Belief Change: Partial Epistemic Priorities and Logical Non-Omniscience

Li Yan Yuan and Aditya K. Ghose and Randy Goebel
Department of Computing Science
University of Alberta
Edmonton, Alberta, Canada T6G 2H1
Internet: yuan@cs.Ualberta.ca and aditya@cs.Ualberta.ca and goebel@cs.Ualberta.ca
Phone : (403)492-7171 and (403)492-5389 and (403)492-2683

Abstract

In many database applications, designers can easily provide at least some information about the relative importance of the information to be stored and manipulated. While of potentially high value, the ordering information is typically only partial. Here we address the issue of updates in such partially ordered situations, which we call *epistemically stratified databases*.

In the current database theory literature, Alchourrón, Gärdenfors and Makinson (AGM) have proposed a collection of rationality postulates that define rational updates to deductively closed databases. We reformulate the AGM framework to accomodate empistemically stratified databases, and to relax the closure requirement. Our immediate goal is to exploit the use of partial ordering information and to relax the logical omniscient flavour of the closure condition. A more ambitious goal is only hinted at; it is motivated by a desire to develop a more general theory of updates that integrates the naturally related ideas in artificial intelligence and deductive databases.

Our approach begins with the definition of a new contraction operator which uses the information provided by the partial order to make a rational choice amongst the various possible outcomes of the update operation. This operator is shown to satisfy a reformulated set of rationality requirements. A second and similar contraction operator based on a logically omniscient framework is shown to satisfy most of the original AGM rationality postulates for contraction.

1 Introduction

The recent literature on both deductive database updates and belief revision and nonmonotonic reasoning have focussed on the problem of how to rationally choose between logically indistinguishable revised theories or databases. The problem is non-trivial, and is of fundamental significance to both the areas of database management and artificial intelligence (AI).

The focus is usually on two aspects of this problem: (1) the precise formalization of an operator which produces an updated database (as given by an update semantics) or a change in belief set (as given by a belief change operator), and (2), the identification of a set of criteria that any "rational" update or belief change operation should satisfy.

Alchourrón, Gärdenfors and Makinson [6], [1], [7] and [5] have proposed a set of postulates (henceforth the AGM postulates) which are claimed to characterized what is essential about every kind of rational database update or belief change operation. Several update semantics have been evaluated with respect to these postulates by Katsuno and Mendelzon in [8]. They conclude that many existing semantics do not satisfy these postulates. Those that do, have other undesirable characteristics; Dalal's semantics [2], for instance, do satisfy AGM postulates, but use a suspicious criterion (the number of propositional letters by which interpretations differ) to decide what to discard and retain in the updated database.

Part of our motivation derives from several lacunae that remain in the existing work. Most research has considered only *epistemically uniform* databases. In most applications, however, different items of knowledge have different levels of epistemic importance. When a choice has to be made regarding what knowledge should be discarded from a database, it makes intuitive sense to discard knowledge with a lower epistemic status.

Fagin, Ullman and Vardi [3] have considered epistemically non-uniform databases, but the semantics they proposed assumed that a pre-specified *total order* on all items of knowledge. Similarly, Gärdenfors and Makinson [7] have proposed a revision operator for belief sets which are totally ordered. But total orders suggest complete information about application domains, and are difficult, if not impossible, to articulate. Partial orders may be all we can expect (cf. Doyle and Wellman [**?**]).

With eventual practical application in mind, another serious consideration is the disciplined creation and maintenance of logical support for the items of knowledge comprising a database. Part of the concern is independent of so-called "foundational" versus "coherence" models of update, where the former suggests that updates preserve premises as is possible and the latter focusses on some form of minimal change. We view this difference as largely an artifact of syntactic update theories, but suggest that the logical relationships, whether considered proof theoretically or semantically, will have to be efficiently manipulated for *any* practical deployment of a update or revision operator.

In the literature that specifically deals with foundational change, Nebel [12], Fuhrmann [4] and Nayak [11] have considered updates in the presence of partial orders, and their approaches have several obvious problems which we address. A major concern is that most of these approaches assume that change is performed by logically omniscient agents. As has been extensively documented in the literature on logics of knowledge and belief, such an assumption can lead to some obviously unintuitive results (e.g., [**?**]).

We examine some of the problems with the logically omniscient framework of deductively closed belief sets on which the AGM postulates are based and, using results obtained in [14], propose a *partial closure* of beliefs as an alternative to full deductive closure. This avoids some of the problems, but requires some adjustment of the AGM postulates to obtain a set of rationality requirements for foundational non-logically omniscient belief change. Our contraction operator for this revised framework uses the extra-logical information provided by the partial ordering to make a rational choice amongst the possible outcomes. Since the information is assumed to be partial, when multiple outcomes arise the operator takes a conservative approach by choosing only those beliefs are common to all possible outcomes. As explained below and unlike some proposals based on closure, this revised approach does not permit the loss of all original information.

We explain how the new operator satisfies our reformulated set of rationality requirements. Furthermore, we show that a similarly defined operator based on the logically omniscient framework of deductively closed sets (which takes a less conservative approach by sanctioning the disjunction of all the possible outcomes) is shown to satisfy all but the *recovery posulate* of the original AGM set.

As per our more general motivation mentioned above, we have found it convenient to express many of the concepts in the terminology of the literature on belief revision in AI. Within the current scope of this paper, we believe that the correspondence between belief sets and deductive databases and between belief change and database updates should be straightforward.

2 Preliminaries

2.1 A Taxonomy of Belief Change

To analyze the complex process of revision in simpler terms, we adopt the classification from [1]. In this scheme, belief change may be of three types:

Expansion: the addition of a new belief A, which is consistent with the existing belief set K; it is denoted by K^{+}_{A}.

Contraction: the retraction of some currently held belief A from the set of existing beliefs K, and is denoted by K^{-}_{A}.

Revision: the addition of a new belief A that is possibly inconsistent with the current set of beliefs K, and is denoted by K^{*}_{A}.

Expansion is the simplest kind of belief change, requiring the trivial addition of a belief to a set of beliefs. Contraction is non-trivial—in general, there may be several possible outcomes of the contraction operation and a rational choice of what to retain in the resulting belief set is required. Revision is a combination of the operations of contraction and expansion: the existing set of beliefs is first contracted with the negation of the new belief, and the result of this is expanded with the new belief, as given by the so-called *Levi identity* shown below:

$$\mathrm{K}^{*}_{A} = (\mathrm{K}^{-}_{\neg A})^{+}_{A}.$$

Note that this assumes a consistent result.

Since revision can be viewed as the composition of the contraction and expansion operations, and since expansion is trivial, we can study contraction as the fundamental operation in belief change.

2.2 Rationality Postulates for Belief Revision

To discuss *rational* belief change, it is important to first define what this "rationality" entails. The postulates proposed by Alchourròn, Gärdenfors and Makinson [1] (the so-called AGM postulates), provide a set of criteria that captures our commonly held intuitions of what rational belief change should be. To substantiate our claim that our belief change operator is rational, we shall show that our operator satisfies these postulates. Different sets of AGM postulates exist for expansion, contraction, and revision. Taking the view that contraction is the most fundamental belief change operation, we shall show that our contraction operators satisfy the AGM postulates for contraction.

Generally speaking, belief revision requires a specification for the inputs and the outputs of a belief revision operator, together with some criteria for the operator's application. The AGM postulates provide criteria for the output of a belief revision process that takes one new proposition and a deductively closed set as input. The representation language is propositional logic; the deductively closed sets denoting beliefs are referred to as *knowledge sets*. In the postulates which follow, K represents the deductively closed set of beliefs currently held, while A and B represent beliefs which are retracted from K. The contraction operation is denoted by K_A^-. Our numbering scheme "n-" refers to the nth postulate for contraction.

(1-) K_A^- is a knowledge set.

(2-) $\mathrm{K}_A^- \subseteq \mathrm{K}$.

(3-) If $\mathrm{K} \not\models A$ then $\mathrm{K}_A^- = \mathrm{K}$.

(4-) If $\not\models A$ then $A \notin \mathrm{K}_A^-$.

(5-) If $A \in \mathrm{K}$ then $\mathrm{K} \subseteq (\mathrm{K}_A^-)_A^+$.

(6-) If $\models A \leftrightarrow B$ then $\mathrm{K}_A^- = \mathrm{K}_B^-$.

(7-) $\mathrm{K}_A^- \cap \mathrm{K}_B^{=} \subseteq \mathrm{K}_{A\wedge B}^-$.

(8-) If $A \notin \mathrm{K}_{A\wedge B}^-$ then $\mathrm{K}_{A\wedge B}^- \subseteq \mathrm{K}_A^-$.

Postulate one (1-) requires the result of contraction to be a consistent deductively closed set of beliefs. Number two (2-) requires that contraction should not result in any new beliefs. The third (3-) says that contracting something that is not already believed has no effect on our beliefs. The fourth postulate (4-) says that unless A is logically valid, contraction is always successful. Five (5-) requires that,when a belief is retracted and then added again, we should be able to recover our original beliefs. Postulate six (6-) requires that if two beliefs are logically equivalent, then contraction of the same set of beliefs with either of them is the same. The seventh (7-) requires that the retraction of a conjunction of beliefs should not retire any beliefs that are common to the retraction of the same belief set with each individual conjunct. The last posulate, eight (8-), requires that, when retracting the conjunct of two beliefs A and B forces us to give up A, then in retracting A, we do not give up any more than in retracting the conjunction of A and B.

2.3 Coherence versus foundational belief revision

It is popular to distinguish foundation from coherence approaches to belief revision, although as hinted above, this distinction may be more an artifact of the method of specification (e.g., on proof theory). Still the classification is common, so we review the distinction here, as a basis for further reference below. The distinction hinges on which components of a belief set can be discarded during the revision operation. The *coherence theory of belief revision* requires that the results revision arises by making minimal changes to the original set. The intuition is one common to theories of knowledge growth in science (e.g., [?]), where "coherence" is acheived by somehow making minimal changes to an existing ist of beliefs. The justification of an individual belief amongst those in a coherent set of beliefs is not its provability w.r.t. to a set of self-evident axioms, but on the extent to which it coheres with all other beliefs.

The *foundational theory of belief revision* requires that every belief be self-evident or have a non-circular, finite sequence of justifications grounded in a set of self-evident beliefs. Under this approach, belief revision involves removing those beliefs that have no satisfactory justification and adding those beliefs that are either self-evident or are justified by a set of self-evident beliefs.

Example 2.3.1: Consider a belief set consisting of the propositions (i) "valve A works," and (ii) "if valve A works, oil flows in the pipeline." A natural consequent of the latter two beliefs is (iii) "oil flows in the pipeline." When propositions (i)-(iii) are revised with the new proposition (iv) "valve A does not work," the coherence theory requires that minimal change be made to the entire corpus of belief. The realization of minimal change is frequently interpreted syntactically, so that the minimum number of propositions is discarded. This would suggest that both (i) and (ii) be retained. However, under the foundational view, (iii) would be viewed as non-self-evident, as its support, proposition (i) is called into question by the new proposition (iv). □

While the above example demonstrates the intended difference between the coherence and foundation approaches, it doesn't make the difference precise. We speculate that semantically defined revision criteria (e.g., select the theory which, after revision, has the minimal model) could likely show that some coherent and foundational theories coincide. Rather than taking up this pursuit here, we continue with our current goal of surveying the current frameworks for AGM-style revision posulates.

The approach taken by Gärdenfors, Alchourron and Makinson [6] [1] in their representation framework of deductively closed belief sets, in their definition of the rationality postulates and in their constructions of the contraction operator is claimed to be based on the coherence view. As Gärdenfors admits [5], the issue of maintaining the connections between premises and the consequences they support (i.e., *reason maintenance*) is totally ignored. However, as the previous example shows, reason maintenance is crucial for maintaining common sense rationality. Furthermore, even if a coherence-based viewed approaches rationality, examining the entire body of beliefs upon each revision to ensure coherence is not computationally viable—the closures of belief sets are typically infinite. It makes more sense to change a finite set of self-evident beliefs from which all other beliefs follow. With these considerations in mind, and even in anticipation of a more general semantically-motivated revision theory, we here restrict out attention to foundational belief revision.

3 Rationality and Logical Omniscience

We have seen in the previous section how the AGM framework of deductively closed sets of beliefs is unintuitive for most real-life applications of belief revision. However, the fact that belief change is not foundational in this framework is not the only problem with it. This framework for belief representation requires *logical omniscience* - the rational agent must believe in all the logical consequences of the beliefs explicitly held or represented. Most real-life rational agents are not logically omniscient - the closure of their beliefs is limited to a proper subset of the set of full logical consequences of their beliefs. How small this subset should be is an open question and one which we shall not try to address. We shall, however, show some of the problems with this logically omniscient framework and shall propose a solution that overcomes some of these.

Example 3.1: Consider the belief set $\{a\}$. The deductive closure of this belief set will contain the belief $\neg c \vee a$. Written differently, this is $c \rightarrow a$. A logically omniscient agent will thus conclude $c \rightarrow a$ on the basis of belief in a. In most real-life situations, this conclusion is neither rational or warranted.□

Consider the following example from [14]:

Example 3.2: Consider the belief set $\{a, b\}$. Let Cn(X) denote the deductive closure of a belief set X. Then Cn($\{a, b\}$)=$\{a, b, a \vee b, a \wedge b, a \rightarrow b, b \rightarrow a, \ldots\}$. When retracting a from this belief set both Cn($\{\neg a, b\}$) and Cn($\{\neg a, \neg b\}$) are maximal consistent subsets which do not imply a. These maximal subsets of the existing belief set thus represent possible outcomes of contraction with minimal change. Intuitively, however, if we believe in both a and b, and wish to retract our belief in a there is no reason to stop believing b. In some sense the set Cn($\{\neg a, \neg b\}$) is not a rational outcome of belief change.□

We are thus led to believe that rational belief change does not require logical omniscience, or the still stronger statement that a rational agent must necessarily be a limited reasoner as opposed to a logically omniscient one. Deductively closed sets are by definition infinite and computationally unrealizable, which is another reason why they should not be used as the framework for belief representation. A suitable belief closure should thus be finitely representable and should include only those beliefs which our view of rationality sanctions. We define the *partial closure* of a set, as distinct from full deductive closure to address some of these problems. The idea of partial closure is not new - it can be traced back to Quine's *prime implicates* . However, we base our definition on the one given in [14].

We recall some definitons first. A *clause* is a set of literals $\{\neg B_1, \ldots, \neg B_n, A_1, \ldots, A_m\}$ which may also be written as:

$$B_1 \wedge \ldots \wedge B_n \rightarrow A_1 \vee \ldots A_m$$

where $m + n \geq 1$. A *theory* is a set of clauses.

Definition 3.1 : The *partial closure* of a theory T, denoted by $T^\star$, is defined as $T^\star = T \cup \{\mu \mid T \vdash \mu$ and $T \not\vdash \mu'$ for any $\mu' \subset \mu$ and there exists no ψ such that both ψ and $\neg\psi$ are in $\mu\}$. Here, $\vdash$ denotes full clausal resolution.□

Example 3.3: $\{a, a \rightarrow b, c, c \rightarrow d\}^{\star}= \{a, a \rightarrow b, b, c, c \rightarrow d, d\}$. □

The partial closure of a theory thus represents the set of minimally derivable clauses which do not contain any tautologies. Thus given a theory $\{a, a \rightarrow b\}$, its partial closure will contain b but not clauses like $c \rightarrow a$ or $d \rightarrow b$ which would have been contained in the full deductive closure of this theory. The partial closure framework cannot make any claim to being a comprehensive solution to the problem of logical omniscience (see [10] for a good discussion of various approaches to addressing the logical omniscience problem in logics of knowledge and belief). By restricting belief representation to the clausal form, by taking derivability under clausal resolution instead of logical implication and by eliminating tautologies, this framework avoids some of the problems of logical omniscience.

In Section 5, we shall define a new version of the AGM postulates for contraction based on the framework of partial closure. But first, we define our contraction operator in the next section.

4 Contracting Epistemically Stratified Belief Sets

A very common situation that occurs during belief revision is that there are a number of alternative belief sets which can potentially constitute the revised belief set. If one were to take a disjunction of these mutually inconsistent belief sets, as in [12], or their intersection, as in [6], one would have to give up most of one's previously held beliefs, with the resulting belief set containing only the consequences of the current epistemic input.

Example 4.1: Consider a belief set represented by a set of formulae from classical propositional logic: $\{a, a \rightarrow b\}$. As a result of an epistemic input $\neg b$, there can be two possible outcomes of the revision operation: $\{a, \neg b\}$ and $\{a \rightarrow b, \neg b\}$. The disjunction, or equivalently, the intersection of these two sets contains only $\neg b$. □

The above example motivates the need for taking into account some extra-logical factors in order to make a rational choice between the possible outcomes of the revision operation. *Epistemic entrenchment* is one such extra-logical factor that has been considered by Gärdenfors and Makinson in [6] and [7], where beliefs that are more epistemically entrenched have greater utility for the purpose of inquiry and decision-making and vice versa. While epistemic entrenchment is delinked from measures of certainty or probability, and is motivated and justified by a complex set of philosophical arguments, we shall give these issues a wide berth and shall consider a much simpler model. In our model, we shall assume that *epistemic priorities* can be assigned to beliefs, without making any commitment to what such a priority assignment should be based on. Bases such as those for epistemic entrenchment, or degress of certainty, or probability are all acceptable for the purpose of assigning epistemic priorities. We then define a framework for belief sets in which a possibly incomplete specification of the relative epistemic priorities amongst beliefs exists.

Definition 4.1: An *epistemically prioritized* belief set is one in which an well-founded partial pre-order $\leq$ amongst the beliefs is specified such that $\alpha \leq \beta$ if and only if β has an epistemic priority that is at least high as that of α. If $\alpha \leq \beta$ and $\beta \not\leq \alpha$ then $\alpha < \beta$. □

Like Gärdenfors and Makinson, our concern here is to utilize an extra-logical ordering on the beliefs to make a rational choice of a revised belief set, given several possible alternatives. We differ in two crucial respects. First, we choose the foundational approach to belief revision. Secondly, we do not make the assumption of total connectivity as in [6] and [7], but permit situations where beliefs are incomparable under the ordering $\geq$. We then define a contraction operator which the extra-logical knowledge provided by the ordering to make a rational choice among the possible outcomes of the contraction operation.

First, we need a framework for foundational belief change. A *belief base*, defined as a finite set of sentences of classical propositional logic, provides such a framework. Finite belief bases have been used by Nebel in [12], but he too considers a total ordering on the propositions of the belief base.

Definition 4.3: A *prioritized belief base* is a finite set of formalae of classical propositional logic for which a partial epistemic prioritization $\leq$ exists. □

Definition 4.4: A *belief set* corresponding to a prioritized belief base Δ is given by $\Delta^\star$. □

An informal description of our contraction operator is as follows:

- For contracting a sentence A from a prioritized belief base Δ, we first identify the maximal consistent subsets of Δ that do not imply A. There can, in general, be more than one such subset. $C_{max}(\Delta, A)$ is the set of all such subsets.
- The operator E_{max} chooses from amongst the subsets contained in $C_{max}(\Delta, A)$ those that contain sentences of higher epistemic priority. Thus there are no subsets in $C_{max}(\Delta, A)$ which dominate, in terms of epistemic content, those in $E_{max}(C_{max}(\Delta, A))$.
- The partial closure of the intersection of the partial closures of all the subsets in $E_{max}(C_{max}(\Delta, A))$ is taken to be the result of the contraction operation. The intuition here is that given a set of possible views of reality described by the partial closures of the elements of $E_{max}(C_{max}(\Delta, A))$, we choose to be conservative and believe only those sentences that are common to all the views.

Definition 4.5: Given a partial pre-order $\leq$, a set of propositional sentences X is said to *dominate* another such set Y if there exists some sentence $x \in (X - Y)$, where $-$ denotes classical set difference, and there exists some sentence $y \in (Y - X)$ such that $x > y$ and it is not true that there exists some sentence $p \in (Y - X)$ and some sentence $q \in (X - Y)$ such that $p > q$. □

The intuition in this definition of dominance is that if some set of sentences dominates another set of sentences, then the former contains beliefs that are, in some obviously comparable way, higher in epistemic priority than those in the latter.

Definition 4.6: The set $C_{max}(\Delta, A)$ of maximal subsets of the irredundant belief base Δ that are consistent with $\neg$A, where A is a propositional sentence, is defined as follows:
$C_{max}(\Delta, A)=\{S \mid S \subseteq \Delta \text{ and } S \not\models A \text{ and for any } S' \text{such that } S \subset S' \subseteq \Delta,\ S' \models A\}$. □

Definition 4.7: The operator E_{max} is defined to select dominant sets of beliefs from a set of such sets. The intuition is that this function is used to select, from the set of those maximal consistent subsets of the belief base which do not imply the belief being contracted, those subsets that retain beliefs of higher epistemic priority. Formally:
$E_{max}(X)=\{S \mid S \in X \text{ and there exists no } S' \text{ such that } S' \text{ dominates } S\}$.□

Definition 4.8: Given a prioritized belief base Δ, a belief set K given by K$=\Delta^\star$ and an epistemic prioritization $\leq$, the contraction of K with some sentence A is given by:

$$K^{\sim}_{A} = (\bigcap(E_{max}(C_{max}(\Delta, A)))^\star)^\star$$

where $\bigcap$X denotes the intersection of all the elements of X, given that X is a set of sets.□

Example 4.3: Consider the prioritized belief base :

$$\Delta = \{a, a \rightarrow b, c, c \rightarrow b\}.$$

We wish to retract the belief b from this belief base. The set maximal subsets of Δ which are consistent with $\neg b$ is given by:

$$\mathrm{C}_{max}(\Delta,\, b)=\{\{a, c\}, \{a, c \rightarrow b\}, \{c, a \rightarrow b\}, \{a \rightarrow b, c \rightarrow b\}\}.$$

Let $a > c > a \rightarrow b$ be the only $>$-ordering relations derivable from the given well-founded partial pre-order $\leq$. Then it is easy to see from the definition of dominance that $\{a, c\}$ and $\{a, c \rightarrow b\}$ dominate all the other elements of $\mathrm{C}_{max}(\Delta,\, b)$, and that they are mutually incomparable with respect to dominance. Thus they are the only epistemically maximal subsets in $\mathrm{C}_{max}(\Delta,\, b)$.

$$\mathrm{E}_{max}(\mathrm{C}_{max}(\Delta,\, b)=\{\{a, c\}, \{a, c \rightarrow b\}\}.$$

The result of the contraction operation is given by the intersection of the partial closures of these two sets.

$$\mathrm{K}_b^{\sim} = \{a, c\}^{\star} \cap \{a, c \rightarrow b\}^{\star} =(\{a\})^{\star}.□$$

This construction of the contraction operator represents a skeptical approach - we only choose to believe in whatever is sanctioned by all possible views of the world that might result when a given belief is retracted, where a view of the world is given by the partial closure of a set of clauses. As we shall see later, this construction corresponds fairly well to our conception of how a non-omniscient, limited reasoner should perform belief change. A series of belief change operations can be viewed as one finite belief base yielding another, while at each step, the partial closure of the current belief base represents the agent's view of the world.

We can also define a somewhat different construction for foundational contraction by logically omniscient agents which are somewhat less skeptical when faced with multiple possible views of the world as a result of belief change by being willing to admit any of those possible views of the world. These agents thus take the disjunction of the different possible views. In our representation framework of partially closed theories, disjunction is defined as the operation ED, or *extended disjunction*, which we define below. We base this definition on a similar one given in [14].

Definition 4.9: Let T_1 and T_2 be any two theories and let $\Pi = \{T_1, \ldots, T_n\}$ be a set of theories.

- $ed(T_1, T_2) = \{T_1 \cap T_2\} \cup \{\alpha_i \vee \alpha_j \mid \alpha_i \in (T_1 - T_2) \text{ and } \alpha_j \in (T_2 - T_1)\}$.
- $\mathrm{ED}(\Pi) = ed(\Pi)$ if cardinality of Π is 2.
  $= ed(T_1,\mathrm{ED}(\Pi - T_1))$ if cardinality of Π is greater than 2.□

Definition 4.10: The foundational contraction operator $-$ for logically omniscient agents is defined as follows:

$K_A^- = Cn(ED(E_{max}(C_{max}(\Delta, A))))$.

where Δ is a prioritized belief base, K is the corresponding belief set given by

K=Cn(Δ), K_A^- denotes the contraction of the belief A from K and Cn(X) is the set of all possible logical consequences of formulae contained in X. □

Example 4.4: Consider the irredundant belief base B given by:

$$\Delta = \{a, a \to b, c, d\}$$

Let $a > a \to b$ be the only $>$-ordering relation deducible from the given well-founded partial pre-order $\leq$.

$C_{max}(\Delta, ((c \wedge d) \vee b)) = \{\{a, c\}, \{a, d\}, \{a \to b, c\}, \{a \to b, d\}\}$.
$E_{max}(C_{max}(\Delta, ((c \wedge d) \vee b))) = \{\{a, c\}, \{a, d\}\}$.
$K_C^- n(c \wedge d) \vee b = \{a, c \vee d\}$.□

We shall see later that this contraction operator satisfies the rationality requirement as defined in the logically omniscient framework of the original AGM postulates.

5 A Revised Set of Rationality Postulates

The AGM postulates define the requirements for rational belief change in logically omniscient agents, as we have seen in Section 3. We also saw how the assumption of logical omniscience can provide unintuitive results, and how the partial closure of an agent's beliefs avoids some of these unintuitive results. In Section 4, we defined the framework of irredundant finite belief bases for foundational belief change. We have thus provided the groundwork for defining the requirements of foundational, non-omniscient, rational belief change.

But first we need to discuss the relevance of AGM postulate (5-), the so- called *recovery postulate*, as a rationality requirement. As the following example shows, neither of our contraction operators $\sim$ and $-$ satisfy postulate (5-).

Example 5.1: Consider the irredundant finite belief base Δ given below:

$$\Delta = \{a, a \to b\}$$

Let $a \to b > a$ be the only $>$-ordering relation derivable from $\leq$.In the case of the operator $\sim$ we take the belief set K= $\Delta^\star$ and in the case of the operator $-$ we take K=Cn(Δ). Then:

$K_b^\sim = \{a \to b\}^\star$.
$K_b^- = Cn\{a \to b\}$.

It is easy to see that postulate (5-) is violated by both contraction operators although the belief change here is obviously rational.□

We therefore conclude that in the presence of an epistemic stratification of beliefs, the recovery postulate is not a requirement for rational contraction. Since an epistemically uniform belief set is a special case of an epistemically stratified one, recovery, in general, is not a rational requirement at all.

We are now in a position to provide a framework and a set of rationality requirements for foundational, non-omniscient belief change. Beliefs are represented by a prioritized finite belief base Δ of propositional sentences in clausal form. Belief closure is given by the partial closure of a theory, so that K= $\Delta^{\star}$. Individual beliefs A and B are represented as clauses. We now reformulate the AGM postulates for contraction in this framework and drop the recovery postulate totally. The revised set of postulates are:

(1∼) $K_A^{\sim}$ is a partially closed theory.

(2∼) $K_A^{\sim} \subseteq K^{\star}$.

(3∼) If $K \not\models A$ then $K_A^{\sim} = K^{\star}$.

(4∼) If $\not\models A$ then $K_A^{\sim} \not\models A$.

(5∼) If $\models A \leftrightarrow B$ then $K_A^{\sim} = K_B^{\sim}$.

(6∼) $K_A^{\sim} \cap K_B^{\sim} \subseteq K_{A\wedge B}^{\sim}$.

(7∼) If $K_{A\wedge B}^{\sim} \not\models A$ then $K_{A\wedge B}^{\sim} \subseteq K_A^{\sim}$.

The changes made to the original postulates are fairly obvious, given the framework of partially closed theories. Postulate (1∼) reflects our commitment to the view that rationality necessarily requires logical non-omniscience. Also, inclusion of a belief set in a deductively closed belief set has been changed to implication of a belief by a partially closed belief set in postulates (4∼) and (7∼).

6 Rationality Results for ∼ and −

The main results of this section are that ∼ satisfies the postulates (1∼)-(7∼) while − satisfies (1-)-(4-) and (6-)-(8-). The following two lemmas are useful in these proofs.

Lemma 1: $C_{max}(\Delta, A\wedge B) = C_{max}(\Delta, A) \cup C_{max}(\Delta K, B)$.

Lemma 2: For any set of theories $\Pi = \{T_1, \ldots, T_n\}$ and for any α, $\Pi \models \alpha$ iff each $i = 1 \ldots n$ $T_i \models \alpha$.

Theorem 1: The contraction operator ∼ satisfies postulates (1∼)-(7∼).

Proof : For postulates (1∼)-(5∼), the proof is trivial. The proof for postulates (6∼) and (7∼) is given in the appendix. □

Theorem 2: The contraction operator − satisfies postulates (1-)-(4-) and (6-)-(8-).

Proof : The proof for postulates (1-)-(4-) and (6-) is trivial. Using Lemma 2, the proof for postulate (7-) closely follows that of (6∼) and that for (8-) closely follows the proof for (7∼). These are not presented in this paper for lack of space.□

7 Related Work

Our approach differs from the work of Alchourrón, Gärdenfors and Makinson [1], [6], [7], [5] in three crucial ways. First, they assume the ordering relation among the sentences of the infinite, deductively closed belief set satisfy the following 5 requirements:

E1 If $\alpha \leq \beta$ and $\beta \leq \gamma$ then $\alpha \leq \gamma$.(Transitivity)

E2 If $\alpha \vdash \beta$ then $\alpha \leq \beta$.(Dominance)

E3 For any α and β, $\alpha \leq \alpha \wedge \beta$ or $\beta \leq \alpha \wedge \beta$.(Conjunctiveness)

E4 If Δ is consistent, $\alpha \notin \mathrm{Cn}(\Delta)$ iff $\alpha \leq \beta$ for all β.(Minimality)

E5 $\beta \leq \alpha$ for all β, only if $\vdash \alpha$.(Maximality)

These requirements can turn out to be too restrictive in general. In our approach, we only require a well-founded partial order, which it more generally applicable. Specifically, E1 - E3 imply that the ordering is a *total* one. Pragmatically, a total ordering on an infinite deductively closed set of sentences is impossible to specify. Secondly, we take the foundational approach to belief change while their approach is coherentist. Thirdly, we define rational belief change for non-omniscient, limited reasoners while they require their rational agents to be logically omniscient.

Nebel [12] considers foundational contraction of finite *belief bases* which are finite sets of propositions considered to represent the set of "basic beliefs". However, he too considers a total ordering on the belief sentences. Also, the definition of *dominance* used may leave certain sets of beliefs incomparable, when they are actually intuitively comparable and are comparable under our definition.

Fuhrmann [4] and Nayak [11] have proposed a foundational contraction operator. When revising with a propositional sentence A, they denote $E(A)$ to be the set of minimal subsets of the belief base which entail A. Given a partial order $\leq$ on the beliefs in the belief base, they denote $R(A)$ to be the set of $\leq$-minimal beliefs in each of the subsets contained in $E(A)$. For a belief base K, contraction is then defined as $Cn(K - R(A))$ where $-$ is taken as the set difference operator. While this contraction operator corresponds to our intuition in most cases, there are situations in which more beliefs are given up than is warranted, as the following example shows.

Example 7.3: Consider the finite belief base given by
$\{a, a \rightarrow b, c, c \rightarrow b\}$
with the following ordering relations:
$a > c \rightarrow b$
$c > a \rightarrow b$
Then $E(b) = \{\{a, a \rightarrow b\}, \{c, c \rightarrow b\}\}$.
$R(b) = \{a, a \rightarrow b, c, c \rightarrow b\}$. We thus have the unintuitive result in which the entire existing belief base is removed. It is clear that the belief base $\{a, c\}$ would be an intuitive result of the contraction. This is precisely the result that our contraction operator provides. The reason Fuhrmann and Nayak's operator give up too much is because no ordering relation exists between the elements of the subsets contained in $E(b)$. The full power of the extra-logical information provided by the ordering $\leq$ has not been brought to bear during the process of contraction.$\square$

Fagin, Ullman and Vardi have proposed update semantics for prioritized databases in [3]. In their framework, the database priorities were represented by numerical tags attached to sentences in the database. However, they too consider total orderings. Willard and Yuan [14] have proposed an update semantics for deductive databases in which rules have higher priority over atomic facts. However, it is obvious that such a prioritization does not hold true in general. Katsuno and Mendelzon have analyzed belief revision in terms of orderings on models [8], but have provided no explicit construction of a belief revision operator. Also, in real-life, orderings are specified by people on the syntactic form of beliefs, and not on models. Rao and Foo [13] have analyzed both the foundational and coherentist approaches to belief revision, but their approach assumes a modal logic with auto-epistemic operators.

8 Appendix

Note : For the puropose of brevity in the proofs that follow, we shall use the term "p-closure" to denote "partial closure".

Lemma 1: $C_{max}(\Delta, A \wedge B) = C_{max}(\Delta, A) \cup C_{max}(\Delta K, B)$.

Proof: By definition, the elements of $C_{max}(\Delta, A \wedge B)$ are maximal consistent subsets of Δ which do not entail $A \wedge B$. Hence, any element of $C_{max}(\Delta, A \wedge B)$ does not entail A or does not entail B but never both. The proof is then trivial.□

Theorem : $K^{\sim}_{A} \cap K^{\sim}_{B} \subseteq K^{\sim}_{A \wedge B}$.

Proof : Assume that there exists $d \in (K^{\sim}_{A} \cap K^{\sim}_{B})$ such that $d \notin K^{\sim}_{A \wedge B}$. We will show that this is not possible.

The possible ways in which some $d \in K^{+}$ can be absent in $K^{\sim}_{A \wedge B}$ are:

Case 1: d is not in the p-closure of any element of $C_{max}(\Delta, A \wedge B)$).

Case 2: d is in the p-closure of some element of $C_{max}(\Delta, A \wedge B)$ but not of any element of $E_{max}(C_{max}(\Delta, A \wedge B))$.

Case 3: d is in the p-closure of some but not all elements of $E_{max}(C_{max}(\Delta, A \wedge B))$.

We shall now analyze each case:

Case 1: $d \models A \wedge B$. Hence $d \notin (K^{\sim}_{A} \cap K^{\sim}_{B})$.

Case 2: Let X_d denote an element of $C_{max}(\Delta, A \wedge B)$ which includes d in its p-closure. In this case there exists some $Y \in C_{max}(\Delta, A \wedge B)$ such that for all X_d, Y dominates X_d. By Lemma 1, $Y \in C_{max}(\Delta, A)$ or $Y \in C_{max}(\Delta, B)$. Assume that $Y \in C_{max}(\Delta, A)$. Two situations are possible:

- There is no $X_d \in C_{max}(\Delta, A)$. Then $d \notin (K^{\sim}_{A} \cap K^{\sim}_{B})$.
- There exists at least one $X_d \in C_{max}(\Delta, A)$. Then Y will dominate X_d, so X_d will not be in $E_{max}(C_{max}(\Delta, A))$. Hence $d \notin (K^{\sim}_{A} \cap K^{\sim}_{B})$.

Case 3: There must exist some $Y \in E_{max}(C_{max}(\Delta, A \wedge B))$ which does not contain d in its p-closure. By Lemma 1, Y must be in $C_{max}(\Delta, A)$ or $C_{max}(\Delta, B)$. We assume $Y \in C_{max}(\Delta, A)$. We must now consider two possible situations:

- $Y \notin E_{max}(C_{max}(\Delta, A))$. Then there must exist some $Z \in C_{max}(\Delta, A)$ such that Z dominates Y. By Lemma 1, $Z \in C_{max}(\Delta, A \wedge B)$. But since Z dominates Y, $Y \notin E_{max}(C_{max}(\Delta, A \wedge B))$. Hence, this situation is impossible.
- $Y \in E_{max}(C_{max}(\Delta, A))$. Since d is not in the p-closure of Y, and since $K^{\sim}_A$ is the intersection of the p-closures of all the members of $E_{max}(C_{max}(\Delta, A))$, $d \notin K^{\sim}_A$. Hence, $d \notin (K^{\sim}_A \cap K^{\sim}_B)$.

Hence, it is not possible for some $d \in (K^{\sim}_A \cap K^{\sim}_B)$ and $d \notin K^{\sim}_{A \wedge B}$. □

Theorem : If $K^{\sim}_{A \wedge B} \not\models A$ then $K^{\sim}_{A \wedge B} \subseteq K^{\sim}_A$.

Proof : Assume there exists a $d \in K^{\sim}_{A \wedge B}$ such that $d \notin K^{\sim}_A$. We will show that this is impossible.

A clause d may not be in $K^{\sim}_A$ in the following three ways:

Case 1: d is not in the p-closure of any element of $C_{max}(\Delta, A)$.

Case 2: d is in the p-closure of some element of $C_{max}(\Delta, A)$ but not in the p-closure of any element of $E_{max}(C_{max}(\Delta, A))$.

Case 3: d is in the p-closure of some but not all elements of $E_{max}(C_{max}(\Delta, A))$.

We shall analyze each of these cases in turn:

Case 1: d must be in all elements of $E_{max}(C_{max}(\Delta, A))$ and these must all be elements of $C_{max}(\Delta, B)$ which are not in $C_{max}(\Delta, A)$, by Lemma 1. It is not possible for $E_{max}(C_{max}(\Delta, A \wedge B))$ to be a singleton because by definition of maximality, every element of $C_{max}(\Delta, B)$ that is not in $C_{max}(\Delta, A)$ must imply A, so that $K^{\sim}_{A \wedge B} \models A$. Let us assume the simplest case when $E_{max}(C_{max}(\Delta, A \wedge B))$ contains exactly two elements, both of which are in $C_{max}(\Delta, B)$ but not in $C_{max}(\Delta, A)$. Let $T = \{C_1, \ldots, C_n\}$ be any subset of the p-closure of the original belief base Δ which implies A. In general, there can be more than one such subset. It is required that the intersection of the p-closures of the two elements of $E_{max}(C_{max}(\Delta, A \wedge B))$ not contain T. Let us assume that one of these two does not contain T but contains a proper subset. By definition of maximality, this set must then be in $C_{max}(\Delta, A)$. But this violates the requirement that each element of $E_{max}(C_{max}(\Delta, A \wedge B))$ must be in $C_{max}(\Delta, B)$ but not in $C_{max}(\Delta, A)$.

Case 2: Let X_d denote those elements of $C_{max}(\Delta, A)$ which contain d. Now if $E_{max}(C_{max}(\Delta, A \wedge B))$ contains no element of $C_{max}(\Delta, A)$, then the proof is the same as in Case 1. Otherwise, if there is an element of $C_{max}(\Delta, A)$ in $E_{max}(C_{max}(\Delta, A \wedge B))$, none of these will be an X_d since all X_d in $C_{max}(\Delta, A)$ are dominated by other elements of $C_{max}(\Delta, A)$. Thus there will be at least one element of $E_{max}(C_{max}(\Delta, A \wedge B))$ which will not contain d in its p-closure. So $d \notin K^{\sim}_{A \wedge B}$.

Case 3: Again, if $E_{max}(C_{max}(\Delta, A \wedge B))$ contains no element of $C_{max}(\Delta, A)$, then the proof is the same as in Case 1 above. If $E_{max}(C_{max}(\Delta, A \wedge B))$ contains even one element of $C_{max}(\Delta K, A)$, then it must contain all elements of $E_{max}(C_{max}(\Delta, A))$. Then there will be a least one element of $E_{max}(C_{max}(\Delta, A \wedge B))$ which will not contain d in its p-closure. Hence, $d \notin K^{\sim}_{A \wedge B}$. □

References

[1] Carlos E. Alchourrón, Peter Gärdenfors and David Makinson. On the logic logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic,* 50:510-530, 1985.

[2] Mukesh Dalal. Investigations into a theory of knowledge base revision : Preliminary Report. In *Proc. of the 7th National(USA) Conf. on Artificial Intelligence,* pages 475-479, 1988.

[3] Ronald Fagin, Jeffrey D. Ullman and Moshe Y. Vardi. On the semantics of updates in databases. In *Proc. of the Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems,* pages 352-365, 1983.

[4] André Fuhrmann. *Relevant logics, modal logics and theory change*, Ph.D. Dissertation, Department of Philosophy and Automated Reasoning Project, Australian National University, 1988.

[5] Peter Gärdenfors. *Knowledge in flux - modelling the dynamics of epistemic states*, Bradford/M.I.T Press, Cambridge (Mass.), 1988.

[6] Peter Gärdenfors. Epistemic Importance and Minimal Changes of Belief.*Australasian Journal of Philosophy*, 62:137-157, 1984.

[7] Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proc. of the Second Conference on Theoretical Aspects of Reasoning about Knowledge,* pages 83-95, 1988.

[8] Hirofumi Katsuno and Alberto O. Mendelzon. A unified view of propositional knowledge base updates. In *Proc. of the 11th International Joint Conference on Artificial Intelligence*, pages 1413-1419, 1989.

[9] Arthur M. Keller and Marianne Winslett Wilkins. On the use of an extended relational model to handle changing incomplete information. *IEEE Trans. on Software Engineering,* SE-11:7, pages 620-633, July 1985.

[10] Gregory L. McArthur. Reasoning about knowledge and belief: a survey. *Computational Intelligence*, 4:223-243, 1988.

[11] Abhaya Nayak. Foundational Belief Change - I. Forthcoming.

[12] Bernhard Nebel. A knowledge level analysis of belief revision. In *Proc. of the First International Conference on Principles of Knowledge Representation and Reasoning*, pages 301-311, 1989.

[13] Anand S. Rao and Norman Y. Foo. Formal theories of belief revision. In *Proc. of the First International Conference on Principles of Knowledge Representation and Reasoning,* 1989.

[14] Leigh Willard and Li Yan Yuan. The revised Gärdenfors postulates and update semantics. In *Proc. of ICDT '90*, pages 409-421, 1990.

[15] Marianne Winslett. Reasoning about action using a possible models approach.In *Proc. of the 7th National(USA) Conf. on Artificial Intelligence,* pages 89-93, 1988.