärdenfors (1984, 1988), Alchourrón, Gärdenfors, and Makinson (1985), and Makinson (1985).

Postulates for revisions:

(K*1) K^*A is a knowledge set
(K*2) $A \in K^*A$
(K*3) $K^*A \subseteq K^+A$
(K*4) If $-A \notin K$, then $K^+A \subseteq K^*A$
(K*5) $K^*A = K_\perp$ only if $\vdash -A$
(K*6) If $\vdash A \leftrightarrow B$, then $K^*A = K^*B$
(K*7) $K^*A\&B \subseteq (K^*A)^+B$
(K*8) If $-B \notin K^*A$, then $(K^*A)^+B \subseteq K^*A\&B$

Since (K*1) - (K*6) do not refer to revisions with respect to compound sentence these postulates will be called the *basic* postulates for revision. It can be shown that in the presence of the basic postulates, the conjunction of (K*7) and (K*8) is equivalent to the following principle (Gärdenfors (1988), principle 3.3.6):

(K*V) $K^*A \vee B = K^*A$ or $K^*A \vee B = K^*B$ or $K^*A \vee B = K^*A \cap K^*B$.

Postulates for contractions:

(K-1) K^-A is a knowledge set
(K-2) $K^-A \subseteq K$
(K-3) If $A \notin K$, then $K^-A = K$
(K-4) If not $\vdash A$, then $A \notin K^-A$
(K-5) $K \subseteq (K^-A)^+A$
(K-6) If $\vdash A \leftrightarrow B$, then $K^-A = K^-B$
(K-7) $K^-A \cap K^-B \subseteq K^-A\&B$
(K-8) If $A \notin K^-A\&B$, then $K^-A\&B \subseteq K^-A$

Again, a motivating idea for these postulates, in particular (K-5), (K-7) and (K-8), is that K^-A represents the *minimal change* of K needed to retract the fact A under the integrity constraint (I). There is an extended discussion of the postulate (K-5) of "recovery" for contraction in Makinson (1987). (K-1) - (K-6) will be called the basic postulates for contractions. In the presence of these, the conjunction of (K-7) and (K-8) can be shown to be equivalent to the following principle (Alchourrón, Gärdenfors, and Makinson (1985), Gärdenfors (1988), principle 3.4.7):

(K-V) Either $K^-A\&B = K^-A$ or $K^-A\&B = K^-B$ or $K^-A\&B = K^-A \cap K^-B$.

We next turn to a study of the connections between revision and contraction functions. A revision of a knowledge set can be seen as a *composition* of a contraction and an expansion. More precisely: In order to construct the revision K^*A, one first contracts K with respect to -A and then expands K^--A by A. Formally, we have the following definition:

(Def *) $K^*A = (K^- {-A})^+A$

That this definition is appropriate is shown by the following result:

Theorem 1: If a contraction function - satisfies (K-1) to (K-6), then the revision function * obtained from (Def *) satisfies (K*1) - (K*6). Furthermore, if (K-7) also is satisfied, (K*7) will be satisfied for the defined revision function; and if (K-8) also is satisfied, (K*8) will be satisfied for the defined revision function.

Conversely, contractions can be defined in terms of revisions. The idea is that a sentence B is accepted in the contraction K^-A if and only if B is accepted both in K and in K^*-A. Formally:

(Def -) $K^-A = K \cap K^* {-A}$

Again, this definition is supported by the following result:

Theorem 2: If a revision function * satisfies (K*1) to (K*6), then the contraction function - obtained from (Def -) satisfies (K-1) - (K-6). Furthermore, if (K*7) is satisfied, (K-7) will be satisfied for the defined contraction function; and if (K*8) is satisfied, (K-8) will be satisfied for the defined contraction function.

Theorems 1 and 2 show that the two sets of postulates for revision and contraction functions are *interchangeable* and a method for constructing one of the functions would automatically, via (Def *) or (Def -), yield a construction of the other function.

It should be noted, however, that the rationality postulates do not uniquely determine a revision or a contraction function. On the other hand, we claim that the postulates (K*1) - (K*8) and (K-1) - (K-8) exhaust what can be said about revisions and contraction in logical and set-theoretical terms only. This means that we must seek further information about the epistemic status of the elements of a knowledge state in order to solve the uniqueness problem. This project will be the topic of next section.

3. EPISTEMIC ENTRENCHMENT

Even if all sentences in a knowledge set are accepted or considered as facts (so that they are assigned maximal probability), this does not mean that all sentences are are of equal value for planning or problem-solving purposes. Certain pieces of our knowledge and beliefs about the world are more important than others when planning future actions, conducting scientific investigations, or reasoning in general. We will say that some sentences in a knowledge system have a higher degree of *epistemic entrenchment* than others. This degree of entrenchment will, intuitively, have a bearing on what is abandoned from a knowledge set, and what is retained, when a contraction or a revision is carried out.

From an epistemological point of view, some may see the notion of epistemic entrenchment as more fundamental than that of contraction. Some may, conversely, see contraction as being more fundamental, and some, finally, may remain sceptical of any such priorization. From a purely formal point of view, the most promising direction is perhaps that which takes the relation of epistemic entrenchment as basic. Accordingly, we begin this section by presenting a set of postulates for epistemic entrenchment which will serve as a basis for a *constructive definition* of appropriate revision and contraction functions.

The guiding idea for the construction is that when a knowledge system K is revised or contracted, the sentences in K that are given up are those having the *lowest* degrees of epistemic entrenchment (Fagin, Ullman and Vardi (1983), pp. 358 ff., introduce the notion of "database priorities" which is closely related to the idea of epistemic entrenchment and is used in a similar way to update knowledge sets. However, they do not present any axiomatization of this notion).

We will not assume that one can quantitatively measure degrees of epistemic entrenchment, but only work with *qualitative* properties of this notion. One reason for this is that we want to emphasize that the problem of uniquely specifying a revision function (or a contraction function) can be solved, assuming only very little structure on the knowledge systems apart from their logical properties. Another, quite different, way of doing this was described by Alchourron and Makinson (1985).

If A and B are sentences in L, the notation $A \leq B$ will be used as a shorthand for "B is at least as epistemically entrenched as A". The strict relation $A < B$, representing "B is epistemically more entrenched than A", is defined as "$A \leq B$ and not $B \leq A$. Note that the relation $\leq$ is only defined *in relation to a given* K -- different knowledge sets may be associated with different orderings of epistemic entrenchment.

Postulates for epistemic entrenchment:

(EE1)	If $A \leq B$ and $B \leq C$, then $A \leq C$	(transitivity)
(EE2)	If $A \vdash B$, then $A \leq B$	(dominance)
(EE3)	For any A and B, $A \leq A\&B$ or $B \leq A\&B$	(conjunctiveness)
(EE4)	When $K \neq K_{\perp}$, $A \notin K$ iff $A \leq B$, for all B	(minimality)
(EE5)	If $B \leq A$ for all B, then $\vdash A$	(maximality)

The justification for (EE2) is that if A logically entails B, and either A or B must be retracted from K, then it will be a smaller change to give up A and retain B rather than to give up B, because then A must be retracted too, if we want the revised knowledge set to satisfy the integrity constraint (I). The rationale for (EE3) is as follows: If one wants to retract A&B from K, this can only be achieved by giving up either A or B and, consequently, the informational loss incurred by giving up A&B will be the same as the loss incurred by giving up A or that incurred by giving up B. (Note that it follows already from (EE2) that $A\&B \leq A$ and $A\&B \leq B$). The postulates (EE4) and (EE5) only take care of limiting cases: (EE4) requires that sentences already not in K have minimal epistemic entrenchment in relation to K; and (EE5) says that only logically valid sentences can be maximal in $\leq$. (The converse of (EE5) follows from (EE2), since if $\vdash A$, then $B \vdash A$, for all B).

We note the following simple consequences of these postulates:

Lemma 3: Suppose the ordering $\leq$ satisfies (EE1) - (EE3). Then it also has the following properties:
(i) $A \leq B$ or $B \leq A$ (connectivity);
(ii) If $B\&C \leq A$, then $B \leq A$ or $C \leq A$;
(iii) $A < B$ iff $A\&B < B$.
(iv) If $C \leq A$ and $C \leq B$, then $C \leq A\&B$.
(v) If $A \leq B$, then $A \leq A\&B$.

Note that in view of (i), $A < B$ may be more simply defined as not $B \leq A$.

The main purpose of this article is to show the connections between orderings of epistemic entrenchment and contraction and revision functions. We will accomplish this by providing two conditions, one of which determines an ordering of epistemic entrenchment assuming a contraction function and knowledge set as given, and the other of which determines a contraction function assuming an ordering of epistemic entrenchment and knowledge set as given. The first condition is:

(C$\leq$) $A \leq B$ if and only if $A \notin K^{-}_{A\&B}$ or $\vdash A\&B$

The idea underlying this definition is that when we contract K with respect to A&B we are forced to give up A or B (or both) and A should be retracted just in case B is at least as epistemically entrenched as A. In the limiting case when both A and B are logically valid, they are of equal epistemic entrenchment (in conformity with (EE2)).

The second, and from a constructive point of view most central, condition gives an explicit definition of a contraction function in terms of the relation of epistemic entrenchment:

(C-) $B \in K^{-}_{A}$ if and only if $B \in K$ and either $A < A \vee B$ or $\vdash A$

Perhaps the best way of motivating this condition (apart from the fact that it "works" in the sense of theorems 4 - 6 below) is to note that if B is in K (and K is consistent), then the epistemic entrenchment of -B will always be less than that of B according to (EE4), so the relation $-B \leq B$ will give no clue as to whether B should be in K^{-}_{A} or not. We have to look for other formulas involving B and A. According to (C$\leq$), $A < B$ is essentially the same as $B \in K^{-}_{A\&B}$. If we replace B by AvB, we get AvB ∈ K^-A&(AvB) iff A < AvB (assuming that A is in K). But K^-A&(AvB) is the same as K^-A. And, given the under-standing that the contraction operation should satisfy (K-5), we have also that -AvB ∈ K^-A; hence, for any B ∈ K, we have AvB ∈ K^-A iff B ∈ K^-A. Putting this together gives: If A ∈ K, then for any B ∈ K, B ∈ K^-A iff A < AvB. (Note that this argument does not stand completely on its own feet, since it presumes (C≤) and the validity of several of the basic postulates for contraction including most conspicuously (K-5)). The case that A ∉ K is handled by noting that if B ∈ K, then A < AvB follows from the postulates for epistemic entrenchment so then (C-) says just K^-A = K as desired.

As mentioned above, one might take the ordering of epistemic entrenchment to be more fundamental than a contraction function or a revision function. Condition (C-) now provides us with a tool for explicitly defining a contraction function in terms of the ordering ≤. An encouraging test of the appropriateness of such a definition is the following theorem, which is the central result of this article:

Theorem 4: If an ordering ≤ satisfies (EE1) - (EE5), then the contraction function which is uniquely determined by (C-) satisfies (K-1) - (K-8) as well as the condition (C≤).

Indirectly, theorem 4 provides us with a consistency proof for the set of postulates for contractions (and thereby also for the postulates for revisions via theorem 2) since it is easy to show, using finite models, that the set (EE1) - (EE5) is consistent.

Conversely, we can show that if we start from a given contraction function and determine an ordering of epistemic entrenchment with the aid of condition (C≤), the ordering will have the desired properties:

Theorem 5: If a contraction function - satisfies (K-1) - (K-8), then the ordering ≤ that is uniquely determined by (C≤) satisfies (EE1) - (EE5) as well as the condition (C-).

(A weaker version of theorems 4 and 5 was proved in a very roundabout way in Gärdenfors (1988) as Theorem 4.28. That proof depended on the results of Grove (1986).)

Theorems 4 and 5 imply that conditions (C-) and (C≤) are *interchangeable* in the following sense: Let **C** be the class of contraction functions satisfying (K-1) - (K-8) and **E** the class of orderings satisfying (EE1) - (EE5). Let C^ be a map from **E** to **C** such that C^(≤) = - is the contraction function determined by (C-) for a given ordering ≤; and let E^ be a map from **C** to **E** such that E^(-) = ≤ is the ordering determined by (C≤) for a given contraction function -. We have as an immediate consequence of theorems 4 and 5 that:

Corollary 6: For all - in **C**, C^E^(-) = -; and for all ≤ in **E**, E^C^(≤) = ≤.

From an epistemological point of view, these results suggest that the problem of constructing appropriate contraction and revision functions can be *reduced* to the problem of providing an appropriate ordering of epistemic entrenchment. Furthermore, condition (C-) gives an *explicit* answer to which sentences are included in the contracted knowledge set, given the initial knowledge set and an ordering of epistemic entrenchment. From a computational point of view, applying (C-) is trivial, once the ordering ≤ of the elements of K is given.

We will conclude the paper by some remarks on the computational aspects of adding an ordering of epistemic entrenchment to the representation of a knowledge system. One important question concerns the *amount of information* that needs to be specified in order to determine an ordering of epistemic entrenchment over a knowledge set K. In all applications, knowledge sets will be *finite* in the sense that the consequence relation |- partitions the elements of K into a finite number of equivalence classes. In algebraic terms, the set of these equivalence classes will be isomorphic to a finite Boolean algebra. This isomorphism is helpful when it comes to implementing a representation of a knowledge set. A finite knowledge set can, for example, be described via its set of *atoms* or via its set of *dual atoms* (which correspond to maximal disjunctions of atoms).

A Boolean algebra with n atoms has 2^n elements and, in general an ordering over a Boolean algebra must be specified for all these elements so that there exist $(2^n)!$ different total orderings of such an algebra (the number of pre-orderings is even larger). However, the postulates (EE1) - (EE5) introduce constraints on the ordering $\leq$ so that the number of orderings satisfying these postulates will be much smaller. The following result shows that the number of orderings over a Boolean algebra with 2^n elements is only n!:

Theorem 7: Let K be a finite knowledge set, and let T be the set of all top elements of K, i.e. all dual atoms of K. Then any two relations $\leq$ and $\leq'$, each satisfying (EE1) - (EE5), that agree on all pairs of elements in T are identical.

The computational interpretation of this result is that in order to specify the ordering of epistemic entrenchment over a knowledge set K containing 2^n elements, and thus a complete contraction function over K according to Theorem 4, one needs only specify the ordering of n elements from K. This means that the information required is *linear* in the number of atomic facts in K.

Appendix: Verification of theorems

Theorem 1: If a contraction function - satisfies (K-1) to (K-6), then the revision function * obtained from (Def *) satisfies (K*1) - (K*6). Furthermore, if (K-7) also is satisfied, (K*7) will be satisfied for the defined revision function; and if (K-8) also is satisfied, (K*8) will be satisfied for the defined revision function.

Proof: (a) Suppose that - satisfies (K-1) - (K-6), and that * is defined by (Def *). Then immediately, by the definition of +, K*A is a knowledge set containing A, giving us conditions (K*1) and (K*2). For (K*3), note that K^--A $\subseteq$ K by (K-2), so by the monotony of |- and the definition of +, $(K^-\text{-A})^+$A $\subseteq$ K^+A, that is K*A $\subseteq$ K^+A as required. For (K*4), suppose -A $\notin$ K. Then by (K-3), K $\subseteq$ K^--A, so by the monotony of |- again, K^+A $\subseteq$ $(K^-\text{-A})^+$A = K*A by (Def *) as required. For (K*5), suppose K*A = $K_\perp$. We want to show |- -A, and by (K-4) it will suffice to show -A $\in$ K^--A. Because K*A = $K_\perp$ we have in particular -A $\in$ K*A = $(K^-\text{-A})^+$A = {B: K^--A U {A} |- B}. Thus K^--A U {A} |- -A, so since |- includes classical logic and is closed under disjunction in the antecedent, K^--A |- -A; so because K^--A is a knowledge set, -A $\in$ K^--A as desired. For (K*6), suppose |- A<->B. Then |- -A <-> -B by classical logic, K^--A = K^--B by (K-6), so K*A = K*B by (Def *) and classical logic.

(b) The derivations of (K*7) and (K*8) from (K-7) and (K-8) respectively in addition to the basic postulates (K-1) to (K-6) for contraction, are a little more complex. They are given in full in Alchourrón, Gärdenfors and Makinson (1985), Observations 3.1 and 3.2.

Theorem 2: If a revision function * satisfies (K*1) to (K*6), then the contraction function - obtained from (Def -) satisfies (K-1) - (K-6). Furthermore, if (K*7) is satisfied, (K-7) will be satisfied for the defined contraction function; and if (K*8) is satisfied, (K-8) will be satisfied for the defined contraction function.

Proof: Suppose that * satisfies (K*1) to (K*6) and that - is defined by (Def -). Then K^-A is the intersection of two knowledge sets which, as is well known, is always a knowledge set, giving (K-1). Clearly, K^-A $\subseteq$ K, giving (K-2). For (K-3), suppose A $\notin$ K. We need to show that K^-A = K, and so given (K-2) already checked, it suffices to show that K $\subseteq$ K^-A = K $\cap$ K*-A, so it suffices to show K $\subseteq$ K*-A. But because A $\notin$ K and K is a knowledge set, --A $\notin$ K so by (K*4) K^+-A $\subseteq$ K*-A; and by the definition of +, clearly $K \subseteq K^+-A$. Putting these together gives $K \subseteq K^*-A$ as desired. For (K-4), suppose $A \in K^-A$; we need to show |- A. Because $A \in K^-A$ we have by (Def -) that $A \in K^*-A$, and clearly by (K*2) also $-A \in K^*-A$. Thus, using (K*1), $K^*-A = K_\perp$. Applying (K*5) gives us |- --A, that is |- A as desired. For (K-5), we need to show that $K \subseteq (K \cap K^*-A)^+A$. Suppose $B \in K$; we need to show that $(K \cap K^*-A) \cup \{A\}$ |- B. Because |- includes classical logic and is closed under disjunction in the antecedent, this is the same as $K \cap K^*-A$ |- -A v B. But because $B \in K$ and K is a knowledge set, $-A \vee B \in K$; and by (K*2) and (K*1) we have also $-A \vee B \in K^*-A$, so $-AvB \in K \cap K^*-A$ and thus $K \cap K^*-A$ |- -A v B as desired. For (K-6) from (K*6) the argument is similar to the one we gave in the reverse direction.

(b) The derivations of (K-7) and (K-8) from their counterparts for * are a little more complex, and are also given in full in Alchourrón, Gärdenfors and Makinson (1985), Observations 3.1 and 3.2.

Theorem 4: If an ordering $\leq$ satisfies (EE1) - (EE5), then the contraction function which is uniquely determined by (C-) satisfies (K-1) - (K-8) as well as the condition (C$\leq$).

Proof: Suppose $\leq$ satisfies (EE1) - (EE5) and that the contraction function - is defined by (C-). The condition (K-2), that is $K^-A \subseteq K$, is immediate from (C-). For (K-1), the argument is more subtle. Suppose K^-A |- C; we need to show $C \in K^-A$. By the assumption of compactness of |-, there are $B_1, \ldots, B_n \in K^-A$ with $B_1 \& \ldots \& B_n$ |- C. To show that $C \in K^-A$ it suffices by rule (C-) to show that $C \in K$ and either $A < AvC$ or |- A. Since $B_1, \ldots, B_n \in K^-A \subseteq K$ as already observed and $B_1 \& \ldots \& B_n$ |- C, and K is a knowledge set, we have $C \in K$. Now suppose not |- A, we need to show that $A < AvC$, that is $A \leq AvC$ and not $AvC \leq A$. By lemma 3 (i), connectivity, it suffices to show the latter. First consider the principal case that $n \geq 1$. Now since $B_1, \ldots, B_n \in K^-A$, we know by condition (C-) that either |- A or else for each B_i, $A < AvB_i$. By supposition, not |- A, so we have $A < AvB_i$ and thus not $AvB_i \leq A$ for each B_i. By lemma 3(ii), this gives us not $(AvB_1) \& \ldots \& (AvB_n) \leq A$, and so by (EE2) and classical logic, not $A \vee (B_1 \& \ldots \& B_n) \leq A$. But $B_1 \& \ldots \& B_n$ |- C, so $A \vee (B_1 \& \ldots \& B_n)$ |- A v C so by (EE2) $A \vee (B_1 \& \ldots \& B_n) \leq A \vee C$ so by transitivity not $AvC \leq A$ as desired. In the limiting case that $n = 0$, so that |- C, we have |- AvC so by (EE2), $D \leq AvC$ for all D; whilst since not |- A, we have by (EE5) that not $D \leq A$ for some D; so that by transitivity not $AvC \leq A$ as desired.

For (K-3), suppose $A \notin K$. We need to show $K \subseteq K^-A$. Let $B \in K$. By rule (C-) it suffices to show $A < AvB$, that is $A \leq AvB$ and not $AvB \leq A$, so by lemma 3 (i) the latter suffices. Since $A \notin K$ we have $K \notin K$. so by (EE4), $A \leq D$ for all D. On the other hand, since $B \in K$, (EE4) also gives us not $B \leq D$ for some D. Hence by transitivity not $B \leq A$ and so by (EE2) and transitivity, not $AvB \leq A$, as desired.

For (K-4) suppose not |- A; we need to show $A \notin K^-A$, which by (C-) is the same as either $A \notin K$ or both not $A < AvA$ and not |- A. Suppose then that $A \in K$. Then since $A \notin K$ and K is a knowledge set, not |- A. Hence it suffices to show not $A < AvA$, for which it suffices to show $AvA \leq A$. But AvA |- A, so by (EE2) $AvA \leq A$ as desired.

For (K-5), suppose $B \in K$; we need to show $K^-A \cup \{A\}$ |- B. By the hypothesis on |-, it suffices to show $-AvB \in K^-A$, i.e. by rule (C-) that $-AvB \in K$ and either $A < Av(-AvB)$ or |- A. Because $B \in K$ and K is a knowledge set, we have the former. For the latter, suppose not |- A. We need to show $A \leq Av(-AvB)$ and not $Av(-AvB) \leq A$, and by lemma 3(i), it suffices to show the latter. But Av(-AvB) is a tautology, so |- Av(-AvB), so by (EE2), $D \leq Av(-AvB)$ for all D, whilst by (EE5) not $D \leq A$ for some D, so that by transitivity, not $Av(-AvB) \leq A$ as desired.

For (K-6), suppose |- A <-> B. We show $K^-A \subseteq K^-B$; the converse is similar. Suppose $C \in K^-A$. Then by rule (C-), $C \in K$ and either $A < AvC$ or |- A. Since |- A <-> B we have |- AvC <-> BvC so using (EE2), $A < AvC$ iff $B < BvC$. Thus we have $C \in K$ and either $B < BvC$ or |- B, so by rule (C-), $C \in K^-B$ as desired.

For (K-7), suppose $C \in K^{-}A \cap K^{-}B$. Then $C \in K$, and either |- A or A < AvC, and also either |- B or B < BvC. We want to show that $C \in K^{-}A\&B$, so by rule (C-) it suffices to show that either |- A&B or A&B < (A&B) v C. Suppose not |- A&B. In the case that |- A, we have |- B <-> A&B, so by (K-6) already verified $K^{-}B = K^{-}A\&B$ and so $C \in K^{-}A\&B$ as desired. Likewise when |- B we have $K^{-}A = K^{-}A\&B$ and so $C \in K^{-}A\&B$. Hence we may assume both not |- A and not |- B. So because $C \in K^{-}A$ and $C \in K^{-}B$ we have not AvC ≤ A and not BvC ≤ B, so using (EE2) and transitivity, not AvC ≤ A&B and not BvC ≤ A&B, so by lemma 3 (ii) not (AvC) & (BvC) ≤ A&B, so by (EE2) and transitivity, not (A&B) v C ≤ A&B, so that A&B < (A&B) v C. Because also $C \in K$, it follows by (C-) that $C \in K^{-}A\&B$ as desired.

For (K-8), suppose $A \notin K^{-}A\&B$. We want to show that $K^{-}A\&B \subseteq K^{-}A$. As limiting cases, note that if |- A then by (K-2) and (K-5) already verified, $K^{-}A\&B \subseteq K \subseteq (K^{-}A)^{+}A = K^{-}A$, and that if $A \notin K$, then by (K-3) already verified, $K^{-}A\&B = K = K^{-}A$, as desired. Hence we may suppose without loss of generality that $A \in K$ and not |- A. Now suppose $C \notin K^{-}A$; we need to show $C \notin K^{-}A\&B$. If $C \notin K$, then we are done, using (K-2) already established, so we suppose $C \in K$. Because $C \notin K^{-}A$ whilst $C \in K$ and not |- A, we have by rule (C-) that not A < AvC, so using lemma 3 (i), AvC ≤ A. Because not |- A, we also have not |- A&B, so by rule (C-), in order to show $C \notin K^{-}A\&B$ it suffices to show (A&B) v C ≤ A&B. For this it will suffice to show A ≤ A&B, for then we have the sequence (A&B) v C |- AvC ≤ A ≤ A&B, so that (A&B) v C ≤ A&B by (EE2) and transitivity. To show that A ≤ A&B, we appeal to the initial hypothesis that $A \notin K^{-}A\&B$. Because $A \in K$ and not |- A&B, we have by (C-) that not A&B < (A&B) v A, so using lemma 3 (i), A |- (A&B) v A ≤ A&B, so A ≤ A&B by (EE2) and transitivity, as desired.

Finally, to verify the condition (C≤), suppose first that A ≤ B whilst $A \in K^{-}A\&B$; we need to show |- A&B. Since $A \in K^{-}A\&B$ we have by rule (C-) that $A \in K$ and either |- A&B or (A&B) < (A&B) v A. But since A ≤ B we have by lemma 3 (v) that A ≤ A&B, so using (EE2) and transitivity, (A&B) v A ≤ A&B, so not A&B < (A&B) v A and hence |- A&B as desired. For the converse, suppose either $A \notin K^{-}A\&B$ or |- A&B. In the latter case, |- B so A ≤ B as required by (EE2). Suppose, then, that not |- A&B. Since $A \notin K^{-}A\&B$ and not |- A&B, we have by rule (C-) that either $A \notin K$ or not A&B < (A&B) v A. The former case gives us A ≤ B as required, by (EE4). The latter case gives us, using lemma 3 (i), A |- (A&B) v A ≤ A&B |- B, so that by (EE2) and transitivity we have again A ≤ B as desired.

Theorem 5: If a contraction function - satisfies (K-1) - (K-8), then the ordering ≤ that is uniquely determined by (C≤) satisfies (EE1) - (EE5) as well as the condition (C-).

Proof: Suppose the contraction function satisfies (K-1) to (K-8), and that the relation ≤ is defined by the rule (C≤). It will be useful to note first that it follows immediately from (K-5) and (K-2) that when |- A, then $K^{-}A = K$. We leave the verification of (EE1), transitivity, until last, as it is considerably more complex than the others.

For (EE2), suppose A |- B. To show A ≤ B, it suffices by rule (C≤) to show that either |- A&B or $A \notin K^{-}A\&B$. But if $A \in K^{-}A\&B$, then by (K-1) and the hypothesis A |- B we have $A\&B \notin K^{-}A\&B$, so by (K-4), |- A&B.

For (EE3), we need to show that either $A \notin K^{-}A\&(A\&B)$ or $B \notin K^{-}B\&(A\&B)$ or |- A&(A&B) or |- B&(A&B). By classical logic and (K-6), it suffices to show either $A \notin K^{-}A\&B$ or $B \notin K^{-}A\&B$ or |- A&B. But if the first two fail, we have using (K-1) that $A\&B \in K^{-}A\&B$, so by (K-4), |- A&B as desired.

For (EE4) suppose, $K \neq K_{\perp}$, and suppose first that $A \notin K$. Then by (K-2) we have $A \notin K^{-}A\&B$, so by the rule (C≤), A ≤ B for all B. For the converse, suppose that A ≤ B for all B. Then in particular A ≤ -A, so by (C≤), either $A \notin K^{-}A\&{-A}$ or |- A&-A. But since $K \neq K_{\perp}$ we have using (K-1) that $A\&{-A} \notin K$ and so by (K-1) again, not |- A&-A, so that $A \notin K^{-}A\&{-A} = K$ by (K-3), as desired.

For (EE5), suppose $B \leq A$ for all B. Then by the rule ($C\leq$), for all B either $\vdash B\&A$ or $B \notin K^{-}A\&B$. Choose B with $\vdash B$, for example a classical tautology. Then by (K-1), $B \in K^{-}A\&B$, so $\vdash B\&A$ and so $\vdash A$ as desired.

For (EE1), suppose for reductio ad absurdum that $A \leq B$, $B \leq C$, but not $A \leq C$. Using ($C\leq$) this gives us: either $\vdash A\&B$ or $A \notin K^{-}A\&B$; either $\vdash B\&C$ or $B \notin K^{-}B\&C$; $A \in K^{-}A\&C$ and not $\vdash A\&C$. First, we note that not $\vdash A\&B$. For if $\vdash A\&B$, then $\vdash B$, so by (K-1) $B \in K^{-}B\&C$ so by our second hypothesis $\vdash B\&C$, so $\vdash A\&C$ contradicting our third hypothesis. Next we note that not $\vdash B\&C$. For if B&C, then $\vdash C$ so that using (K-6), $K^{-}A\&C = K^{-}A$, so using our third hypothesis, $A \in K^{-}A$; so by (K-4) $\vdash A$, so $\vdash A\&C$ again contradicting our third hypothesis.

Since not $\vdash A\&B$ and not $\vdash B\&C$ we have: $A \notin K^{-}A\&B$, $B \notin K^{-}B\&C$, but $A \in K^{-}A\&C$. We obtain a contradiction from this triad by first using (K-7) to show $A \in K^{-}A\&B\&C$, and then using (K-8) twice to show the opposite. For the first leg, note that A&B&C is truth-functionally equivalent to (A&C)&(-AvB), so by (K-7) to show $A \in K^{-}A\&B\&C$ it suffices to show $A \in K^{-}A\&C$ and also $A \in K^{-}\text{-}AvB$. The former holds by hypothesis. For the latter, note that the former implies by (K-2) that $A \in K$ so that by (K-5) $K^{-}\text{-}AvB \cup \{\text{-}AvB\} \vdash A$, so using the assumption that $\vdash$ includes truth-functional logic and is closed under disjunction of the antecendent, $K^{-}\text{-}AvB \vdash A$; so that by (K-1), $A \in K^{-}\text{-}AvB$ as desired. This shows that $A \in K^{-}A\&B\&C$.

For the second leg, to show $A \notin K^{-}A\&B\&C$ it suffices, given the information that $A \notin K^{-}A\&B$, to show that $K^{-}A\&B\&C \subseteq K^{-}A\&B$. Hence by (K-8) it suffices to show $A\&B \notin K^{-}A\&B\&C$. But since not $\vdash A\&C$ we have not $\vdash A\&B\&C$, and so by (K-4), $A\&B\&C \notin K^{-}A\&B\&C$. Consequently, using (K-1), either $A\&B \notin K^{-}A\&B\&C$ or $C \notin K^{-}A\&B\&C$. In the first case we are done, so suppose $C \notin K^{-}A\&B\&C$. Then using (K-1) we have $B\&C \notin K^{-}A\&B\&C$, so by (K-8) again, $K^{-}A\&B\&C \subseteq K^{-}B\&C$. Combining this with the information that $B \notin K^{-}B\&C$ gives us $B \notin K^{-}A\&B\&C$ so using (K-1), $A\&B \notin K^{-}A\&B\&C$ in this case too, and we are done.

Finally, to verify the condition (C-), suppose $B \in K^{-}A$. Then we have $B \in K$ immediately by (K-2). Now suppose not $\vdash A$; we want to show $A < AvB$, i.e. $A \leq AvB$ and not $AvB \leq A$. We have the former by (EE2) already verified. For the latter, observe that since not $\vdash A$ we have not $\vdash (AvB)\&A$ and since $B \in K^{-}A$ we have by (K-1) and (K-6) that $AvB \in K^{-}(AvB)\&A$, so by the rule ($C\leq$), not $AvB \leq A$ as desired. For the converse, suppose $B \in K$ and either $A < AvB$ or $\vdash A$. If $\vdash A$, then $K^{-}A = K$ as remarked at the beginning of the proof so since $B \in K$ we have $B \in K^{-}A$ as required. On the other hand, in the case that $A < AvB$ we have that not $AvB \leq A$, so by ($C\leq$), $AvB \in K^{-}(AvB)\&A = K^{-}A$. Moreover, by (K-5), since $B \in K$, $K^{-}A \cup \{A\} \vdash B$ so using (K-1) we have $\text{-}AvB \in K^{-}A$. Putting these together with (K-1) again, we have $B \in K^{-}A$ in this case too, as required.

Theorem 7: Let K be a finite knowledge set, and let T be the set of all top elements of K, i.e. all dual atoms of K. Then any two relations $\leq$ and $\leq'$, each satisfying (EE1) - (EE5), that agree on all pairs of elements in T are identical.

Proof: Suppose that $\leq$ and $\leq'$ are two relations over propositions, each satisfying (EE1) to (EE5). Suppose that $\leq$ and $\leq'$ agree on all pairs of elements in T; we want to show that they are identical. We show that whenever $A \leq B$ then $A \leq' B$; the converse is similar. Suppose $A \leq B$. We begin by disposing of the limiting cases that $A \notin K$, $B \notin K$, $\vdash B$, or $\vdash A$. In the case that $A \notin K$, condition (EE4) applied to $\leq'$ gives us immediately $A \leq' B$ as desired. In the case that $B \notin K$, condition (EE4) applied this time to $\leq$ gives us $B \leq C$ for all C. Since by hypothesis $A \leq B$, we thus have $A \leq C$ for all C, so by (EE4) again applied to $\leq$, $A \notin K$, which gives us the successfully verified first case again. In the case that $\vdash B$, condition (EE2) applied to $\leq'$ gives us immediately that $A \leq' B$ as desired. In the case that $\vdash A$, condition (EE2) applied to $\leq$ gives us $C \leq A$ for all C. Since by hypothesis $A \leq B$, we thus have $C \leq B$ for all C, so by (EE5), $\vdash B$ which gives us the successfully verified third case again.

This leaves us with the principal case that $A,B \in K$, not $\vdash A$ and not $\vdash B$. In this case, boolean considerations tell us that $\vdash A \leftrightarrow (A_1 \& \ldots \& A_n)$ where $1 \leq n$ and all the A_i are top elements of K. Likewise, $\vdash B \leftrightarrow (B_1 \& \ldots \& B_m)$ where $1 \leq m$ and all the B_j are top elements of K. Since $A_1 \& \ldots \& A_n \vdash A$ we have by (EE2) applied to $\leq$ that $A_1 \& \ldots \& A_n \leq A$ so by lemma 3 (ii) applied to $\leq$, $A_i \leq A$ for some $i \leq n$. Also, since $B \vdash B_1 \& \ldots \& B_m$ we have $B \vdash B_j$ for all $j \leq m$, so by (EE2) applied to $\leq$, $B \leq B_j$. Thus, since by hypothesis $A \leq B$, we have $A_i \leq A \leq B \leq B_j$; so that $A_i \leq B_j$ for all $j \leq m$. But since by supposition $\leq$ and $\leq'$ agree on top elements, this gives us $A_i \leq' B_j$ for all $j \leq m$. Thus by lemma 3 (iv) applied to $\leq'$, we have $A_i \leq' B$. Since $A \vdash A_i$ we thus finally have by (EE2) and transitivity that $A \leq' B$ as desired.

References

Alchourrón, C., Gärdenfors, P. and Makinson, D. (1985): "On the logic of theory change: Partial meet contraction functions and their associated revision functions," *Journal of Symbolic Logic 50*, pp. 510-530.

Alchourrón, C. and Makinson, D. (1985): "On the logic of theory change: Safe contraction," *Studia Logica 44*, pp. 405-422.

Fagin, R., Ullman, J.D., and Vardi, M.Y. (1983): "On the semantics of updates in databases," *Proceedings of Second ACM SIGACT-SIGMOD*, Atlanta, pp. 352-365.

Fagin, R., Kuper, G.M., Ullman, J.D., and Vardi, M.Y. (1986): "Updating logical databases," *Advances in Computing Research 3*, pp. 1-18.

Foo, N. and Rao, A. (1986): "DYNABELS", manuscript, Department of Computer Science, Sydney University.

Gärdenfors, P. (1984): "Epistemic importance and minimal changes of belief," *Australian Journal of Philosophy 62*, pp. 136-157.

Gärdenfors, P. (1988): *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, forthcoming as a Bradford Book, MIT Press, Cambridge, Mass.

Ginsberg, M. (1986): "Counterfactuals," *Artificial Intelligence 30*, pp. 35-79.

Grove, A. (1986): "Two modellings for theory change," *Auckland Philosophy Papers*, 1986:13.

Makinson, D. (1985): "How to give it up: A survey of some formal aspects of the logic of theory change," *Synthese 62*, pp. 347-363 and *68*, pp. 185-186.

Makinson, D. (1987): "On the status of the postulate of recovery in the logic of theory change," *The Journal of Philosophical Logic 16*, pp. 383-394.

Martins, J. and Shapiro, S. (1986): "Theoretical foundations for belief revisions," in J.Y. Halpern (ed.) *Theoretical Aspects of Reasoning About Knowledge*, Morgan Kaufmann Publishers, Los Altos, pp. 383-398.

Reiter, R. (1984): "Towards a logical reconstruction of relational database theory," in M.L. Brodie, J. Mylopoulos and J. Schmidt (eds.), *On Conceptual Modelling: Perspectives from Artificial Intelligence, Databases, and Programming Languages*, Springer-Verlag, Berlin, pp. 191-233.