UNIVERSITY OF ALBERTA

Logical Aspects of Belief Change

BY

Aditya Kumar Ghose

A thesis submitted to the Faculty of Graduate Studies and Research in partial fullfillment of the requirements for the degree of Master of Science.

DEPARTMENT OF COMPUTING SCIENCE

Edmonton, Alberta
Fall 1991

ISBN 0-315-70039-4

UNIVERSITY OF ALBERTA

RELEASE FORM

NAME OF AUTHOR: Aditya Kumar Ghose

TITLE OF THESIS: Logical Aspects of Belief Change

DEGREE: Master of Science

YEAR THIS DEGREE GRANTED: 1991

Aditya Kumar Ghose

c/o Prof. A. K. Ghose

Indian School of Mines

Dhanbad, 826004

India

October 11, 1991

UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recomended to the Faculty of Graduate Studies and Research for acceptance a thesis entitled Logical Aspects of Belief Change submitted by Aditya Kumar Ghose in partial fulfillment of the requirements for the degree of Master of Science.

Randy Goebel

Li Yan Yuan

Jia You

Bernard Linsky

To the memory of
Prof. Subodh Kumar Ghose
and
Prof. Sisirkumar Ghose

Abstract

Belief change, the process by which a rational agent makes the transition from one belief state to another, appears to be the basis for most intelligent activity. Most existing studies of belief change have assumed a framework of infinite deductively closed sets of beliefs over which a total ordering has been specified. This approach suffers from several problems: besides the well-known problems of logical omniscience, the framework is computationally unrealizable and reason maintenance is ignored. To address some of these problems, we define a belief change framework that is foundational, logically non-omniscient, and which allows a partial specification of the relative epistemic priorities amongst beliefs. Two belief change operators are studied in the context of this framework. When the requirement that every belief state be consistent is relaxed, we obtain a new framework which allows us to look at reasoning with default logic as a process of belief change. Such an approach appears to have several advantages: some obvious problems with default logic are avoided, and no additional machinery is needed to revise default theories.

Acknowledgements

I would like to thank my supervisor Randy Goebel for supporting and guiding me and for providing encouragement when I needed it the most. I must thank Li Yan Yuan for providing significant technical inputs. Thanks also to Jia You and Bernard Linsky. A heartfelt thank you to Abdul Sattar, Pavan Sikka, Suryanil Ghose and Narendra Ravi, Also to Charles Truscott, Andrew Edgar and others in the AI group.

Contents

1 Introduction **1**

1.1 Motivation 1

1.2 Outline 3

2 A Framework for Belief Change **4**

2.1 A Taxonomy of Belief Change 4

2.2 Rationality Postulates for Belief Change 5

2.3 Contraction Operators 6

2.4 Coherence vs. Foundational Belief Change 8

2.5 Rationality and Logical Omniscience 10

2.6 A Revised Set of Rationality Postulates 12

3 Changing Epistemically Prioritized Beliefs **14**

3.1 Contracting Prioritized Belief Sets 14

3.2 Rationality Results for $\sim$ and $-$ 20

3.3 Related Work 22

4 Default Reasoning as Belief Change **25**

4.1 A Framework for Belief Change 25

4.2 Reiter's Default Logic : Some Unintuitive results 29

4.3 Default Reasoning as Belief Change 31

4.4 Relation to PJ-Default Logic 36

4.5 Related Work 38

5 Conclusion **40**

5.1 Contributions 40

5.2 Directions for Future Work 41

List of Figures

1 Foundational Belief Change 10
2 Backgorund Belief Sets 26

Chapter 1

Introduction

1.1 Motivation

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 [9], [1], [11] and [8] have proposed a set of postulates (henceforth the AGM postulates) which are claimed to characterize 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 [13]. They conclude that many existing semantics do not satisfy these postulates. Those that do, have other undesirable characteristics; Dalal's semantics [3], 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 [6] have considered epistemically non-uniform databases, but the semantics they proposed assumed a pre-specified *total order* on all items of knowledge. Similarly, Gärdenfors and Makinson [11] have proposed a revision operator for belief sets which are totally ordered. But total orders require complete information about application domains; such information is difficult, if not impossible, to articulate. Partial orders may be all we can expect.

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 possible, and the latter concentrates 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 [21], Fuhrmann [7] and Nayak [20] have considered updates in the presence of partial orders, and we address several obvious problems with their approaches. 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., [19]).

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 [26], 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. We examine how this operator, and a similar one based on a logically omniscient framework, measure up to the standards of rationality set forth in the AGM postulates.

The belief change framework that we develop turns out to be an interesting one, not merely because it captures some aspects of limited reasoners changing beliefs in the presence of incomplete information, but because a modified version (which relaxes the requirement that a belief state be consistent) is a very elegant, and useful, way of looking at default reasoning. We show how default logic can be translated into this belief change framework, and how this solves some of the problems associated with default logic. Thus, unlike Brewka [2], Nebel [22] and Gärdenfors and Makinson [10] [18], who merely suggest that it is possible to look at default reasoning as belief change, we claim that it is actually better to do so.

1.2 Outline

Chapter 2 examines the currently popular AGM framework for belief change and proposes some modifications to avoid the problems of logical omniscience and to make the framework foundational. The rationality postulates are revised to reflect these changes. In Chapter 3, the notion of an *epistemically prioritized set* is defined to permit a partial specification of the relative epistemic priorities amongst beliefs. Two belief change operators which utilize this partial prioritization information are defined, one based on a non-omniscient framework and another on an omniscient one. The question of how far these operators measure up to the requirements for rationality is examined. Chapter 4 presents a modified belief change framework in which a belief state is not required to be consistent, and explains how default reasoning can be viewed as belief change in this framework. Chapter 5 outlines the contributions of this study and identifies directions for future work.

Chapter 2

A Framework for Belief Change

Most of the current research into belief change takes as a starting point the work of Alchourrón, Gärdenfors and Makinson [1]. Their approach, which we shall henceforward refer to as the *AGM framework*, provides a useful, if idealized, model for studying belief change. In this chapter, we shall study the AGM framework in some detail, identify some shortcomings, and modify our paradigm accordingly.

2.1 A Taxonomy of Belief Change

In the AGM framework, belief change may be of three types: AGM framework

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* [17] shown below:

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

Contraction can similarly be viewed as a revision operation, as given by the *Harper identity* [12]shown below:

$$K^{-}_{A} = K^{*}_{A} \cap K.$$

Given that expansion is trivial, and given that revision and contraction can be defined in terms of each other, we could choose to make either operation the basis of our study of belief change. We choose contraction because it is, in some ways, simpler to study and because it appears to be more intuitive to view contraction and expansion as the fundamental operations, and revision as a composition of these.

2.2 Rationality Postulates for Belief Change

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. Different sets of AGM postulates exist for expansion, contraction, and revision which can be related via the equivalences mentioned in the previous section. Since we choose to study belief change in terms of contraction, we shall concentrate only on 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 producing the result of a belief revision process, based on the input of one new proposition and a deductively closed set. Their 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-) $K_A^- \subseteq K$.

(3-) If $K \not\models A$ then $K_A^- = K$.

(4-) If $\not\models A$ then $A \notin K_A^-$.

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

(6-) If $\models A \leftrightarrow B$ then $K_A^- = K_B^-$.

(7-) $K_A^- \cap K_B^- \subseteq K_{A \wedge B}^-$.

(8-) If $A \notin K_{A \wedge B}^-$ then $K_{A \wedge B}^- \subseteq K_A^-$.

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 postulate, 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 Contraction Operators

Based on these rationality postulates, a number of possible contraction operators have been studied in [1] and [8]. *Full meet contraction* and *maxichoice* *contraction* represent limiting cases; in the former, too little is retained and in the latter, too much. *Partial meet contraction* represents a more reasonable middle ground.

Each of these contraction operators require the identification of maximal subsets of the deductively closed belief set K which do not imply the belief A being contracted.

Definition 1:$K \downarrow A = \{B \subseteq K \mid B \not\models A$ and if $B \subset C \subseteq K$ then $C \models A\}$.

Full meet contraction represents the situation where an agent takes into account all possible outcomes of the contraction, given by the elements of $K \downarrow A$, and conservatively chooses to believe in whatever is common to all these possible outcomes.

Definition 2:*Let K_A^- denote the full meet contraction of A from K.*

$$K_A^- = \begin{cases} \bigcap(K \downarrow A) & \textit{if} \not\models A \\ K & \textit{otherwise} \end{cases}$$

The problem with full meet contraction is that it removes too much information. It can be shown that when a belief set K is revised using a revision operator based on full meet contraction (via the *Levi identity*), only the consequences of the new belief are retained.

Maxichoice contraction assumes that there exists a selection function C that selects one element of $K \downarrow A$. The outcome of maxichoice contraction is the selected maximal (with respect to set inclusion) subset of K that does not imply A.

Definition 3:*Let K_A^- denote the maxichoice contraction of A from K.*

$$K_A^- = \begin{cases} C(K \downarrow A) & \textit{if} \not\models A \\ K & \textit{otherwise} \end{cases}$$

The problem with maxichoice contraction is that it retains too much information. It can be shown that revision based on maxichoice contraction (via the *Levi identity*) results in *complete* belief sets, i.e., a belief set that contains z or $\neg z$ for every z, even though the original belief set contained neither.

Partial meet contraction is a more reasonable operation representing the middle ground between maxichoice and full meet contraction. We assume now that there exists a selection function S which selects some subset of

$K \downarrow A$. The result of partial meet contraction is the intersection of the sets selected by S.

Definition 4: *Let K_A^- denote the partial meet contraction of A from K.*

$$K_A^- = \begin{cases} S(K \downarrow A) & \text{if} \not\models A \\ K & \text{otherwise} \end{cases}$$

Full meet contraction satisfies all the AGM postulates for contraction. Any maxichoice and partial meet contraction operator satisfies the first six of the AGM postulates for contraction. If there exists a preference relation over all subsets of K such that the selection function S selects the "best" or "most preferred" ones from $K \downarrow A$, then partial meet contraction can be shown to satisfy the first seven of the AGM postulates for contraction. If this preference relation is transitive, then partial meet contraction satisfies all eight of the AGM postulates for contraction.

2.4 Foundational Belief Change

It is popular to distinguish foundation from coherence approaches to belief revision, although this distinction may be more an artifact of the method of specification (e.g., of proof theory). The distinction hinges on which components of a belief set can be discarded during the belief change operation. The *coherence theory of belief change* requires that minimal changes be made to the original set of beliefs. The intuition is one common to theories of knowledge growth in science where "coherence" is achieved by somehow making minimal changes to an existing set of beliefs. The justification of an individual belief amongst those in a coherent set of beliefs is not its provability with respect 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 5: 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 might result in the retention of both (ii) and (iii), giving the set containing (ii), (iii), and (iv). The result is obviously counter-intuitive, since our belief in (iii) was contingent on our belief in (i). 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. However, given that our approach is essentially syntax-based, the distinction is still useful.

The approach taken by Gärdenfors, Alchourron and Makinson [9], [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 based on the coherence view. As Gärdenfors admits [8], the issue of maintaining the connections between premises and the consequences they support (i.e., *reason maintenance*) is largely 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 change.

To modify the AGM framework from a coherentist to a foundational one, we define a *belief base* to be a finite set of propositional formulae. The belief base is taken to represent the finite set of self-evident beliefs. Corresponding to each belief base is a *belief set* which represents its *belief closure*. For now, we shall take belief closure to be full deductive closure, but in the next section we shall define a notion of *partial closure* which appears to be more appropriate.

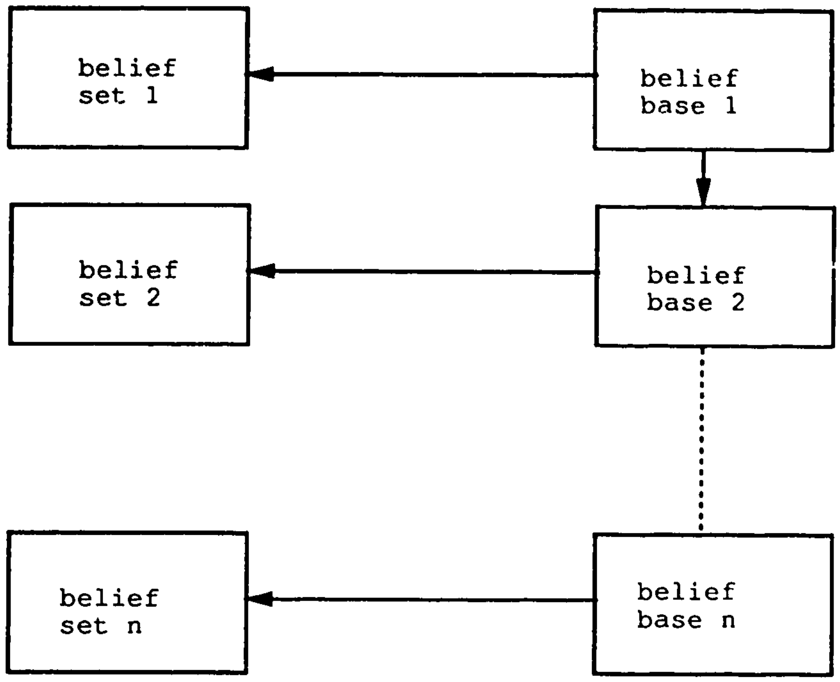

Figure 1: Foundational Belief Change

Figure 2.1 explains the relationship between belief bases and belief sets. Foundational belief change may be viewed as one set of self-evident beliefs yielding another. Corresponding to each such set of self-evident beliefs is a belief set representing the total set of beliefs held at that point of time.

2.5 Rationality and Logical Omniscience

In the previous section we argued that the AGM framework of deductively closed sets of beliefs is unintuitive for most real-life applications of belief change. However, the fact that belief change is not foundational in this framework is not the only problem with it. It also 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 logical consequences of their beliefs. What 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 6: 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 [26]:

Example 7: Consider the belief set $\{a, b\}$. Let Cn(X) denote the deductive closure of a belief set X. Then $\mathrm{Cn}(\{a,b\})=\{a, b, a \vee b, a \wedge b, a \rightarrow b, b \rightarrow a, \ldots\}$. When retracting a from this belief set both $\mathrm{Cn}(\{b\})$ and $\mathrm{Cn}(\{\neg b \vee a\})$ 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 this sense, the set $\mathrm{Cn}(\{\neg b\})$ 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 even stronger position 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 further prohibits their use as a framework for belief representation. A suitable belief closure should at least be finitely representable, and 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 implicants* . However, we base our definition on the one given in [26].

We first recall some definitions. We first restrict ourselves to the language of propositional logic. 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 8 : *The partial closure of a theory T, denoted by T^*, is defined as $T^* = T \cup \{\mu \mid T \vdash \mu$ 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 9: $\{a, a \rightarrow b, c, c \rightarrow d\}^{*} = \{a, a \rightarrow b, b, c, c \rightarrow d, d\}$. □

The partial closure of a theory thus represents the set of minimally derivable clauses which do not contain any tautologies. So, for example, 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 [19] for a good discussion of various approaches to addressing the logical omniscience problem in logics of knowledge and belief). But, 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.

2.6 A Revised Set of Rationality Postulates

We have, till now, suggested two significant changes to the AGM framework: we first suggest that a foundational approach may be more appropriate than a coherentist one, and we then identify some problems with full deductive closure which seem to be avoided with partial closure. It is necessary, at this point, to change the AGM rationality postulates to accomodate these changes in the framework.

Consider a finite belief base Δ of propositional clauses. The corresponding belief set is given by the partial closure of the belief base, given by $K = \Delta^{*}$, where Δ^{*} denotes the partial closure of Δ. Let $K^{\sim}_{A}$ denote the partially closed belief set that is obtained by contracting the belief A from the belief set K. Following the style of the original AGM postulates, we define the rationality postulates in terms of belief sets instead of belief bases. The revised set of rationality postulates for contraction are as follows:

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

(2∼) $K^{\sim}_{A} \subseteq K$.

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

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

(5∼) If $K \models A$ then $K \subseteq (K^{\sim}_{A})^{+}_{A}$.

(6$\sim$) If $\models A \leftrightarrow B$ then $K^{\sim}_{A} = K^{\sim}_{B}$.

(7$\sim$) $K^{\sim}_{A} \cap K^{\sim}_{B} \subseteq K^{\sim}_{A \wedge B}$.

(8$\sim$) If $K^{\sim}_{A \wedge B} \not\models A$ then $K^{\sim}_{A \wedge B} \subseteq K^{\sim}_{A}$.

The changes made to the original postulates are fairly obvious, given the framework of partially closed theories. Postulate (1$\sim$) 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$\sim$), (5$\sim$) and (7$\sim$).

Chapter 3

Changing Epistemically Prioritized Beliefs

In the previous chapter, we modified the AGM framework to make it foundational and we introduced the notion of partially closed belief sets. We shall now deviate from the AGM framework in another crucial way: we shall permit a partial specification of relative epistemic priorities amongst beliefs, as opposed to a total ordering that the AGM framework requires. In the context of these changes, we shall define two new contraction operators and evaluate how they measure up to the requirements for rationality.

3.1 Contracting Prioritized Belief Sets

A common situation during belief revision is that a number of alternative belief sets can potentially constitute the revised belief set. If one were to take a disjunction of these mutually inconsistent belief sets, as in [21], or their intersection, as in [9], one would have to give up most of one's previously held beliefs, with the result that the new belief set would contain only the consequences of the current input.

Example 10: Consider a set of beliefs represented by: $\{a, a \rightarrow b\}$. As a result of an 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 considering 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 [9] and [11], where beliefs that are more epistemically entrenched have greater utility for the purpose of inquiry and decision-making and vice versa. Epistemic entrenchment is detached 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 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 degrees of certainty, or probability are all acceptable for the purpose of assigning epistemic priorities. We then define a framework which permits a possibly incomplete specification of the relative epistemic priorities amongst beliefs.

Definition 11: *An epistemically prioritized belief set is one in which a 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, we want to use the extra-logical ordering on the beliefs to rationally choose amongst several possible alternative belief sets. But 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 [9] and [11], but permit situations where beliefs are incomparable under the ordering $\geq$. In other words, we want to proceed with our choice given only partial information about epistemic priority.

As before, we define a *belief base* to be a finite set of propositional clauses. Note that finite belief bases have also been used by Nebel in [21], but he too considers a total ordering on the propositions of the belief base.

Definition 12: *A prioritized belief base is a finite set of propositional clauses for which a partial epistemic prioritization $\leq$ exists.*

As noted above, our contraction operator intends to exploit whatever epistemic priority information exists, including simple statements like $\alpha \leq \beta$, and their algebraic consequences.

Definition 13: *A belief set corresponding to a prioritized belief base Δ is given by Δ^*.*

Note that our definition refers to two sets of sentences: an original finite set of sentences, the belief *base*, and the partial closure of that belief base, which we call the belief *set*. In what follows below, it is important to distinguish these.

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} selects from amongst $C_{max}(\Delta, A)$ those sets that contain sentences of higher epistemic priority; note that no set in $C_{max}(\Delta, A)$ dominates, in terms of epistemic content, any set in $E_{max}(C_{max}(\Delta, A))$. If it turns out that every element of $C_{max}(\Delta, A)$ is dominated by some other element, then the entire set $C_{max}(\Delta, A)$ is taken to be the result of applying the operator E_{max}.

- 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 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 conservatively believe only those sentences common to all possible views.

The input to this contraction operation consists of a sentence to be removed, and a belief base which has been provided directly or obtained from the previous belief change step. The contraction operation yields a new contracted belief base and a corresponding belief set which is the the partial closure of the belief base.

Definition 14: *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* $p > q$.

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

Definition 15: *The set* $C_{max}(\Delta, A)$ *of maximal subsets of the belief base* Δ *that do not imply* A, *where* A *is a propositional sentence, is defined as*

follows:
$C_{max}(\Delta, A)=\{S \mid S \subseteq \Delta$ *and* $S \not\models A$ *and for any* S'*such that* $S \subset S' \subseteq \Delta$, $S' \models A\}$.

It is easy to see that C_{max} is identical to the operation $\downarrow$ defined in Chapter 2, except that it is now defined in the context of finite belief bases.

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.

Definition 16: *If for every* $s \in X$ *there exists some other* $s' \in X$ *such that* s' *dominates* s *, then* $E_{max}(X) = X$. *Otherwise* $E_{max}(X) = \{S \mid S \in X$ *and there exists no* $S' \in X$ *such that* S' *dominates* $S\}$

Definition 17: *Given a prioritized belief base* Δ, *a belief set* K *given by* $K = \Delta^*$ *and an epistemic prioritization* $\leq$, *the contraction of* K *with some sentence* A *is given by:*

$$K_A^\sim = \begin{cases} \bigcap(E_{max}(C_{max}(\Delta, A)))^* & \textit{if} \not\models A \\ K & \textit{otherwise} \end{cases}$$

For notational convenience, if X *is a set of sets of clauses, we take* X^* *to denote the set of partial closures of each of its elements, and* $\bigcap X$ *to denote the intersection of its elements.*

Note that the result K in the above definition will also be a belief set (i.e., a partially closed belief base), since the intersection of a set of partially closed sets remains partially closed. We shall now define the contracted belief set to also be the new belief base. We shall use the same symbol $\sim$ to denote the result $\Delta_A^\sim$ of contracting a sentence A from a prioritized belief base Δ without ambiguity; the context shall indicate whether it operates on belief bases or belief sets.

Definition 18: *The contraction of a sentence* A *from a prioritized belief base* Δ *yields a new prioritized belief base* $\Delta_A^\sim$ *defined as follows:*

$$\Delta_A^\sim = \begin{cases} \bigcap(E_{max}(C_{max}(\Delta, A)))^* & \textit{if} \not\models A \\ \Delta & \textit{otherwise} \end{cases}$$

If $\leq$ *is the partial ordering relation defined on* Δ, *then the subset of* $\leq$ *that is defined on* $\Delta_A^\sim$ *is the prioritization relation that applies to* $\Delta_A^\sim$.

It is easy to see that the belief set that results from a contraction operation represents the partial closure of the contracted belief base. The problem with this definition is that it promotes some some non-self-evident beliefs to the status of self-evident beliefs (these are the beliefs that are contained in the partial closures but not in the original belief base). We can describe belief change in this case to be ***stepwise foundational.*** In a single belief change step, given a set of self-evident beliefs, the result does not contain any non-justified beliefs. This property, however, may not hold over a series of belief change steps.

Example 19: 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:

$$C_{max}(\Delta,\ b)=\{\{a,c\}, \{a, c \rightarrow b\}, \{c, a \rightarrow b\}, \{a \rightarrow b, c \rightarrow b\}\}.$$

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 $C_{max}(\Delta,\ b)$, and that they are mutually incomparable with respect to dominance. Thus they are the only epistemically maximal subsets in $C_{max}(\Delta,\ b)$.

$$E_{max}(C_{max}(\Delta, b)) = \{\{a,c\}, \{a, c \rightarrow b\}\}.$$

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

$$K_b^{\sim} = \{a,c\}^{\star} \cap \{a, c \rightarrow b\}^{\star} = \{a\}.$$
$$\Delta_b^{\sim} = \{a\} \ \square$$

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 [26].

Definition 20: *Let* T_1 *and* T_2 *be any two theories and let* $\Pi = \{T_1, \ldots, T_n\}$ *be a set of theories.*

- $ed(T_1, T_2) = \{T_1 \cap T_2\} \cup \{\alpha_i \vee \alpha_j \mid \alpha_i \in (T_1 - T_2) \textit{ and } \alpha_j \in (T_2 - T_1)\}$.
- $ED(\Pi) = \begin{cases} ed(\Pi) & \textit{if cardinality of } \Pi \textit{ is 2} \\ ed(T_1, ED(\Pi - T_1)) & \textit{if cardinality of } \Pi \textit{ is greater than 2} \end{cases}$

The foundational contraction operator — for logically omniscient agents is defined below. Since we consider logically omniscient agents, belief closure is now full deductive closure.

Definition 21: *Let* Δ *be a prioritized belief base and* K *be the corresponding deductively closed belief set given by* $K = Cn(\Delta)$. *Then the contraction of a sentence* A *from* K, *given by* K_A^-, *is defined as follows:*

$$K_A^- = \begin{cases} Cn(ED(E_{max}(C_{max}(\Delta, A)))) & \textit{if} \not\models A \\ K & \textit{otherwise} \end{cases}$$

Definition 22: *Let* Δ *be a prioritized belief base. The prioritized belief base* Δ_A^- *that results from contracting the sentence* A *from* Δ *is defined as follows:*

$$\Delta_A^- = \begin{cases} ED(E_{max}(C_{max}(\Delta, A))) & \textit{if} \not\models A \\ \Delta & \textit{otherwise} \end{cases}$$

If $\leq$ *is the partial ordering relation defined on* Δ, *then the subset of* $\leq$ *that is defined on* Δ_A^- *represents the partial ordering that prioritizes the new belief base* Δ_A^-.

It is easy to see that the resultant logically omniscient belief base K_A^- represents the deductive closure of the resultant belief base Δ_A^-. Belief change in this case is not merely stepwise foundational as it was with the previous operator, but fully foundational in the sense that the foundational property holds over any number of belief change steps.

Example 23: Consider the prioritized belief base B given by:

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

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

$C_{max}(\Delta, ((c \wedge d) \vee b)) = \{\{a, c\}, \{a, d\}, \{a \rightarrow b, c\}, \{a \rightarrow b, d\}\}$.
$E_{max}(C_{max}(\Delta, ((c \wedge d) \vee b))) = \{\{a, c\}, \{a, d\}\}$.
$K^{-}_{(c \wedge d) \vee b} = Cn\{a, c \vee d\}$
$\Delta^{-}_{(c \wedge d) \vee b)} = \{a, c \vee d\}$. □

3.2 Rationality Results for $\sim$ and $-$

In this section, we shall examine the extent to which the two contraction operators $\sim$ and $-$ satisfy the rationality postulates for belief change. The following two observations wil be useful in the proofs that follow.

Lemma 24: $C_{max}(\Delta, A \wedge B) = C_{max}(\Delta, A) \cup C_{max}(\Delta, B)$.

Lemma 25: *For any p,* $ED(T_1, \ldots, T_i, \ldots, T_n) \models p$ *iff* $T_i \models p$ *for* $1 \leq i \leq n$.

We shall first consider the operator $\sim$. Since belief closure in this case is partial closure, we shall use the revised postulates (1$\sim$)-(8 $\sim$) to evaluate this operator.

Theorem 26: *The contraction operator $\sim$ satisfies postulates (1$\sim$)-(4$\sim$) and (6$\sim$) and (7$\sim$).*

Proof : For postulates (1$\sim$)-(4$\sim$) and (6$\sim$) the proof is trivial. The proof for postulate (7$\sim$) is given in the appendix. □

That postulates (5$\sim$) and (8$\sim$) are not satisfied by the operator $\sim$ is not surprising. As we shall see, this is a consequence of the foundational, logically non-omniscient framework and the fact that partial specification of epistemic priorities are permitted, rather than any fundamental shortcoming of the operator itself. Let us consider first postulate (5$\sim$), the so-called *recovery postulate.* This postulate requires that if we add a belief that we had earlier contracted, we should get back at least the original set of beliefs. Superficially, that is an absolutely valid requirement. However, as the following example shows, a perfectly rational approach that makes maximum use of the available information can result in situations where this requirement is not met.

Example 27: Consider the prioritized belief base Δ given below:

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

Let $a \rightarrow b > a$ be the only $>$-ordering relation derivable from $\leq$.

$$K_b^{\sim} = \{a \rightarrow b\}$$
$$(K_b^{\sim})_b^{+} = \{a \rightarrow b, b\}$$

As a result of restricting our attention exclusively to finite belief bases, once a belief a has been discarded, there is no way of retrieving it. The recovery postulate is thus not satisfied. □

We shall now consider an example of a situation where postulate (8$\sim$) is not satisfied by the operator $\sim$.

Example 28: Consider the prioritized belief base Δ given below.

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

Let $a \rightarrow b > c \rightarrow d$ and $c > a$ be the only $>$-ordering relations derivable from $\leq$.

$$C_{max}(\Delta, b \wedge d) = \{\{a, c, c \rightarrow d\}, \{a \rightarrow b, c, c \rightarrow d\}, \{a, a \rightarrow b, c\}, \{a, a \rightarrow b, c \rightarrow d\}\}$$
$$E_{max}(C_{max}(\Delta, b \wedge d)) = \{\{a, a \rightarrow b, c\}, \{a \rightarrow b, c, c \rightarrow d\}\}$$
$$K_{b \wedge d}^{\sim} = \{a \rightarrow b, c\}$$
$$C_{max}(\Delta, b) = E_{max}(C_{max}(\Delta, b)) = \{\{a, c, c \rightarrow d\}, \{a \rightarrow b, c, c \rightarrow d\}\}$$
$$K_b^{\sim} = \{c, c \rightarrow d, d\}$$

Thus, although $K_{b \wedge d}^{\sim} \not\models b$, $K_{b \wedge d}^{\sim}$ is not contained in $K_b^{\sim}$. Postulate (8$\sim$) is violated.□

Once again, although the contraction operation makes perfectly valid choices using all the available information, this rationality postulate is not satisfied. Interestingly, Willard and Yuan report in [26] that their revision operator, which also uses a form of partial closure, and assumes a fixed partial prioritization in which rules of the form $a \rightarrow b$ have higher priority over facts of the form a, does not the satisfy the corresponding eighth postulate for revision. Katsuno and Mendelzon [13] suggest that the eighth postulate for revision is too strong if the ordering of interpretations is a partial and not a total one (their considerations were model-theoretic).

We shall now evaluate the operator $-$ in terms of the original AGM postulates for contraction, since the closure in this case is full deductive closure. The results we obtain are surprisingly similar to those for $\sim$.

Theorem 29: *The contraction operator $-$ satisfies postulates (1-)-(4-) and (6-) and (7-).*

Proof : The proof for postulates (1-)-(4-) and (6-) is trivial. The proof for (7-) is given in the appendix. □

Consider the following examples of situations where postulates (5-) and (8-) are violated.

Example 30: Consider the belief base of Example 27.

$$K_b^- = Cn(a \rightarrow b)$$
$$(K_b^-)_b^+ = Cn(a \rightarrow b, b)$$

The recovery postulate is thus not satisfied. □

Example 31: Consider the belief base of Example 28.

$$K_{b \wedge d}^- = Cn(a \rightarrow b, c, a \vee \neg c \vee d)$$
$$K_b^- = Cn(c, c \rightarrow d, d)$$

Postulate (8-) is not satisfied. □

The above results lead us to suspect that these requirements are perhaps too strong.

3.3 Related Work

Our approach differs from the work of Alchourrón, Gärdenfors and Makinson [1], [9], [11], [8] 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, $\alpha \notin \text{Cn}(\Delta)$ 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 [21] 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 [7] and Nayak [20] have proposed a foundational contraction operator. When revising with a propositional sentence A, they define $E(A)$ 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 32: 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.□

Fagin, Ullman and Vardi have proposed update semantics for prioritized databases in [6]. 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 [26] 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 [13], 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 [24] have analyzed both the foundational and coherentist approaches to belief revision, but their approach assumes a modal logic with auto-epistemic operators.

Chapter 4

Default Reasoning as Belief Change

We begin this chapter by introducing a more general framework for belief change by relaxing the requirement that a belief state represent a consistent set of beliefs. Armed with a belief change framework which permits possibly inconsistent sets of beliefs, we propose a method for performing default reasoning as belief change and identify the benefits that this approach has to offer.

4.1 A Framework for Belief Change

A common feature of most of the existing work on belief change, including the approach we have taken in previous chapters, is that the process is viewed as a transition from one belief state to another. A belief state is a consistent view of the world, which, in the AGM framework, and in our modified framework, translates to a consistent set of propositional formulae. We will now look at belief change in a somewhat different manner. We shall relax the consistency requirement so that a belief set is now treated not as a consistent set of beliefs but instead as a possibly inconsistent set of beliefs that may be potentially held (we shall call this the *background belief set* together with information on the relative epistemic priorities amongst the beliefs. Belief change is then a transition from one *background belief set* to another. Corresponding to each belief state (given by a background belief set), a consistent view of the world (given by a consistent set of propositional formulae, called the *belief set*) can be identified. Figure 4.1 below presents this idea graphically.

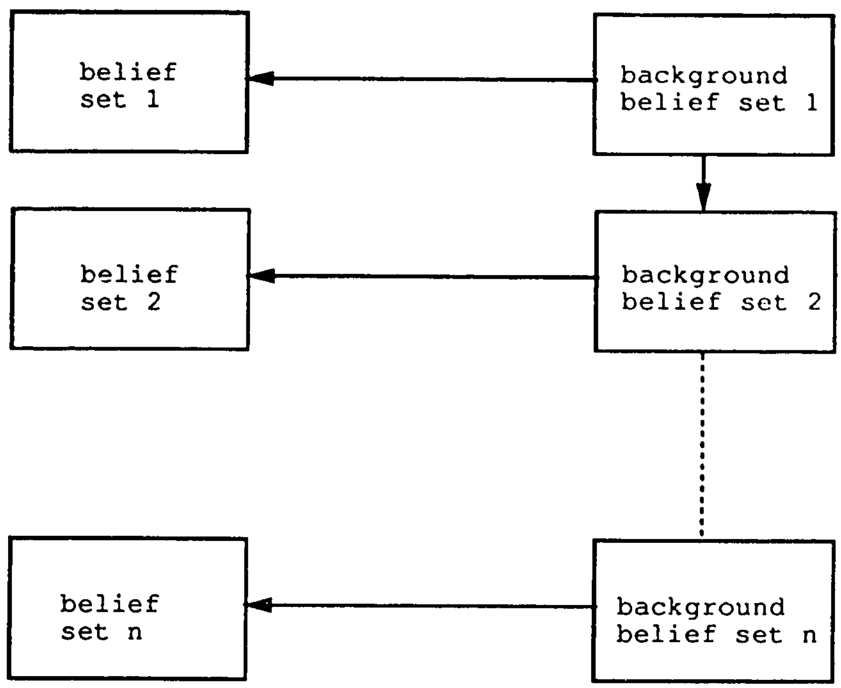

Figure 2: Backgorund Belief Sets

This framework subsumes the foundational framework that has been described earlier. If the background belief set is restricted to be finite and consistent this framework coincides with the foundational one.

Yuan, You and Bissonauth have considered a similar framework in [29]. They propose a scheme in which two versions of a knowledge base are maintained: the actual, possibly inconsistent knowledge base which contains "historical" information, and the subset of this knowledge base that is actually visible to the user and that is used for query processing. Their "historical" knowledge base corresponds, in our terminology, to the background belief set, while their "visible" knowledge base corresponds to our belief set. Significantly, they show that their update semantics satisfies all the rationality requirement sof the AGM postulates.

In the AGM framework, belief change is classified into three different kinds of operations: *expansion*, *revision* and *contraction*. In the following, we examine the counterparts of each of these operations in our framework:

- *Expansion* in the AGM framework is the operation of adding a new belief which is consistent with all existing beliefs. In this framework, expansion involves adding the new belief to the background belief set in such a way that the new belief is guaranteed to appear in the consistent belief set identified from this background belief set. What this essentially means is that the new belief is assigned the highest possible priority, so that every consistent view of reality that obtains its beliefs from this background belief set is guaranteed to contain the new belief

- *Revision* in the AGM framework involves adding a new belief that is possibly inconsistent with existing beliefs. In the current framework, this turns out to be the same as expansion. The new belief is added to the background belief set in such a way that the new belief is guaranteed to appear in the corresponding belief set. As before, this means assigning the highest possible priority to the new belief.

- *Contraction* in the AGM framework involves retracting a currently held belief. In the current framework, this corresponds to identifying a consistent belief set which does not contain the belief being contracted from the background belief set.

As with our earlier framework, we allow a partial specification of the relative epistemic priorities amongst beliefs. We, however, relax the requirement that beliefs be represented in clausal form. We now recall some definitions in slightly reformulated form.

Definition 33: *An epistemically prioritized set is a set of beliefs (represented as propositional formulae) in which a 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* $\alpha < \beta$.

Definition 34: *A background belief set is a possibly inconsistent set of propositional formulae.*

Definition 35: *A belief set is a consistent set of propositional formulae.*

The following is an informal description of our contraction operator. The approach is essentially the same as that for operators $\sim$ and $-$ defined in Chapter 3. We shall see later that the operations of *expansion* and *revision* in the AGM framework can be expressed as special cases of this contraction operator. Let Δ be the current background belief set, i.e., the set of propositional beliefs from which a consistent view of the world may be constructed. Let a propositional formula A denote the belief that is to be contracted.

- The first step is to identify maximal (with respect to set inclusion) subsets of Δ which are consistent and which do not imply A. Each of these subsets can now be referred to as a *belief set* since they are consistent sets of beliefs. This set of belief sets is given by $c_{max}(\Delta, A)$.

- The second step is to identify the *epistemically maximal* belief sets from the elements of $c_{max}(\Delta, A)$. This is given by $e_{max}(c_{max}(\Delta, A))$.

- Each of the belief sets contained in $e_{max}(c_{max}(\Delta, A))$ qualifies as a valid view of the world that may be held by an agent. The question of which of these beliefs sets, or which combination of these belief sets should be taken to actually represent the agent's view of the world at that point may be addressed in different ways. A credulous reasoner would pick any one of the belief sets as the valid one. A skeptical reasoner, on the other hand, would take the intersection of the belief sets. For our present purposes, it is not important to commit ourselves to any one approach. Since our purpose is to show the relationship between default reasoning and belief change, it is sufficient to show the relation between the multiple views of the world sanctioned by a knowledge base containing defaults and the multiple consistent views of the world given by the elements of $e_{max}(c_{max}(\Delta, A))$. We shall therefore assume that there exists some operator F such that $F(e_{max}(c_{max}(\Delta, A)))$ which identifies the belief set that constitutes the agent's actual view of the world.

Definition 36: *The set $c_{max}(\Delta, A)$ of maximal subsets of the background belief set Δ which are consistent and which do not imply A, where A is a propositional sentence, is defined as follows:*
$c_{max}(\Delta, A)=\{S \mid S \subseteq K$ *and* S *is consistent and* $S \not\models A$ *and for any* S'*such that* $S \subset S' \subseteq K,\ S' \models A\}$.

Definition 37: $e_{max}(X)=\{S \mid S \in X$ *and there exists no* S' *such that* S' *dominates* $S\}$.

The notion of *dominance* used here is the same as that defined earlier.

Definition 38: *The contraction of a belief A from a background belief set* Δ, *given by* $CONTRACT(\Delta, A)$, *is defined as follows:*
$CONTRACT(\Delta, A) = F(e_{max}(c_{max}(\Delta, A)))$.

Appropriate versions of F will yield the operators $\sim$ and $-$ defined in Chapter 3.

It is easy to see how the operations of *expansion* and *revision* (which coincide in this framework) can be expressed in terms of the operators c_{max} and e_{max} that we have already defined. Both operations involve adding the new belief to the background belief set as beliefs with the highest priority (this ensures that they are included in every resultant belief set), and identifying the consistent belief sets from this background belief set.

Definition 39: *Let* $EXPAND(\Delta, A)$ *and* $REVISE(\Delta, A)$ *denote the expansion and revision, respectively, of a background belief set* Δ *with a new belief A.*
$EXPAND(\Delta, A) = REVISE(\Delta, A) = F(e_{max}(c_{max}(\Delta \cup \{A\}, \{\})))$. *Note that the new belief A is assigned a priority higher than that of any other belief in* Δ.

4.2 Reiter's Default Logic : Some Unintuitive results

Default logic seeks to formalize reasoning with incomplete information by augmenting first-order logic with defeasible rules of inference called *default rules*. The general form of a default rule is :

$$\frac{\alpha(\vec{x}):\beta(\vec{x})}{\gamma(\vec{x})}$$

$\alpha(\vec{x})$ is called the *prerequisite*, $\beta(\vec{x})$, the *justification* and $\gamma(\vec{x})$ the *conclusion*. Informally, a default rule may be interpreted as the statement "if, for some instance $\vec{a}$ of $\vec{x}$, $\alpha(\vec{a})$ is provable and $\beta(\vec{a})$ is consistent with what is known, then conclude by default that $\gamma(\vec{a})$". A *default theory* is a pair (W, D), where W is a first-order theory while D is a set of default rules. We shall consider only propositional default theories here, primarily because most existing work on belief change, to which we shall attempt to relate default logic, considers only propositional beliefs.

Before moving on to identify some of the shortcomings of classical default logic as defined by Reiter [25], we shall quickly review some useful results.

An *extension* E of a default theory is the set of beliefs that may be held about a world about which the formulae in W are known and to which the default rules in D are applicable.

Definition [25] 40: *If S is a set of propositional sentences, and $\Gamma(S)$ is the smallest set satisfying the following properties :*

1. $W \subset \Gamma(S)$.
2. $\Gamma(S) \subseteq Cn(\Gamma(S))$.
3. *if* $\frac{\alpha:\beta}{\gamma} \in D$ *and* $\alpha \in \Gamma(S)$ *and* $\neg\beta \notin S$ *then* $\gamma \in \Gamma(S)$.

then E is an extension of the default theory (W,D) iff $\Gamma(E) = E$.

Theorem [25] 41: *Let (W,D) be a propositional semi-normal default theory. Define:*
$E_0 = W$.
$E_{i+1} = Cn(E_i \cup \{\gamma \mid \alpha : \beta/\gamma \in D \text{ and } \alpha \in E_i \text{ and } \neg\beta \notin E\}$, *for* $i \geq 0$
Then E is an extension for (W,D) iff $E = \bigcup_{i=0}^{\infty} E_i$.

Recall that a *normal* default is one where the justification and conclusion are the same, while a *semi-normal* default is one where the justification implies the conclusion.

Default logic is a useful formalism for analyzing defeasible knowledge and a considerable body of work exists in which this has been studied in some detail [25] [5] [16] [14]. There are, however, some well-known examples of situations where default logic provides fairly unintuitive results. In the remainder of this section we shall look at these situations in greater detail.

We essentially identify three classes of problems:

- This is the well-known problem of disjunctive defaults. Consider a default theory (W, D) where $W = \{p \vee r\}$ and $D = \{\frac{p:q}{q}, \frac{r:q}{q}\}$. It is intuitively obvious that we would like to conclude q, given that we know that the prerequisite of at least one of the two defaults is true and that it is consistent to assume the justification in either case. Reiter's default logic, however, gives a single extension which does not contain q. Notice that we also cannot reason about defaults. There is no explicit connection, for example, between $\frac{p:q}{q}$, $\frac{r:q}{q}$ and $\frac{p \vee r:q}{q}$, although it is intuitively obvious that such a connection exists and that establishing such a connection can be useful.

- Default logic sometimes sanctions conclusions which appear to be unwarranted. Consider the default theory (W, D) where $W = \{\neg b \vee \neg d\}$ and $D = \{\frac{:(a \wedge b)}{a}, \frac{:(c \wedge d)}{c}\}$. This theory has a single extension in which both a and c are true. The conclusions are intuitively unwarranted since it is clear that it is not consistent to assume both b and d at the same time.

- Yet another situation where default logic provides unintuitive results is given in the following example. Consider the default theory (W, D) where $W = \{\}$ and $D = \{\frac{:(a \wedge b)}{a}, \frac{:\neg b}{\neg b}\}$, which has a single extension in which b is false. Intuitively, however, this default theory sanctions two mutually incompatible views of the world; a is true in one and b is false in the other. Two extensions are thus expected, one containing a and the other containing $\neg b$.

The underlying problem is that defaults are not interpretable sentences, and there is no way to semantically relate the sentences of W with the defaults. Default logic therefore has to rely on proof-theoretic devices to establish a relation.

Revising default theories is a non-trivial problem [27] and the concept of rational default theory revision remains ill-defined. Rational revision of belief sets, where a belief set is thought to be a set of propositional formulae representing the beliefs of an agent, is, on the other hand, a problem that has been studied in considerable detail (see [1],[8], [13], [21] for a representative subset). In the next section we shall show that a simple approach to the problem of default reasoning exists, which involves viewing default reasoning as the process of belief set revision (or *belief revision*, as it is normally called). It shall be seen that not only does this approach avoid all of the problems with default logic mentioned above, we do not need any additional machinery to revise default theories.

4.3 Default Reasoning as Belief Change

The basis of our approach to viewing default reasoning as belief change is to treat defaults as sentences. We translate a default theory (W, D) to an *epistemically prioritized background belief set* $\Delta_{(W,D)}$. We then retract the disjunction of all possible invalid situations (an invalid situation being when the consequent of a default is concluded when the negation of the justification is provable) for each default in D from Δ. In other words, we construct consistent views of the world, which do not imply any invalid situation. We shall then investigate the relationship between the multiple possible belief sets that are obtained while performing this contraction and the extensions obtained from the default theory (W, D). That some relationship should exist is only expected, given that both default reasoning and belief change involve identifying consistent views of the world in the presence of inconsistent and incomplete knowledge.

Definition 42: *The Δ-translation of a default theory (W,D), given by $\Delta_{(W,D)}$, is defined as follows:*
$\Delta_{(W,D)} = Cn(W) \cup \{\alpha \rightarrow \gamma \mid \frac{\alpha:\beta}{\gamma} \in D\}$
with the following epistemic priority relations:

$$w > (\alpha \rightarrow \gamma)$$

for each $w \in Cn(W)$ and each sentence of the form $\alpha \rightarrow \gamma$ obtained from defaults of the form $\frac{\alpha:\beta}{\gamma}$ in D.

We now define the contraction operation that we believe is the closest counterpart, in the framework of belief change, of the process of obtaining extensions of a default theory.

Definition 43: *The j-contraction of a background belief set $\Delta_{(W,D)}$, i.e., the contraction of invalid combinations of beliefs for each default rule in D, is given by :*

$$e_{max}(c_{max}(\Delta_{(W,D)}, \bigvee(\gamma_i \wedge \neg\beta_i))) \text{ for all } i \text{ such that } Cn(W) \not\models \gamma_i$$

where $D = \bigcup_{i=1}^{n}\{\frac{\alpha_i:\beta_i}{\gamma_i}\}$ and n is the cardinality of D.

Note that in order to view the process of extension computation as a belief change process, we still need to refer to the default rules.

Example 44: This is the classical problem involving the default *birds fly*. Let the propositions B and F represent the facts that an individual x is a bird and the individual x flies, respectively. Consider a default theory (W, D) where:

$$W = \{B\}$$
$$D = \{\frac{B:F}{F}\}$$

The Δ-translation of this default theory is given by :

$$\Delta_{(W,D)} = Cn(B) \cup \{B \rightarrow F\}$$

where:

$$p > (B \rightarrow F)$$

for each $p \in Cn(B)$.

The default theory (W, D) has a single extension given by $Cn(B, F)$. The j-contraction of $\Delta_{(W,D)}$ yields a single belief set containing both B and F. If W is augmented with the belief $\neg F$, (W, D) has a single extension given by $Cn(B, \neg F)$. The j-contraction of $\Delta_{(W,D)}$ yields a single belief set containing both B and $\neg F$.□

Example 45: Consider the disjunctive default problem mentioned earlier. The default theory (W, D) is given by :

$$W = \{p \vee r\}$$
$$D = \{\frac{p:q}{q}, \frac{r:q}{q}\}$$

The Δ-translation of this default theory is given by:

$$\Delta_{(W,D)} = Cn(p \vee r) \cup \{p \rightarrow q, r \rightarrow q\}$$

where

$$x > (p \rightarrow q)$$
$$x > (r \rightarrow q)$$

for every $x \in Cn(p \vee r)$.

While the default theory yields a single extension which does not contain q, it is easy to see that the j-contraction of $\Delta_{(W,D)}$ yields a belief set which contains q in its deductive closure.□

Example 46: Consider the default theory (W, D) given by:

$$W = \{\neg b \vee \neg d\}$$
$$D = \{\frac{:(a \wedge b)}{a}, \frac{:(c \wedge d)}{c}\}$$

The Δ-translation of the default theory is given by:

$$\Delta_{(W,D)} = Cn(\neg b \vee \neg d) \cup \{a, c\}$$

where

$x > a$
$x > c$

for every $x \in Cn(\neg b \vee \neg d)$. The default theory has a single extension which contains both a and c. On the other hand, the j-contraction of $\Delta_{(W,D)}$ is given by:

$$e_{max}(c_{max}(\Delta_{(W,D)}, (\neg b \wedge a) \vee (\neg d \wedge c))$$
$$= \{Cn(\neg b \vee \neg d) \cup \{a\}, Cn(\neg b \vee \neg d) \cup \{c\}\}$$

It is easy to see that while the default theory provides unintuitive results, the two belief sets that result from the j-contraction of $\Delta_{(W,D)}$ correspond precisely to our intuition that we should believe in any one of a and c, but not both. □

Example 47: Consider the default theory given by:

$$W = \{\}$$
$$D = \{\frac{:(a \wedge b)}{a}, \frac{:\neg b}{\neg b}\}$$

The Δ-translation of this default theory is given by:

$$\Delta_{(W,D)} = \{a, \neg b\}$$

There are no epistemic prioritization relations amongst the elements of $\Delta_{(W,D)}$. The default theory sanctions a single extension containing $\neg b$, although our intuition dictates that there should be two extensions, one containing a, and the other containing $\neg b$. The j-contraction of $\Delta_{(W,D)}$ provides such a result, as seen below:

$$e_{max}(c_{max}(\Delta_{(W,D)}, (\neg b \wedge a) \vee (b \wedge \neg b)))$$
$$= \{\{a\}, \{\neg b\}\}$$

□

The preceding examples suggest that not only is it possible to look at default reasoning as belief change, the belief change approach actually avoids many of the unintuitive results obtained from default logic. The original default theory is translated to a background belief set in which facts and default rules have the same uniform representation as propositional sentences. Default rules are converted to simple material implications connecting the prerequisite and the conclusion, which closely corresponds to the intuition underlying a default rule: assume that the prerequisite implies the conclusion if it is consistent to do so. Such an assumption becomes inconsistent if the negation of the justification is provable. That this intuition is captured by expressing defaults as beliefs in a background belief set is not surprising. Every belief in a background belief set has the same intuition: assume the belief if it is consistent to do so. The problem is that the background belief set obtained by the Δ-translation of a default theory does not represent, in the same uniform representation, the conditions that block the application of a default. These conditions have to be represented meta-theoretically as the beliefs that are contracted from a background belief set during the process of identifying the possible views of the world sanctioned by the corresponding default theory. To see why this is not just an unintuitive device for obtaining the desirable results, consider the following informal prescription for representing a knowledge base containing *default knowledge* in this belief change framework and and reasoning with it. First, the knowledge is represented as propositional sentences. Then, known invalid combinations are contracted from this background belief set to obtain consistent possible views of the world. Viewed in this fashion, a belief change approach to default reasoning appears to be more intuitive than even default logic.

An important fallout of the translation that we have proposed is that we can now reason with the contrapositive of a default rule. While it is possible to argue both for and against the desirability of this, we do not view this a major drawback. Whereas Reiter's default logic requires that contrapositives be explicitly stated as default rules, the belief change view requires that contrapositives be explicitly disabled in situations where they are not desired.

Example 48: Consider the previous example involving the default *birds fly*. Assume that it is known that an individual x does not fly, represented by the proposition $\neg F$, but nothing is known regarding whether the individual is a bird or not. This is represented by the default theory (W, D) as follows:

$$W = \{\neg F\}$$
$$D = \{\frac{B:F}{F}\}$$

The corresponding Δ-translation is given by:

$$\Delta_{(W,D)} = Cn(\neg F) \cup \{B \rightarrow F\}$$

where

$p > (B \rightarrow F)$

for every $p \in Cn(\neg F)$. While the only extension of (W, D) would contain just $\neg F$, the j-contraction of $\Delta_{(W,D)}$ would yield a belief set containing both $\neg F$ and $\neg B$ in its deductive closure. To prevent the contrapositive from being applied we would have to explicitly block it. This would mean contracting $\neg F \rightarrow \neg B$ as an explicitly inadmissible condition. The result then would be a belief set which contains only $\neg F$ in its deductive closure.□

4.4 Relation to PJ-Default Logic

In this section, we shall look at a variant of default logic called PJ-default logic developed by Delgrande and Jackson [4] to address precisely the same problems with Reiter's default logic that we pointed out earlier. Not surprisingly, we shall suggest that our view of default reasoning as belief change corresponds exactly to PJ-default logic in the sense that the multiple possible views of the world that result from a j-contraction of the background belief set $\Delta_{(W,D)}$ are the same as the extensions of the PJ-default theory (W, D).

PJ-default logic is the combination of two separate variants of Reiter's default logic: P-default logic (or *prerequisite-free* default logic) and J-default logic (or *justification* default logic). P-default logic addresses the problem of the requirement for proving the antecedent of a default rule being too strong, which results in the disjunctive default problem mentioned earlier. In P-default logic, normal defaults of the form $\frac{\alpha:\beta}{\beta}$ are replaced with prerequisite-free defaults of the form $\frac{:\alpha\supset\beta}{\alpha\supset\beta}$. Semi-normal defaults of the form $\frac{\alpha:\beta\wedge\gamma}{\beta}$ are replaced by prerequisite-free default rules of the form $\frac{(\alpha\supset\beta)\wedge\gamma}{(\alpha\supset\beta)}$.

P-default logic is stronger than Reiter's default logic in that it sanctions more conclusions, as shown by the following two theorems.

Theorem [4] **49:***Let (W,D) be a normal default theory and let (W,D') be the theory where* $\frac{\alpha:\beta}{\beta} \in D$ *iff* $\frac{:(\alpha\supset\beta)}{(\alpha\supset\beta)} \in D'$. *If E is an extension of (W,D) then there is an extension of (W,D'), E', such that* $E \subseteq E'$.

Theorem [4] **50:***Let (W,D) be a semi-normal default theory and let (W,D') be the theory where* $\frac{\alpha:(\beta\wedge\gamma)}{\beta} \in D$ *iff* $\frac{:(\alpha\supset\beta)\wedge\gamma}{(\alpha\supset\beta)} \in D'$. *If E is an extension of (W,D) then there is an extension of (W,D'), E', such that* $E \subseteq E'$.

J-default logic seeks to strengthen the scope of consistency for the justification of a default rule. The weaker notion of consistency in Reiter's default logic results in situations where a default theory (W, D) given by $W = \{\neg b \vee \neg d\}$ and $D = \{\frac{:(a \wedge b)}{a}, \frac{:(c \wedge d)}{c}\}$ sanctions a single extension containing both a and c, or situations where a default theory (W, D) given by $W = \{\}$ and $D = \{\frac{:(a \wedge b)}{a}, \frac{:\neg b}{\neg b}\}$ sanctions a single extension in which b is false. J-default logic requires the set of justifications used in the specification of an extension to be consistent, instead of each individual justification. Thus a *J- extension* is defined to consist of two sets of formulae: E_J, a set of justifications, and E_T, the subset of E_J assumed to be true.

Definition [4] **51:***Let (W,D) be a semi-normal default theory. Define:*
$E_0 = (E_{J0}, E_{T0}) = (Cn(W), Cn(W))$
$E_{i+1} = (E_{J_{i+1}}, E_{T_{i+1}}) = (Cn(E_{J_i} \cup \{\beta \wedge \gamma\}), Cn(E_{T_i} \cup \{\beta\}))$
where

$i \geq 0$,
$\frac{\alpha:(\beta \wedge \gamma)}{\beta} \in D$,
$\alpha \in E_{T_i}$,
$\neg(\beta \wedge \gamma) \notin E_{J_i}$.

Then E is a J-extension for (W,D) iff
$E = (E_J, E_T) = (\bigcup_{i=0}^{\infty} E_{J_i}, \bigcup_{i=0}^{\infty} E_{T_i})$.

The following theorem shows that J-default logic is weaker than Reiter's default logic in the sense that some of the unwarranted conclusions sanctioned by Reiter's default logic can no longer be obtained.

Theorem [4] **52:***Let (W,D) be a semi-normal default theory. If E is an extension of (W,D) obtained using Reiter's definition of an extension, then there is a J-extension $E' = (E'_J, E'_T)$ of (W,D) such that $E'_T \subseteq E$.*

PJ-default logic is a combination of P-default logic and J-default logic in that only prerequisite-free defaults are permitted and extensions are obtained by using the definition of J-extensions given above. PJ-extensions can thus be defined as follows.

Definition [4] **53:***Let (W,D) be a prerequisite-free semi-normal default theory. Define:*
$E_0 = (E_{J_0}, E_{T_0}) = (Cn(W), Cn(W))$
$E_{i+1} = (E_{J_{i+1}}, E_{T_{i+1}}) = (Cn(E_{J_i} \cup \{\beta \wedge \gamma\}), Cn(E_{T_i} \cup \{\beta\}))$
where

$i \geq 0$,
$\frac{:(\beta \wedge \gamma)}{\beta} \in D$,

$\neg(\beta \wedge \gamma) \notin E_{J_i}$.

Then E is a PJ-extension for (W,D) iff
$E = (E_J, E_T) = (\bigcup_{i=0}^{\infty} E_{J_i}, \bigcup_{i=0}^{\infty} E_{T_i})$.

It appears that the PJ-extensions of a PJ-default theory are precisely the belief sets obtained by contracting known invalid situations from the corresponding background belief set. However, since a formal proof has not yet been arrived at, we leave this as a conjecture.

Conjecture 54: *Let (W,D) be a PJ-default theory and let $\Delta_{(W,D)}$ be its translation into a background belief set (i.e., a Δ-translation). Then for each belief set B in the j-contraction of $\Delta_{(W,D)}$, there exists a PJ-extension $E = (E_J, E_T)$ of (W,D) such that $Cn(B) = E_T$, and for every PJ-extension (E_J, E_T) of (W, D), there exists a belief set B such that $E_T = Cn(B)$.*

It is important, at this point, to take a step back and look at why this result is important. Having seen how a belief change approach to default reasoning can avoid some of the counterintuitive results obtained from Reiter's default logic, we can actually establish a direct correspondence between a variant of default logic and our approach. Once again, the advantages of our approach become clear: not only can the unintuitive results of Reiter's default logic be avoided in our approach, we do not need the additional machinery to update default theories in our approach.

4.5 Related Work

Given the strong similarity between the process of identifying maximal consistent subsets of the background belief set and the process of theory formation, the applicability of the THEORIST framework [23] for belief change using possibly inconsistent background belief sets should come as no surprise. Viewing the THEORIST framework as a framework for default reasoning, Brewka [2] has identified the close similarity between a variant of the framework which permits stratifying hypotheses into priority classes (and does away with *facts*, so that everything is potentially refutable) and belief change. Almost identical results have been reported by Nebel [22], together with some results on the computational complexity of the membership problem in revised belief sets. Gärdenfors [10] has examined the relation between *expectation inference operations*, which essentially involves a variant of the maximal consistent subset approach, and belief revision. He presents a generalized notion of nonmonotonic entailment; given a set of formulae called *expectation set* Δ, β is said to be nonmonotonically entailed by α if $\alpha \rightarrow \beta$ is a consequence of every maximal consistent subset of Δ that does not imply $\neg\alpha$ that is selected by some selection function S. The obvious limitation in this approach is that the *expectation set* Δ is required to be consistent. Gärdenfors and Makinson [18] have also proposed a translation of the AGM postulates for belief change [1] into a similar set of postulates for nonmonotonic entailment.

Thus, although default reasoning and belief change have been related, no relationship has been established between reasoning with default logic and belief change. Also, while the existing work seems to suggest that it is possible to view default reasoning as belief change, it was never clear as to why it might actually be better to do so.

Chapter 5

Conclusion

5.1 Contributions

The main contributions of this work are threefold:

- A framework for foundational, logically non-omniscient belief change in the presence of a partial specification of relative epistemic priorities amongst beliefs is defined. More specifically, this approach offers the following advantages:
  - Since the approach is foundational, *reason maintenance* is realized. This also makes the approach computationally realizable; no reference has to be made to infinite deductively closed sets.
  - *Partial closure*, as opposed to full deductive closure, avoids many undesirable beliefs in the closure, as well as unwarranted outcomes of belief change.
  - Previous approaches required that a total ordering on the infinite deductively closed belief sets be specified. This framework works with a partial specification of epistemic priorities.
- Two constructions for contraction operators are provided, which satisfy most requirements for rationality.
- A framework for viewing default reasoning as belief change is defined. More specifically:

- We propose a method of translating a default theory into a belief change framework such that many of the problems associated with reasoning with default logic are avoided.
- We claim, unlike previous authors in this area, that not only is it possible to look at default reasoning as belief change, it may actually be better to do so.
- We suggest that an exact correspondence may exist between our view of default reasoning as belief change and a variant of default logic called PJ-default logic.

5.2 Directions for Future Work

This study should prove to be a useful starting point for studies into belief change in logically non-omniscient reasoners. In this dissertation, logical non-omniscience has only been studied in terms of partial closure. Several other approaches to the problem of logical omniscience exist, all of which merit study from the perspective of belief change.

An obvious direction in which the study of the relationship between belief change and default reasoning could develop is to examine whether this approach to default reasoning satisfies the *cumulativity requirement* (informally, this states that if both x and y are nonmonotonically entailed by W, then x should be nonmonotonically entailed by $W \cup y$). It seems that cumulativity is satisfied, but a formal study remains to be done.

Chapter 6

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

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 a clause $d \in (K^{\sim}_{A} \cap K^{\sim}_{B})$ such that $d \notin K^{\sim}_{A \wedge B}$. We will show that this is not possible.

The possible ways in which some $d \in K$, where K is a partially closed set of clauses, 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 $d \notin K^{\sim}_{A \wedge B}$. □

Lemma 2: *For any p, $ED(T_1, \ldots, T_i, \ldots, T_n) \models p$, iff $T_i \models p$ for $1 \leq i \leq n$.*

Proof: This follows from the definition of $ED(\{T_1, \ldots, T_i, \ldots, T_n\})$.□

Theorem: $K^{-}_A \cap K^{-}_B \subseteq K^{-}_{A \wedge B}$.

Proof: As before, let us assume that there exists a clause $d \in K^{-}_A \cap K^{-}_B$ such that $d \notin K^{-}_{A \wedge B}$. We shall show that this is impossible.

Consider the possible ways in which a clause d may not be in $K^{-}_{A \wedge B}$:

Case 1: d is not a consequence of any element of $C_{max}(\Delta, A \wedge B)$.

Case 2: d is a consequence of some element of $C_{max}(\Delta, A \wedge B)$, but is not a consequence of any element of $E_{max}(C_{max}(\Delta, A \wedge B))$.

Case 3: d is a consequence of some elements of $E_{max}(C_{max}(\Delta, A \wedge B))$, but is not a consequence of $\mathrm{ED}E_{max}(C_{max}(\Delta, A \wedge B))$.

The proof for cases 1 and 2 are similar to those for the operator ' '. For Case 3, there must exist some $Y \in E_{max}(C_{max}(\Delta, A \wedge B))$ such that $Y \not\models d$. By Lemma 1, Y must be in $C_{max}(\Delta, A)$ or $C_{max}(\Delta, B)$. We assume, $Y \in C_{max}(\Delta, A)$. Two possible situations might exist:

- $Y \notin E_{max}(C_{max}(\Delta, A))$. 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.

- $Y \in E_{max}(C_{max}(\Delta, A))$. Then, by Lemma 2, $d \notin K_A^-$. Hence, $d \notin K_A^- \cap K_B^-$. □

Bibliography

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

[2] Gerhard Brewka. Belief revision in a framework for default reasoning. In *Proc. of the Konstanz Workshop on Belief Revision*, pages 206- 222, 1989.

[3] 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.

[4] J. P. Delgrande and W. K. Jackson. Default logic revisited. In *Proc. of the Second International Conference on the Principles of Knowledge Representation and Reasoning*, pages 118-127, 1991.

[5] David Etherington. *Reasoning with Incomplete Information.* Research Notes in Artificial Intelligence. Morgan Kaufmann, Los Altos, CA, 1988.

[6] 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.

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

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

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

[10] Peter Gärdenfors. Nonmonotonic inferences based on expectations: A preliminary report. In *Proc. of the Second International Conference on the Principles of Knowledge Representation and Reasoning*, pages 585-590, 1991.

[11] 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.

[12] W. L. Harper. Rational belief change, Popper functions and counterfactuals. *Synthese*, 30:221-262, 1975.

[13] 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.

[14] H.A. Kautz and B. Selman. Hard problems for simple default logics. In *Proc. of the First International Conference on the Principles of Knowledge Representation and Reasoning*, pages 189-197, 1989.

[15] 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.

[16] Kurt Konolige. On the relation between default theories and autoepistemic logic. In *Proc. of IJCAI-87*, pages 394-401, 1987.

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

[18] David Makinson and Peter Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In *Proc. of the Konstanz Workshop on Belief Revision*, pages 185- 205, 1989.

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

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

[21] 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.

[22] Bernhard Nebel. Belief revision and default reasoning: Syntax-based approaches. In *Proc. of the Second International Conference on Principles of Knowledge Representation and Reasoning*, pages 417-428, 1991.

[23] David Poole, Romas Aleliunas and Randy Goebel. THEORIST: A logical reasoning system for defaults and diagnosis. Technical Report, University of Waterloo, 1985.

[24] 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.

[25] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81-132, 1980.

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

[27] M.A. Williams and N.Y. Foo. Nonmonotonic dynamics of default logics. Technical Report 385. Basser Dept. of Computer Science, University of Sydney, 1990.

[28] 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.

[29] Li Yan Yuan, Jia You and Ashvanee Bissonauth. A retrospective update semantics for epistemically stratified knowledge bases. Forthcoming.