UNIVERSITY OF ALBERTA

A Syntax-Based Approach to Iterated Belief Change

BY

©

Pablo Oscar Hadjinian

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

DEPARTMENT OF COMPUTING SCIENCE

Edmonton, Alberta
Fall 1993

ISBN 0-315-88421-5

UNIVERSITY OF ALBERTA

RELEASE FORM

NAME OF AUTHOR: Pablo Oscar Hadjinian

TITLE OF THESIS: A Syntax-Based Approach to Iterated Belief Change

DEGREE: Master of Science

YEAR THIS DEGREE GRANTED: 1993

(Signed) Pablo Hadjinian

Pablo Oscar Hadjinian
Mendez 2485, Dto. B
Capital Federal (1437)
República Argentina

Date: August 31, 1993

"You can't always get what you want,
but if you try some time you get what you need".
Mick Jagger, Rolling Stones.

UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled **A Syntax-Based Approach to Iterated Belief Change** submitted by Pablo Oscar Hadjinian in partial fulfillment of the requirements for the degree of Master of Science.

. . . .

Dr. Jia-Huai You (Supervisor)

. . . .

Dr. Bernard Linsky (External)

. . . .

Dr. Randy Goebel (Examiner)

Dr. Jonathan Schaeffer (Chair)

Date: . . . .

The author dedicates this dissertation to the brave
ARMENIAN fighters in NAGORNO-KARABAGH.

Abstract

Most existing formalizations treat belief change as a single step process, and ignore several problems that become important when a theory, or belief state, is revised over several steps. This dissertation identifies these problems, and argues for the need to retain all of the multiple possible outcomes of a belief change step, and for a framework in which the effects of a belief change step persist as long as is consistently possible. To demonstrate that such a formalization is indeed possible, we develop a framework which uses the language of PJ-default logic [5] to represent a belief state, and which enables the effects of a belief change step to persist by propagating *belief constraints*. Belief change in this framework maps one belief state to another, where each belief state is a collection of theories given by the set of extensions of the PJ-default theory representing that belief state. Belief constraints do not need to be separately recorded; they are encoded as clearly identifiable components of a PJ-default theory. The framework meets the requirements for iterated belief change that we identify and the belief change operators satisfy a set of properties which are equivalent to the AGM rationality postulates [1].

Acknowledgements

I would like to thank my supervisor Jia-Huai You for his guidance and moral support not only during my research but also while taking courses. I am also pleased to acknowledge Randy Goebel for his moral, intellectual and financial support. Special thanks to my "Guru Friend" Aditya Ghose for his genius, friendship, and listening when I needed it. Together, we spent many sleepless but fruitful nights trying to enrich our knowledge. The rest of the AI group deserves my thanks as well. Alan, Shurjo, Zhong, and others were good friends and always made time in the lab more enjoyable. Thanks also to the other member of my examinig committee, Bernard Linsky, who provided constructive criticism which improved this document. This thesis would not have been possible without the love, patience and support of my fiancé Claudia Negruzzi. Finally, my warm thanks to my friend Jorge Gismondi and my family for their moral and emotional support throughout my M.Sc. program.

Contents

List of Figures

Chapter 1

Introduction

Belief change, the process by which a rational agent acquires new beliefs or retracts previously held ones, is a crucial component of intelligent behavior.

Consider a knowledge base represented by a set of sentences of some logic, say propositional logic. As our perception of the world described by the knowledge base changes, the knowledge base must be modified. It is possible that the new acquired knowledge contradicts the existing information. As we want to keep the knowledge base consistent, we have to revise it. Then, some of the original beliefs need to be retracted. We do not want to give up all the conflicting beliefs. That would be an unnecessary loss of information. The problem is that logical considerations are not enough to tell us which beliefs to give up. A more complicated problem arises when we also consider the logical consequences of the beliefs in the knowledge base. We need to decide which of them to retract.

Example 1 *Suppose your knowledge base contains the following information* *codified in some language $\mathcal{L}$:*

α : *Hripsime goes on flight 704.*

β : *Flight 704 departs from Edmonton at 3 pm.*

γ : *Flight 704 arrives in Toronto at 8 pm.*

If the knowledge base is coupled with a program that has the ability to make logical inferences from the given information, the following fact is derivable from α, β *and* γ *:*

ϵ : *Hripsime will be in Toronto at 8 pm.*

Now, suppose we find that Hripsime happens not to be in Toronto at 8 pm. As we want to incorporate $\neg\epsilon$ *to the knowledge base, the information needs to be revised in order to solve the inconsistency.* □

From the *philosophy* community, pioneering work has been done by Alchourrón, Gärdenfors and Makinson ([1] [20]). They proposed a set of rationality postulates, (known as the AGM postulates), which have been a widely used criterion for evaluating belief change operators. These postulates do not assume any specific representation of the knowledge base. Actually, knowledge bases are modelled as deductively closed sets of sentences in a certain language. Katsuno and Mendelzon [17] have given a model-theoretic characterization of revision schemes that satisfies the AGM postulates in the propositional case.

1.1 Motivation

Most existing formalizations view belief change as a single step process, mapping from one set of beliefs to another [1], [9], [22]. Several new challenges emerge, however, when one considers belief change over several steps. The process of belief change can, in general, have many possible outcomes. With a commitment to producing one unique outcome, most existing formalizations use some form of ordering on the beliefs to select some subset of the possible candidate theories, as in [1], [9], or combine all of the candidate theories to obtain the unique outcome, as in [22]. Both approaches have undesirable consequences. In the first case, the choices determine which beliefs should be held in the revised belief state. This is unintuitive. Although an agent needs to choose one candidate theory in order to have a consistent set of beliefs to act upon, the revised belief state should be determined only by the previous belief state plus the new evidence, and not by the choices encoded in the operator used to perform the revision. In the second case, too much information is lost as a result of combining the mutually incompatible outcomes into a single theory. A related issue is the question of belief persistence. Common sense dictates that the effects of belief change persist as long as there is no evidence to the contrary. However, as a result of not considering iterated belief change, the effects of belief change do not persist over iterated steps in most existing formalizations. Previously contracted beliefs reappear without any justification – a situation we refer to as the problem of spurious beliefs.

1.2 Results Reported in this Dissertation

The results reported in this thesis are based on those in [12] and respresent joint work. Motivated by the pitfalls in the state of the art, we develop an alternative framework which obviates the need for selecting amongst the multiple possible outcomes of a belief change step, does not discard potentially useful information, and formalizes our intuition of belief persistence.
Since choosing some subset of the set of outcomes can cause problems, and since combining multiple outcomes into a single theory can result in too much information being discarded, we argue that an appropriate representation of a belief state is as a collection of theories, as long as this collection can be represented in a compact fashion. Belief change thus becomes a mapping from one collection of theories to another.

Given that each belief state corresponds to a possibly inconsistent and incomplete picture of the world, nonmonotonic formalisms are obvious candidates for representing such a state of affairs. In this dissertation, we use the language of PJ-default logic to represent each belief state. However, any suitable nonmonotonic formalism could be used as the belief representation language using the general principles we describe here. The collection of theories at each belief state is thus the set of extensions of the corresponding PJ-default theory. We implement the persistence of belief change by propagating belief constraints. These are of two types: constraints which enforce necessary belief in some facts and constraints which enforce necessary disbelief in others.

The intuition behind our update operator is that if the new information logically contradicts the original set of beliefs, then its effects are two-fold [26]. A subset of the original belief set is not affected at all; these remain as part of the set of facts of the corresponding PJ-default theory. Another subset of the original belief set is contradicted, or brought into question; these beliefs are demoted to the status of tentative beliefs. This is achieved by removing these beliefs from the set of facts and incorporating them as default rules. A three-fold characterization of beliefs is thus introduced.

For a PJ-default theory (W, D) representing a belief state, the set of facts W corresponds to the *necessary beliefs*, the set of consequents of the default rules in D corresponds to the set of *tentative beliefs*, while the negations of the conjuncts in the "semi-normal" part of the justification of each PJ-default rule correspond to the set of *necessary disbeliefs*.

By addressing pragmatic concerns, such as belief persistence over iterated belief change steps, our framework represents a first step towards bridging the gap between theory and practice in this area. Our approach does not suffer from the problem of spurious beliefs. By factoring out the process of choosing amongst the multiple possible outcomes of a belief change step (typically based on orderings on beliefs [9], [22], [23]), our approach permits the dynamic prioritization of beliefs. In other words, different orderings could be used at different times to choose amongst the multiple possible outcomes without causing inconsistencies or unreasonable results from the belief change process.

Our approach can also be viewed as an account of how a nonmonotonic theory evolves into a more accurate representation of the world through the process of belief change.

1.3 Dissertation Outline

This dissertation is organized as follows. Chapter 2 describes some of the existing formalizations of belief change. In Chapter 3, we argue that after the belief change operation, the resulting state should be independent of any choices made by the belief change operator, and that the effects of belief change should persist. Chapter 4 provides the details of our alternative framework. Chapter 5 discusses the relation of our work to other accounts of belief change. Finally, in Chapter 6, we give our conclusions and directions for future research.

Chapter 2

Belief Change: Existing Frameworks

One of the fundamental aspects of intelligent behavior is the revision of beliefs in the view of new information. In this chapter we review some preliminaries. The *AGM* postulates are also introduced. Based on these rationality postulates, we study how to model *belief states.* Finally, we present belief change operators for belief states modelled as *belief sets* and *belief bases.*

2.1 Fundamentals

When performing a belief revision operation, the objective is to map one belief state to another in the presence of new information, no matter what type the belief change is. Alchourrón, Gärdenfors and Makinson have undertaken a systematic study of the dynamics of belief change, resulting in what is currently known as the AGM framework for belief change (see [1], [21], and [9] for a representative subset of their work). Their taxonomy of belief change operations involves three processes:

Expansion: in which a new belief is added to the belief system together with the logical consequences of the addition (notwithstanding the larger set obtained is consistent).

Contraction: in which an existing belief is retracted without adding any new facts. Also, so that the resulting belief system be closed under logical consequences it may be possible that some other beliefs are also retracted.

Revision: in which a new belief that is inconsistent with the existing beliefs is added, and, in order to maintain consistency in the resulting belief system, some of the old beliefs are retracted. Note that in the rest of this dissertation, we have used the term revision in the general sense, interchangeably with update and belief change– its use in the specific sense of this definition will be clear from the context.

To confront the problem of belief change there exist two general strategies to follow. We can formulate explicit *constructions* of the revision process and formulate *postulates* for such constructions. For us, computer scientists, the ultimate goal is to develop "algorithms" to solve the belief revision problem.

Nevertheless, to be able to know whether an algorithm is successful or not, it is necessary to know what an "appropriate" revision function is. And there is where the *rationality postulates* play the role of evaluating the revision operators. In the following sections we present the AGM rationality postulates which, although idealized, provide a useful model for studying belief revision.

First, in the next section, we investigate how to model a belief state. The rest of this chapter assumes propositional logic as the representation language L.

2.2 Belief State Modelling

Choosing a belief state representation is a crucial step due to the fact that it has a direct impact on the way the belief change operators will be handled. Here, we present two methods, one of which will be used later when we present our framework.

1. One way of modelling a belief state is as a *belief set*, ("knowledge set" in the AGM framework), where, if S is a set of sentences of the language L, A is a belief based on S *if* $A \in Cn(S)$, with $Cn(S) = \{A : S \vdash A\}$ as the set of all logical consequences of S.
   It is important to note that believing in a sentence does not mean that it has any form of justification. This form of representing belief states has been considered in [7] and [9], among others. It has the benefit of handling facts, logical integrity constraints and derivation rules in a uniform way. Of course, we are going to face problems when it comes down to implementing a system, because there are in general infinitely many logical consequences to consider.

2. Other authors ([20], [14], [8]) have argued that some of our beliefs have no independent status and that when we perform revisions we do it on some *finite base* of the belief set. Hence, the suggestion of modelling

the belief state as a belief base, (finite set of sentences). Instead of introducing belief change operations for belief sets, it is supposed that a belief set has an associated belief base and that belief change functions are defined on them. Note, that this way of representing a belief state introduces a more *fine-grained* structure since we can have two bases B_k and C_k such that $Cn(B_k) = Cn(C_k)$ but $B_k \neq C_k$. This means that *syntax is relevant*, that is, the meaning of the knowledge base that results from a belief change operation is not independent of the syntax of the original knowledge base.

For example, consider the following scenario from Hansson [14].

> ...suppose that on a public holiday you are standing in the street of a town that has two hamburger restaurants. Let us consider the subset of your belief set that represents your beliefs about whether or not each of these two restaurants is open.
>
> When you meet me, eating a hamburger, you draw the conclusion that at least one of the restaurants is open $(A \vee B)$. Further seeing from a distance that one of the two restaurants has its lights on, you believe that this particular restaurant is open (A). This situation can be represented by the set of beliefs $\{A, A \vee B\}$.
>
> When you have reached the restaurant, however, you find a sign saying that it is closed all day. The lights are only turned on for the purpose of cleaning. You now have to include the

negation of A, i.e. $\neg A$ into you belief set. The revision of $\{A, A \vee B\}$ to include $\neg A$ should still include $A \vee B$, since you still have reasons to believe that one of the two restaurants is open.

In contrast, suppose you had not met me or anyone else eating a hamburger. Then your only clue would have been the lights from the restaurant. The original belief system in this case can be represented by the set $\{A\}$. After finding out that this restaurant was closed, the resulting set should not contain $A \vee B$, since in this case you have no reason to believe that one of the restaurants is open.

This example illustrates the need to differentiate in some epistemic contexts between the set $\{A\}$ and the set $\{A, A \vee B\}$. Note that the closure of $\{A\}$ and $\{A, A \vee B\}$ are identical. Therefore, if we assume our knowledge base to be closed under logical consequence, the distinction between $\{A\}$ and $\{A, A \vee B\}$ also cannot be made.

"If we were to implement a system it seems that belief bases would be preferable as intuitively they would be easier to handle. This is one of the reasons why we use the second type of belief state modelling for the alternative framework presented in Chapter 4".

Another question that arises when modelling belief states is: *Should we keep track of the justifications for each belief ?*

In one extreme we have the *coherence approach*, where the justifications for each belief are not recorded. The focus is instead on the *logical* structure of the beliefs, as long as the beliefs cohere in the present state. The objectives of the coherence theory are to maintain *consistency* and to make *minimal* changes. Belief sets, as introduced above, fall into this category.

On the other extreme we have the *foundational approach*, where every belief is self-evident or is supported by a structure of self-evident beliefs. Beliefs with no justification are not accepted. Justifications are *prima facie* defeasible. The foundations theory rejects any principle of *conservatism*. Also, justifications cannot legitimately be *circular* and an *infinite* chain of justifications is also disallowed. Belief bases where all beliefs have independent justification give rise to a moderate form of foundational theory. The so-called truth maintenance systems, ([6], [4]), maintain explicit records of justifications.

2.3 The AGM Rationality Postulates

In general, both contraction and revision have multiple possible consistent outcomes. The AGM framework defines various methods of combining these outcomes to obtain a single consistent resultant belief state. The AGM postulates provide a criterion for producing the result of a belief revision process, based on the input of a new belief and a deductively closed set of existing beliefs. In the AGM framework, the representation language is propositional logic and the deductively closed sets denoting beliefs are referred to as *knowledge sets.*

We present below the postulates for contraction. 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^-, revision by K_A^*, and expansion by K_A^+. In 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 $A \notin K$, then $K_A^- = K$.

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

(5-) $K \subseteq (K_A^-)_A^+$.

(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. Postulates (7-) and (8-) together enforce the requirement of *informational economy* - the requirement that updating beliefs should result in as few changes to the existing body of beliefs as possible.

The operations of contraction and revision can be defined in terms of each other. Let K_A^* and K_A^+ denote, respectively, the revision and expansion of K with A. The *Levi identity* [19] defines revision in terms of contraction as shown below:

$$K_A^* = (K_{\neg A}^-)_A^+$$

The *Harper identity* [16] defines contraction in terms of revision as shown below:

$$K_A^- = K_{\neg A}^* \cap K$$

The postulates for contraction and revision are thus equivalent. If a given contraction operator satisfies the postulates for contraction, the corresponding revision operator (defined in terms of the Levi identity) must satisfy the postulates for contraction. Based on the AGM rationality criteria we study some contraction operators in the next section. We choose contraction because it is, in some aspects, easier to study and more intuitive to view revision as a composition of contraction and expansion.

2.4 Contraction Operators

In this section we introduce an explicit modelling of contraction functions for *belief sets* and *belief bases.* These contraction operators will be based on the rationality postulates studied before.

2.4.1 Belief Set Contraction

The problem to be addressed is how to define the contraction K_x^- where K is a belief set and x is a propositional sentence. As these operators are based on the idea of identifying maximal consistent subsets, let the *removal* of x from A, denoted by $A \downarrow x$, be defined as follows:

$$A \downarrow x = \{B \subseteq A \mid B \nvdash x, \forall C : B \subset C \subseteq A \Rightarrow C \vdash x\}$$

That is, the removal of x from A results in the set of maximal consistent subsets of the deductively closed set K that do not imply the sentence x.

Maxichoice contraction [1] is a first approach towards modeling a contraction function where there exists a selection function γ that picks out an element of $K \downarrow x$ for any K and any x whenever $K \downarrow x$ is nonempty. We then define K_x^{-m} as follows:

$$\textit{(Maxichoice)} \quad K_x^{-m} = \begin{cases} \gamma(K \downarrow x) & \text{if} \not\models x \\ K & \text{otherwise} \end{cases}$$

Maxichoice satisfies postulates $(1-)-(6-)$ but in general produces contractions that are *too large*, i.e., a belief sets that contain either y or $\neg y$ for every y, even though they were not part of the original belief set. Maxichoice contraction functions create maximal belief sets. We then define K_x^{-f} as follows:

Full meet contraction [1] is a second tentative solution where the propositions retained after the contraction operation are those common to all of the maximal consistent subsets in $K \downarrow x$.

$$\textit{(Meet)} \quad K_x^{-f} = \begin{cases} \bigcap(K \downarrow x) & \text{if } \not\models x \\ K & \text{otherwise} \end{cases}$$

It is easy to show that any full meet contraction function satisfies (1-) – (6-). The drawback is that in general the resulting contracted belief sets are far *too small.*

Partial meet contraction [1] is a representative belief revision operator that satisfies all of the AGM rationality postulates. Let S be a *selection function* that selects a nonempty subset of $K \downarrow x$ (provided $K \downarrow x$ is nonempty, $\emptyset$ otherwise). The partial meet contraction operator $-_p$ is defined as follows:

$$\textit{(Partial Meet)} \quad K_x^{-p} = \begin{cases} \bigcap S(K \downarrow x) & \text{if } \not\models x \\ K & \text{otherwise} \end{cases}$$

Let $M(K)$ stand for the family of all the sets $K \downarrow x$, where x is any proposition in K that is not logically valid. Let $\leq$ be a relation defined on $M(K)$. Let:

$$S(K \downarrow x) = \{K' \in K \downarrow x \mid K'' \leq K' \text{ for all } K'' \in K \downarrow x\}$$

If S is defined in this manner, S is said to be *relational* over K, and any partial meet contraction operator defined using some relational selection function is called a *relational* partial meet contraction operator. If the relation $\leq$ is transitive, the partial meet contraction operator is said to be ***transitively relational***. Any transitively relational partial meet contraction (revision) operator satisfies all the AGM postulates for contraction (revision) [9].

Hence, we have found a way of connecting the rationality postulates with a general way of modelling contraction functions. The disadvantage is that the computational cost involved in determining the maximal consistent subsets is overwhelming. This is one of the reasons, as we will show when we present our framework, for choosing to define belief revision operators on finite set of sentences as opposed to infinite ones.

Epistemic Entrenchment is a contraction operator which uses a special class of total orderings defined on the entire language, to decide which beliefs to retain, is shown to satisfy all eight of the AGM postulates for contraction, and hence to be equivalent to partial meet contraction (given the result from [1] showing that any contraction operator that satisfies (1-) through (8-) is a partial meet contraction operator and vice versa).

2.4.2 Belief Base Contraction

Motivations to define belief change on belief bases are twofold. *First*, operators defined on belief bases are computationally viable (they do not have to operate on infinite sets). *Second*, belief change operations on belief bases permit reason maintenance, while those on belief sets do not. Intuitively, an agent stores the explicit beliefs in a belief base. Some authors demand them to be finite ([22], [3]), and some do not ([8], [23]). If bases are allowed to be infinite, it is supposed that there exists some other base that generates them.

So, each belief state K is associated with a base B such that $K = Cn(B)$. How the changes are to be performed is driven by the structure of the base. Nevertheless, many approaches do not identify a belief base after the contraction making *iterated belief change* impossible. The main idea when revising belief bases is that *syntax* is important. This is not always possible if belief change operations are performed on an infinite set of sentences. So, it looks like a model theoretic characterization is not appropriate.

For example, belief bases $B = \{A, B\}$ and $B' = \{A \wedge B\}$ generate the same belief set K. But, contracting A from them give different results.
As it is shown later, we propose an approach that only appeals to finite bases. The constructive models introduced in the previous section have also been thoroughly studied in the context of belief bases. Here, as it is not our main objective, we are going to present special cases of these operators instead of studying the theory of *base contraction* in full generality. To study belief change over belief bases with full generality see [14], [15].

For example, Nebel [22] starts defining the base contraction operator $\sim$, analogous to *full meet contraction*, as follows:

$$B \sim x = \begin{cases} \bigvee_{C \in (B \downarrow x)} C & \text{if } \not\vdash x \\ B & \text{otherwise} \end{cases}$$

where B is a belief base. Then, by taking the initial and final belief sets of the AGM framework to be the deductive closures of B and $B \sim x$ respectively, Nebel shows that the $\sim$ operator satisfies postulates (1-)-(4-) and (6-), but not (5-), (7-) and (8-).
A somewhat modified operator $\simeq$ defined as:

$$B \simeq x = \begin{cases} (\bigvee_{C \in (B \downarrow x)} C) \wedge (B \vee \neg x) & \text{if } \not\vdash x \\ B & \text{otherwise} \end{cases}$$

satisfies the *recovery postulate* (contraction postulate (5-)). To prove that the (7-) postulate is as well satisfied, Nebel shows that his base contraction operator $\simeq$ yields a base whose deductive closure is equal to the result of partial meet contraction of the deductive closure of the original belief base, using a specific *relational* selection function that he defines (this follows from a result in [9] which states that any relational partial meet contraction satisfies postulates (1-) through (7-)). Postulate (8-), however, is not satisfied. Finally, following Nebel, ([22]), we present a special case of *partial meet contraction*, called *prioritized base contraction*, subscribed also by [7] and [3] among others. A prioritized belief base is a pair $< B, \preceq >$ where B is a belief base and $\preceq$ is a *complete preorder* with maximal elements over B. This prioritization can be represented by a sequence $< B_1, B_2, \ldots, B_n >$ where the $B_i's$ are the equivalence classes of B under $\approx$, and $x \preceq y$ for $x \in B_i$ and $y \in B_j$ iff $i \leq j$. If the preorder is complete, $\preceq$ is a linear order on the degrees of epistemic relevance $B_i's$. Nebel calls this relation *epistemic relevance.*
If $< B, \preceq >$ is a prioritized belief base; then B' is a *preferred element* of $(B \downarrow A)$ iff for every B'' and every i s.t. $B' \cap (B_i \cup B_{i+1} \cup \ldots \cup B_n) \subseteq B'' \cap (B_i \cup B_{i+1} \cup \ldots \cup B_n)$ it holds that $B'' \vdash A$. So, B' is a maximal consistent subset from B at each priority level, s.t. A is not implied. To contract A from the theory K with base B, we take $\bigcap(Cn(B' :\ B'$ is a preferred element of $(B \downarrow A)$.

A selection function S that chooses elements of the belief base in this manner can be defined as relational. Let $K = Cn(B)$ and define the relation $\ll$ over $(K \downarrow A)$ as follows,

$$K' \ll K'' \ iff \ K' \cap B \subset K'' \cap B,$$

where S selects those B' which are maximal in $(B \downarrow A)$ under the following preference relation between arbitrary sets of sentences:

$G \ll G'$ iff there is an i s.t. $G \cap B_i \subset G' \cap B_i$ and for every $j > i$, $G' \cap B_i = G' \cap B_j$

Actually, because of the problems with the recovery postulate, he presents this result for revisions rather than for contractions. Nebel also proves that his epistemic relevance revisions in general satisfy the first seven AGM revision postulates but not the eighth.

Chapter 3

Iterated Belief Change

Few existing studies consider iterated belief change. Instead, their focus is on formalizing the single step mapping from one belief state to another. In this chapter we identify the *problem of spurious beliefs* and we argue with full generality about the conditions that a *belief system* should observe so that *iterative* belief change is possible. The material presented in this chapter also appears in [12].

3.1 Principles for Iterated Belief Change

As a result of considering belief change as a single step process, in most existing formalizations, contracted beliefs often re-appear after several belief change steps, not because of the new evidence accumulated since the contraction warrants such a belief, but as a consequence of certain choices made by the belief change operator based on priorities on beliefs or epistemic entrenchment. We call this, the *problem of spurious beliefs.*

We argue that every belief change operator should satisfy the following *Principle of Irrelevance of Choice*:

> **(I1)** The belief state after one or more belief change steps is uniquely determined by the initial belief state and the evidence accumulated over the belief change steps, independent of any prioritization on the beliefs.

Let $F(B_0, E) = f(f(\ldots f(B_0, e_1))\ldots), e_{n-1}), e_n)$ denote the composition of n belief change steps starting with belief state B_0. Here, f denotes the unit step belief change operator. Without making any commitment to a representation language for evidence, let e_i denote the evidence obtained (from the world) in belief change step i, such that, set-theoretically, the accumulated evidence over n steps is denoted by $E = e_1 \cup e_2 \cup \ldots \cup e_n$. Let $f(B_{n-1}, e_n) = B_n$. Then, the principle of irrelevance of choice states that $F(B_0, E) = B_n$. The point to be noted is that the mapping from B_0 to B_n is determined uniquely and entirely by B_0 and the accumulated evidence E, and is independent of extraneous information, such as prioritization on the beliefs. Intuitively, the justification for this postulate is fairly obvious. A priority relation on beliefs does not represent grades of truth (our argument would not hold if that were the case); it merely represents information that could be brought to bear when a choice has to be made. Belief change should be driven entirely by belief inputs from the environment, and should not be influenced by choices encoded within the belief change operator, when these are not based on the truth status of the beliefs.

A related issue is the question of *persistence* of belief change. We claim that belief change operators must satisfy the following *persistence postulate*:

> **(I2)** The effects of belief change persist until there is new evidence to indicate otherwise.

A contracted belief should not re-appear in a belief set unless there is new evidence accumulated since the contraction that justifies this. Similarly, a newly added belief should remain in the belief set until new evidence obtained since the addition warrants its removal. Operators that do not satisfy this postulate suffer from the problem of spurious beliefs. Informally, a spurious belief is one which reappears in a belief set one or more steps after it has been contracted even though the evidence accumulated since the contraction step does not warrant such a belief. The examples in this section involve an assumption of static worlds. That is, we focus on *revisions* (as opposed to updates in dynamic worlds) in the sense of the distinction identified by Katsuno and Mendelzon in [18].

Example: Let the initial belief state only consist of the belief that it is snowing. Thus $K_0 = \{snowing\}$. We are going to analyze two cases. *First*, let the new evidence indicate that the belief that it is *cold* needs to be retracted, possibly because a warm weather system has moved in. Since *cold* is not a consequence of K_0, no change needs is made to the belief state. Thus $K_1 = \{snowing\}$. Let the next item of information supplied to this agent indicate that *snowing* $\rightarrow$ *cold*. If the new belief state is $K_2 = \{snowing, snowing \rightarrow cold\}$, then the belief that it is cold will be one of its consequences. We claim that the renewed belief in *cold* is inappropriate, since there is nothing in the new information obtained since *cold* was contracted that indicates that it should be believed in again.

Second, let us assume that as a result of a warm weather system we revise the belief state K_0 with the new belief $\neg cold$. As *cold* is not a consequence of K_0, the belief change operation is an expansion. So, $K_1 = \{snowing, \neg cold\}$. If, as before, the new piece of evidence indicates that $snowing \rightarrow cold$ is accepted then K_1 has to be revised to incorporate the new belief. There are two possible outcomes as a result of the revision: $\{snowing, snowing \rightarrow cold\}$ and $\{snowing \rightarrow cold, \neg cold\}$. Assume that the new belief state is $K_2 = \{snowing, snowing \rightarrow cold\}$, where *cold* is a consequence of K_2. As before, we argue that there is not enough new evidence that justifies believing *cold*.

□

Note that in the first part of the example the problem of spurious beliefs is the result of not propagating disbeliefs while, in the second part, it happens as a result of preferring one of the outcomes of the revision over the others. The intuition of what constitutes a spurious belief can be formalized as follows:

Definition 1 *Let $K_0 = K_x^-$ denote the belief state obtained as a result of contracting the belief x from the belief state K. Let $K_0' = \{\}$ denote a belief state containing no beliefs. Let both K_0 and K_0' be put through the same sequence of n belief change steps using operator f, such that the resultant belief states are K_n and K_n' respectively. We say that operator f suffers from the problem of spurious beliefs if $x \in K_n$ and $x \notin K_n'$.*

Current formalizations satisfy the persistence postulate over a single step, but this result does not carry through over iterated belief change steps. As the examples below demonstrate, both the contraction operators described in the previous section suffer from the problem of spurious beliefs.

Example: Consider the AGM formalization of belief change. Let the initial belief set be $K = Cn\{a, a \rightarrow b\}$. Let new evidence from the environment necessitate the contraction of b. Let $S(K \downarrow b) = \{Cn\{a \rightarrow b\}\}$. Thus, $K' = K_b^{-p} = Cn\{a \rightarrow b\}$. Let the next belief change step be the expansion of this belief set with a. Trivially, the expanded belief set $K_a'^{+} = Cn\{a, a \rightarrow b\}$. Notice that $b \in K_a'^{+}$. Notice also that the only new evidence obtained since the contraction of b is confirmation for the belief a already held at the time b was contracted. It is easy to see that there is nothing in the new evidence to warrant believing in b again, and that b is therefore a spurious belief. Thus, partial meet contraction violates the principle of irrelevance of choice, since the belief b reappears only as a consequence of the selection function's choice of $Cn\{a \rightarrow b\}$ over other elements of $K \downarrow b$ (if $Cn\{a\}$ had been chosen, the belief b would not reappear). Since belief change (the contraction of b) does not persist inspite of there being no evidence to the contrary, the persistence postulate is also violated. □

Example: Consider Nebel's formalization of belief base revision [22]. Let the initial belief base be $B_0 = \{a \rightarrow b\}$. The base contraction of b from B_0 yields the belief base $B_1 = \{a \rightarrow b, b \rightarrow B_0\}$, where we take B_0 to denote the conjunction of its elements. Later expansion with a yields the base $B_2 = \{a, a \rightarrow b, b \rightarrow B_0\}$. Notice that $b \in Cn(B_2)$. As with the previous example, there is nothing in the new evidence obtained since the contraction of b to warrant believing in b again. Clearly, Nebel's formalization of belief change also suffers from the problem of spurious beliefs. In this case, the problem stems not from choice, but from not keeping an explicit record of belief change steps. □

In the next chapter, we present an alternative formalization of belief change to address these concerns.

Chapter 4

An Alternative Framework for Belief Revision

Having in mind our discussion of the previous chapter, we present an alternative framework for belief revision that uses the language of PJ-default logic to represent a a belief state. First, the motivations and some useful definitions are given. Then, the actual framework is presented, together with an analysis of its advantages and characteristics. Note that most of the text in this chapter appears in [12].

4.1 Motivation

Our account of belief change is motivated by the following observations:

- A belief state is best represented as a collection of theories. There are several different aspects of the belief change process that suggest this.

The principle of irrelevance of choice requires that belief change should be independent of any choice made amongst the possible outcomes of a belief change step. Given that choosing amongst the outcomes may be inappropriate, or even impossible, it could be argued that all outcomes should be retained, provided that there exists a compact and elegant way of representing these multiple possible outcomes. Such representation languages exist; nonmonotonic formalisms are immediate candidates for compactly representing the possibly inconsistent and incomplete picture of the world that each belief state corresponds to. Justifications for such an approach also exist in cognitive science. It appears that humans are able to posit multiple mutually inconsistent explanations for the same observation, suggesting that they are able to draw on many different theories in the same belief state.

- The persistence postulate suggests that an explicit record of belief change steps be maintained. We argue for the maintenance of such a record by propagating a set of *belief constraints*. Two kinds of constraints need to be maintained: a set of formulas which must necessarily be believed and a set of formulas which must necessarily be disbelieved.

- Beliefs originally held to be true can become tentative as a result of belief change. This can happen if a belief is contradicted or brought into question (the intuition is that a belief becomes questionable if it is not in every possible outcome of the belief change step) by the new evidence. In syntactically-oriented nonmonotonic formalisms, this can be viewed as a process of demotion from a fact to a default.

4.2 PJ-Default Logic

Commonsense reasoning is in general *not monotonic.* We generally draw conclusions which later are given up in the light of new information. The typical example is the flying ability of birds. If we know that Tweety is a bird, we tend to draw the conclusion that it flies, since birds typically fly. If later, we are given the information that Tweety is a penguin we are going to withdraw the former conclusion that Tweety flies, without withdrawing any of our former premises. Such forms of reasoning which allow additional information to invalidate old conclusions are called *nonmonotonic.* Several are the formalisms proposed to deal with nonmonotonic reasoning (for a survey see [13]). Note that the conclusion that Tweety flies was reached by default, as originally there was no evidence of Tweety being a penguin. *Default Logic* [25], is a consistency based nonmonotonic formalism aimed at providing a formal framework for default reasoning. To accomplish this, default logic extends first order logic by introducing a type of nonmonotonic inference rule called defaults.

Definition 2 *A default δ is any configuration of the form* $\frac{\alpha:\beta_1,\ldots,\beta_n}{\gamma}$
where

$-\alpha, \beta_1, \ldots, \beta_n, \gamma$ *are sentences*

$-\alpha$ *is called the* **prerequisite** *of the default* δ

$-\beta_1, \ldots \beta_n$ *is called the* **justification** *of the default* δ

$-\gamma$ *is called the* **consequent** *of the default* δ.

Example: Hence, that birds typically fly can be expressed by the default

$$\frac{Bird(x) : Fly(x)}{Fly(x)}$$

which is to be interpreted as:

"If x is a bird and if it is consistent that x can fly then infer that x can fly".□

Definition 3 *A default theory, is an ordered pair,* (W, D)*, consisting of a set of defaults,* D*, and a set of first–order formulae* W*.*

PJ-default logic is a variant of default logic in which default rules are restricted to be prerequisite-free and semi-normal, which means that a PJ-default rule is of the form $\frac{:\beta}{\gamma}$ such that $\beta \models \gamma$). PJ-extensions can be defined as follows [5]:

Definition 4 *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

$E = (E_J, E_T) = (\bigcup_{i=0}^{\infty} E_{J_i}, \bigcup_{i=0}^{\infty} E_{T_i}).$

PJ-default logic addresses several classes of problems with Reiter's default logic, including cases where default logic is too weak, preventing the derivation of "reasonable " conclusions (the disjunctive default problem, for instance), and cases where default logic is too strong, permitting the derivation of unwanted conclusions. Additionally, PJ-default logic has several attractive properties, such as the guaranteed existence of extensions. Ghose and Goebel [11] have earlier defined a belief change framework in which a belief state is represented as a potentially inconsistent set of sentences, together with a partial order on these sentences. An operator is defined that identifies consistent subsets of this set of sentences, that respect the partial order as well as some set of disbelief constraints. A translation from PJ-default theories to this framework is defined such that the process of identifying PJ-default extensions is shown to be equivalent to the process of identifying consistent subsets of sentences using the operator mentioned above, with a partial order which assigns higher priority to sentences obtained from W (given a PJ-default theory (W, D)) than sentences obtained from D and with the set of disbelief constraints consisting of the conjunction of the negations of the justifications of each PJ-default rule.

4.3 Basic Definitions

Our choice of PJ-default theories to represent belief states has the additional advantage that belief constraints do not need to be represented separately but can be incorporated as part of a PJ-default theory. The set of facts W of a PJ-default theory corresponds to the beliefs that the agent is constrained to necessarily hold, since W will be contained in every consistent belief set corresponding to that belief state (i.e., every PJ-extension). Since every PJ-default rule is of the form $\frac{:\beta \wedge \gamma}{\beta}$, $\neg\gamma$ corresponds to the set of beliefs that the agent is constrained to necessarily disbelieve. For example, if told that a certain belief has been contradicted or made questionable, the process of demoting that belief to the status of a tentative belief involves removing a formula from W and adding a new PJ-default rule containing this formula as its consequent to D. Discredited beliefs are thus never totally discarded in our framework in anticipation of future situations in which these beliefs could be consistently held again. Belief change in our framework thus involves mapping one PJ-default theory to another. The possible consistent belief sets that may be held in a given belief state corresponds to the extensions of the PJ-default theory representing that belief state. The process is depicted graphically in Figure 1.

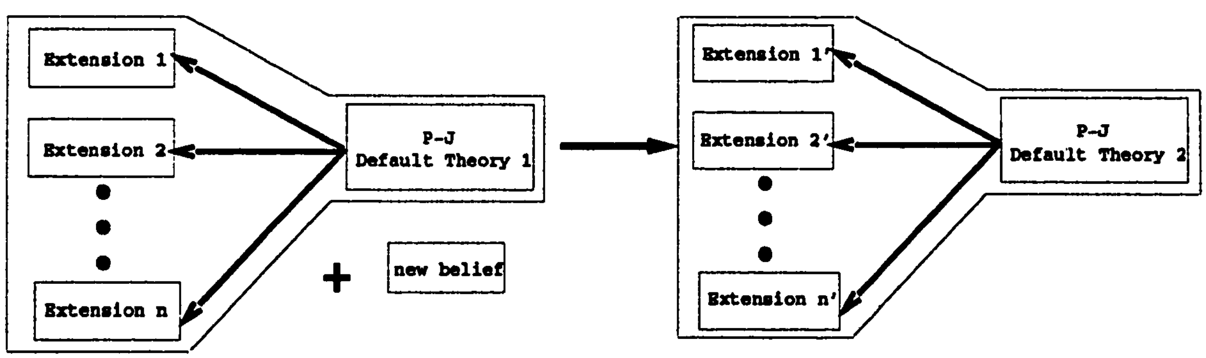

Figure 4.1: Belief State Transformation

Definition 5 *The* `necessary belief set` BC_{belief}, *for a given belief state, is a set of sentences such that for every* Th_i, *such that* Th_i *is a theory that may be consistently held in that belief state, and* $\forall x : x \in BC_{belief}$, $Th_i \models x$.

Definition 6 *The* `necessary disbelief set` $BC_{disbelief}$, *for a given belief state, is a set of sentences such that for every* Th_i, *such that* Th_i *is a theory that may be consistently held in that belief state, and* $\forall x : x \in BC_{disbelief}$, $Th_i \not\models \neg x$.

Thus, while the necessary belief set consists of beliefs that must be held, the necessary disbelief set consists of the *negations* of the beliefs that must not be held. We require that every element of BC_{belief} and $BC_{disbelief}$ be represented as a clausal theory. The reason for this is that we want disbelief constraints not to be treated as syntactinc units which are enabled or disabled as whole units, but in terms of the individual components. We also assume that some syntactic convention exists for distinguishing the elements of BC_{belief} from those of $BC_{disbelief}$. Here we do as follows: a constraint enforcing necessary belief in $a \rightarrow b$ shall be written simple as $\{\{\neg a, b\}\}$ while a constraint enforcing necessary disbelief in $a \vee b$ shall be written as $-\{\{\neg a\}, \{\neg b\}\}$.

Definition 7 *The set of* `belief constraints` $BC = BC_{belief} \cup BC_{disbelief}$ *at any given belief state consists of the currently applicable belief constraints.*

Definition 8 *The* `constraint prioritization` *for a given set of belief constraints* BC *is a total order* $\prec$ *defined on the elements of* BC.

The intuition behind defining a prioritization of the belief constraints is to reflect the temporal order in which the belief change steps were performed, and hence the order in which each belief constraint was obtained. The order in which evidence is obtained is crucial in determining the outcome of a series of belief change steps. For instance, if one were to obtain evidence that a is to be believed, followed by further evidence that a is to be believed, followed by evidence that a is to be disbelieved, one would reach a belief state in which a is not believed. Now, consider the results of permuting this order of belief change steps, such that the evidence first requires that a be believed, then requires that a be disbelieved, and finally that a be believed again. In this case, a would be believed in the final belief state.

Note that by requiring a total order on the set of belief constraints, we are not reverting back to what we originally set out to change: a formalization of belief change that involves making choices based on priorities. *First*, this total ordering has to be defined on a finite set of belief constraints whose size can typically be expected to be much smaller than that of the total set of beliefs that may be potentially held. As we shall see later in this chapter, the size of the set of belief constraints does not grow monotonically with the new evidence. *Second*, this ordering does not determine which outcome is selected; it merely determines how far the effects of a belief change step should persist. Note, that existing formalizations already use such an ordering by requiring that every belief change operation succeed (AGM contraction postulate (2-)), although this is not obvious on account of their formalization being restricted to a single step. The new belief is thus implicitly assigned a higher priority than all existing beliefs. *Third*, the total ordering that we require is automatically generated by the order of belief change steps. On the other hand, information on the prioritization of beliefs (such as epistemic entrenchment in the AGM framework) can be difficult, and even impossible, to obtain.

4.4 Belief Change Operation: a Two Step Process

In this section we present in detail how the actual belief revision operation is performed. Let BC_{old} *and* BC_{new} be the current and new set of belief constraints respectively, let Ω be the new piece of evidence, and let (W, D) and (W', D') be the current and updated PJ-default theories, that result from the belief change operation, respectively. The two step belief change process is as follows:

1. **Belief Constraints Update**
   $BC_{old} \times \Omega \rightarrow BC_{new}$
2. **$\Gamma-$PJ-default theory update**
   $(W, D) \times BC_{New} \rightarrow (W', D')$

where (W, D) and (W', D') are PJ-default theories, $BC_{old}\ and BC_{new}$ stand for sets of belief constraints, before and after having performed the belief change operations, and Ω denotes the new belief constraint which could either be a new belief to be added, (necessary believe), or a belief to be contracted (necessary disbelieve).

4.4.1 Updating the Set of Belief Constraints

In our model, belief change involves the incorporation of a new belief constraint Ω. If Ω is a necessary belief constraint, then the formula Ω must be a consequence of every possible outcome of the belief change operation. If Ω is a necessary disbelief constraint, then every outcome of the belief change operation must be consistent with Ω, (since Ω represents the negation of the belief that cannot be held). The set of constraints needs to be updated to incorporate the new belief constraint, specially if the new belief constraint is incompatible with the existing set of constraints. We define this notion of incompatibility formally below.

Definition 9 *A belief constraint Ω is said to be* `compatible` *with a set of belief constraints BC (we write $\Omega \cup BC$ is compatible) if and only if the following conditions hold:*

1. *For every x_i such that $x_i \in BC_{belief}$, $\bigwedge_i x_i \not\models \neg\Omega$.*
2. *For every x_i such that $x_i \in BC_{disbelief}$, $\bigwedge_i \neg x_i \not\models \neg\Omega$.*

We also need to be able to identify subsets of an individual belief constraint (viewing each constraint as a set of clauses) that are compatible with a set of belief constraints. The operator $\uparrow$ that identifies such subsets is defined below.

Definition 10 *Let bc be a belief constraint and BC be a set of belief constraints. Then*

$$bc \uparrow BC = \{x \subseteq bc \mid (x \cup BC \text{ is compatible}) \wedge (\forall x' \text{ such that } x \subset x' \subseteq bc \text{ and } x' \cup BC \text{ is incompatible}\}$$

Updating the set of belief constraints involves taking the new constraint and adding as many of the previous constraints (or parts of previous constraints) as can be compatibly added, starting with the most recent constraint and working backwards to the earlier ones. We shall assume the following notational convention for referring to individual belief constraints: $\forall bc_i \in BC, bc_j \prec bc_k \leftrightarrow k < j$. In other words, a more recent constraint will have a smaller subscript. In the following, we also assume that Θ represents Ω in the appropriate syntactic form, i.e., if Ω is a necessary disbelief constraint, then $\Theta = -\Omega$, while $\Theta = \Omega$ otherwise.

$$\begin{aligned} BC_{new} = \; & \Theta \cup \{Y \subseteq BC_{old} \mid Y = \textstyle\bigcup_{i \geq 1} Y_i, \\ & \forall i \geq 1(((Y_i = \{bc_i\}) \wedge (bc_i \cup \Theta \cup (\textstyle\bigcup_{j=1}^{i-1} Y_j) \text{ is compatible})) \\ & \vee((Y_i = \textstyle\bigcap(bc_i \uparrow (\bigcup_{j=1}^{i-1} Y_j \cup \Theta)) \wedge (bc_i \cup \Theta \cup (\textstyle\bigcup_{j=1}^{i-1} Y_j) \text{ is incompatible})))\} \end{aligned}$$

Example: Let $BC_{old} = \{\{\{a\}\}, -\{\{\neg b\}\}, \{\{\neg a, b\}\}\}$ and the temporal ordering of these constraints, from most recent to least recent, be $\{\{a\}\} \succ \{\{\neg a, b\}\} \succ -\{\{\neg b\}\}$. Then $BC_{new} = \{\{\{a\}\}, \{\{\neg a, b\}\}\}$. Since the two most recent necessary belief constraints $\{\{a\}\}$ and $\{\{\neg a, b\}\}$ are together not compatible with an older necessary disbelief constraint $-\{\{\neg b\}\}$, the older constraint is discarded. If the temporal ordering had been $\{\{a\}\} \succ -\{\{\neg b\}\} \succ \{\{\neg a, b\}\}$, the updated set of belief constraints would have been $BC_{new} = \{\{\{a\}\}, -\{\{\neg b\}\}\}$.□

Henceforth, for our framework, the operation of contraction, denoted by $-$, will be treated as adding a necessary disbelief constraint, while the operations of expansion and revision, denoted by $+$ and $*$ respectively, will be treated as adding a new necessary belief constraint. The operation of "undoing a contraction" (notice that this is not the same as revision with the contracted belief) has received very little attention in existing frameworks for belief change. Our framework permits the undoing of a contraction operation; it involves the removal of the relevant necessary disbelief constraint. It is not possible to define such an operation in the AGM framework, or using Nebel's formalization of belief change.

4.4.2 Updating the PJ-default Theory

The PJ-default theory (W', D') representing the new belief state can be obtained given the new belief constraints BC_{new} and the previous belief state (W, D). The important steps in this transformation are:

- Beliefs whose status become *tentative* as a result of the belief change step are demoted to the level of defaults.
- *Necessary disbeliefs* are enforced by placing the negations of these beliefs in the justifications of PJ-default rules.
- *Necessary beliefs* are enforced by making such beliefs facts, as opposed to default rules.

The new PJ-default theory (W', D') is computed as follows:

$$
\begin{aligned}
W' &= BC_{belief_{new}} \\
D' &= \{\frac{:\delta_i \wedge (\bigwedge \gamma_j)}{\delta_i} \mid \delta_i \in (W - W') \textit{ and } \gamma_j \in BC_{disbelief_{new}}\} \cup \\
&\quad \{\frac{:\beta_i \wedge (\bigwedge \gamma_j)}{\beta_i} \mid \frac{:\beta_i \wedge \phi_i}{\beta_i} \in D \textit{ and } \gamma_j \in BC_{disbelief_{new}}\}
\end{aligned}
$$

Definition 11 *A* Γ*-PJ-default theory* (W, D) *is one in which every default* $d \in D$ *is of the form* $\frac{:\beta_i \wedge \Gamma}{\beta_i}$.

Notice that every PJ-default theory representing a belief state in our framework is a Γ-PJ-default theory, i.e., the "semi-normal" part of the justification of every default rule is identical. Notice also that belief constraints do not need to be explicitly represented but are clearly identifiable parts of a Γ-PJ-default theory; the set of necessary belief constraints corresponds to W while the set of necessary disbelief constraints corresponds to the set consisting of the negation of each conjunct in the "semi-normal" part (the Γ) of every Γ-PJ-default rule. A necessary disbelief constraint that is not retained in the updated set of constraints is discarded entirely. However, a necessary belief constraint that is removed from the set of belief constraints is still retained in the theory as a default rule. We require that the dummy default rule $\frac{:\top}{\top}$ be an element of D for every (W, D) representing a belief state. This is to enable us to record necessary disbelief constraints even if there are no other elements in D. **Example:** Let the initial belief state be given by the Γ-PJ-default theory (W, D) where $W = \{a, a \rightarrow b, c, c \rightarrow d\}$ and $D = \{\frac{:\top}{\top}\}$. Notice that there are no necessary disbelief constraints at this point. Assume that the elements of W were obtained in a single belief change step. Thus, there is a single necessary belief constraint and $BC_{belief} = \{\{\{a\}, \{\neg a, b\}, \{c\}, \{\neg c, d\}\}\}$.

The belief state that results from the contraction of b, denoted by $(W, D)^-_b$, is computed as follows:

- **Belief Constraints Update**

$$BC_{belief_{new}} = \{c, c \rightarrow d\}$$
$$BC_{disbelief_{new}} = \{b\}$$

- **Default Theory Update**

$$W' = \{c, c \rightarrow d\}$$
$$D' = \{\frac{: a \wedge \Gamma}{a}; \frac{: (a \rightarrow b) \wedge \Gamma}{(a \rightarrow b)}\}, \ where \ \Gamma = \{\neg b\}$$

Note that, as a result of the contraction, a and $(a \rightarrow b)$ become *tentative beliefs*. Further, if we get new evidence that a has to be believed again, it is set theoretically added to W, but b is still disbelieved.□

As a $\Gamma-$PJ-default theory is a PJ-default theory, every $\Gamma-$PJ-extension E can be computed using **Definition 4**. That is,
$E = (E_J, E_T) = (\bigcup_{i=0}^{\infty} E_{J_i}, \bigcup_{i=0}^{\infty} E_{T_i})$.

Being that E_J, the "semi-normal" part of the justification, is the same for every default they need not be recomputed for every extension. Due to this, we give an alternative way of computing the extensions of a $\Gamma-$PJ-default theory.

Definition 12 *The set of* **active defaults** *for a $\Gamma-$PJ-default theory (W, D) is defined as:* $\Delta_{(W,D)} = \{\frac{:\beta}{\alpha} \in D$ *and* $Cn(W \cup \{\alpha\} \cup \Gamma)$ *is consistent*$\}$.

Definition 13 *The set of* **consequents** *of active defaults is defined as:* $Cons_{(W,D)} = \{\alpha \mid \frac{:\beta}{\alpha} \in \Delta_{(W,D)}\}$.

Definition 14 *The set of* **tentative beliefs** M *of a* Γ*–PJ-default theory* (W, D) *is defined as:* $M_{(W,D)} = \{(e_i \setminus Cn(W))$ *s.t.* $e_i \in E(W, D)\}$.

Theorem 1 *Let* $E_{(W,D)}$ *be the set of extensions of a* Γ*–PJ-default theory* (W, D). *We then have that:* $E_{(W,D)} = \{Cn(W \cup \delta_i) \mid \delta_i \in (Cons_{(W,D)} \downarrow \Gamma)\}$.

Proof: The proof is given in the appendix.□

By representing each belief state as a Γ-PJ-default theory, we have factored out the question of which theory (extension) to choose as the currently operative set of beliefs from the process of belief change. The user, or agent, could thus use a variety of techniques to actually pick the currently operative set of beliefs, including priorities on beliefs (such as epistemic entrenchment in the AGM framework) or even use the temporal order in which the currently tentative beliefs were obtained (given that the temporal order on the full beliefs is already being used for identifying the new set of belief constraints). Our framework would, in fact, permit a dynamic prioritization of beliefs; different orderings could thus be used at different times to select theories without causing inconsistencies or unreasonable outcomes from the belief change process.

The intuitions formalized in the operators presented above represent a somewhat strict set of requirements for iterated belief change. It might be reasonable under some circumstances to relax these requirements. Our account of belief change has the advantage of being general enough to provide a framework within which a variety of intuitions regarding belief change can be formalized. For instance, if one wished to disregard the temporal order in which belief constraints were obtained while updating the set of belief constraints, this would correspond to using an empty constraint prioritization relation. A belief change operator that does not propagate belief constraints could be modeled by restricting necessary disbeliefs constraints to be only the current disbelief constraint, in the case that the current belief change step is one of contraction, and to be empty otherwise and by making the constraint prioritization relation empty. Restricting the set of defaults, the set of necessary disbelief constraints and the constraint prioritization relation to be empty would make our operator identical to Nebel's base contraction operator. A belief change step caused by evidence from a dubious source could be modeled by placing the new belief constraint lower in the constraint prioritization (instead of putting it at the top).

In the next chapter we study the properties of our framework and we relate it to some of the existing formalizations.

Chapter 5

Properties and Related Work

In the previous chapter we introduced an alternative framework for iterated belief change. Here, we study its properties as well as its relation to some of the existing formalizations. These properties and relations are largely based on those that appear in [12].

5.1 Properties

It has become popular in recent times to evaluate every new belief change operator against the yardstick of rationality provided by the AGM postulates for belief change, primarily because they seem to be the best formalization of the consensus view on the requirements that an ideal belief change operator should satisfy. Our formalization cannot, however, be evaluated using the AGM postulates directly, for the following two reasons. *First*, the AGM postulates consider transitions between belief states represented as a single deductively closed theory. Our operator maps between belief states represented as collections of theories (the multiple possible extensions of the PJ-default theories). *Second*, since the AGM postulates consider belief change as a single step process, it is difficult to evaluate "rationality" over iterated belief change steps. It is possible, however, to reformulate the postulates to enable us to articulate a set of similar properties about our framework.

It turns out that our operator satisfies the reformulated versions of the AGM contraction postulates (1-) through (8-). We interpret postulate (1-), as one way of articulating the following *principle of categorical matching* stated by Gärdenfors and Rott in [10]:

> The representation of a belief state after a belief change has taken place should be of the same format as the representation of the belief state before the change.

For postulates (2-) through (8-), we reformulate every condition on knowledge sets to apply to every extension of the PJ-default theory representing a belief state. For postulates (7-) and (8-) we can actually prove a stronger condition. This is not surprising as we do not throw away any piece of information after each contraction. $E(W, D)$ refers to the set of extensions of the $\Gamma-$PJ-default theory (W, D).

Definition 15 *Let* (W, D) *and* (W', D') *be* $\Gamma-$*PJ-default theories. We say that* $E(W, D) \subseteq E(W', D') \;\leftrightarrow\; \forall e_i \in E(W, D),\; \exists e'_j \in E(W', D')$ *s.t.* $e_i \subseteq e'_j$.

Theorem 2 *The contraction operator* $-$ *satisfies:*

(1-) The principle of categorical matching.

(2-) $E(((W,D)^{-}_{A}) \subseteq E(W,D)$.

(3-) *If* $\forall e \in E(W,D)\ e \not\models A$, *then* $E((W,D)^{-}_{A} = E(W,D)$.

(4-) *If* $\not\models A$, *then* $\forall e : (e \in E(W,D)^{-}_{A}) \rightarrow (e \not\models A)$.

(5-) *If* $\exists \in E(W,D)$ *s.t.* $A \in e$, *then* $E(W,D) \subseteq E(((W,D)^{-}_{A})^{+}_{A})$.

(6-) *If* $\models A \leftrightarrow B$ *then* $E((W,D)^{-}_{A}) = E((W,D)^{-}_{B})$.

(7-)–(8-) *If* $\forall e : (e \in E((W,D)^{-}_{A \wedge B})) \rightarrow (e \not\models A)$ *then* $E((W,D)^{-}_{A \wedge B}) = E((W,D)^{-}_{A})$.

Proof: The proof of properties (2-)–(7-) is given in the appendix.□

Once more, it is important to point out that the *contraction* and *undo* operators are different. It is easy to see why the *undo* operator does not satisfy the reformulated version of postulate (2-). As a consequence of retracting a necessary disbelief constraint, a PJ–default rule that could not earlier be fired, (because of the presence of the necessary disbelief constraint that is now discarded), is enabled. Then, a new formula which was not in any of the extensions for the earlier belief state could thus appear in some extension of the updated belief state.

A similar set of properties can be identified for our revision operator $*$.

Theorem 3 *The revision operator* $*$ *satisfies:*

(1) The principle of categorical matching.*

(2)* $\forall e : (e \in E((W,D)^{*}_{A})) \rightarrow (e \models A)$.

(3)* $E((W,D)^*_A \subseteq E((W,D)^+_A$.

(4)* *If* $\forall e \in E(W,D)\ e \not\models \neg A$, *then* $E((W,D)^+_A \subseteq E((W,D)^*_A$.

(6)* *If* $\models A \leftrightarrow B$, *then* $E((W,D)^*_A) = E((W,D)^*_B)$.

(7)–(8*)* *If* $\forall e : (e \in E((W,D)^*_{A\vee B})) \rightarrow (e \not\models \neg A)$ *then* $E((W,D)^*_{A\vee B}) = E((W,D)^*_A)$.

Proof: The proof of these properties is immediate using the Levi identity defined below.□

In a result analogous to the *Levi identity* [19], we can show that the operations of revision and contraction in our framework can be defined in terms of each other.

Conjecture 1 $E((W,D)^*_A) = \bigcup_i Cn(Z_i \cup \{A\})$ *where* $Z_i \in E((W,D)^-_{\neg A})$.

The properties listed above indicate that the behaviour of our belief change operators is reasonable over a single belief change step. The final result in this section indicates that our operators behave reasonably over iterated belief change steps as well, in the sense that they propagate the effects of a belief change step as far as is consistently possible.

Conjecture 2 *Let* $(W_0, D_0) = (W,D)^-_x$ *and* $(W'_0, D'_0) = \{\{\}, \{\frac{:\top}{\top}\}\}$. *Let* (W_n, D_n) *denote the belief state obtained after n belief change steps (where each belief change step involves the application of operator* $*$ *or* $-$*) from* (W_0, D_0). *Let* (W'_n, D'_n) *denote the belief state obtained from* (W'_0, D'_0) *after the same n belief change steps. Then, for every n,* $\exists e : (e \in E(W_n, D_n)) \wedge (e \models x)$ *if and only if* $\exists e' : (e' \in E(W'_n, D'_n)) \rightarrow (e' \models x)$.

5.2 Related Work

Over a single step, and starting with a deductively closed theory, our framework is identical to the AGM framework in the sense that the outcomes generated by our framework are identical to the choices available to the AGM *selection function*.

Theorem 4 *For a Γ-PJ-default theory (W, D), $E((W, D)_A^-) = K \downarrow A$ if $W = K$ where K is a deductively closed theory, $D = \{\frac{:\top}{\top}\}$ and given an empty constraint prioritization relation.*

Proof: The proof is given in the appendix.□

Thus, if we were to start with a deductively closed theory as the set of facts, and an empty set of defaults, then the set of extensions of the default theory obtained after contracting a belief A would correspond precisely to the set of possible outcomes that the selection function in partial meet contraction. Whereas partial meet contraction requires that a choice is actually made, we do not require any choices, but retain all the multiple possible outcomes compactly represented as a PJ-default theory.

The following theorem shows how our approach relates to Nebel's base contraction operator [22].

Theorem 5 *Let (W, D) be a Γ-PJ-default theory with $W = B$, where B is a finite belief base, and $D = \{\frac{:\top}{\top}\}$. Then $Cn(B \simeq A) = Cn((\bigvee E((W, D)_A^-)) \wedge (B \vee \neg A))$ given an empty constraint prioritization relation.*

Proof: The proof is given in the appendix.□

As with our framework, a belief that becomes suppressed as a result of a contraction operation can be recovered in Nebel's framework when the belief state is revised with the contracted belief. However, our framework permits an explicit operation to undo a contraction, which can also result in beliefs being recovered. Such an operation is not possible in Nebel's framework.

With their commitment to producing a unique outcome for the belief change operation, both the AGM and Nebel frameworks render too many potentially useful beliefs unusable (notice that they are not actually discarded, but can be recovered later under certain circumstances); in the AGM framework, this is a consequence of taking the intersection of the selected outcomes, while in Nebel's case, this is a consequence of taking the disjunction of every possible outcome. Our framework retains every possible outcome at all times, and thus does not suffer from this problem.

Brewka [3] shows how belief revision can be viewed as a simple process of adding new information to theories represented in his *preferred subtheories framework*, which in turn is a generalization of the THEORIST framework introduced by Poole, Goebel and Aleliunas in [24]. The THEORIST framework envisages a knowledge base comprising of a set of closed formulas that are necessarily true, called *facts*, and a set of possibly open formulas that are tentatively true, called *hypotheses.* Default reasoning in this framework involves identifying *maximal scenarios* (or extensions), where a scenario consists of the set of facts together with some subset of the set of ground instances of the hypotheses which is consistent with the set of facts. The framework can be augmented with *constraints*, which are closed formulas such that every THEORIST scenario is required to be consistent with the set of constraints. Brewka's framework of preferred subtheories differs from THEORIST in two significant ways. First, *facts* are done away with, making every formula in the knowledge base potentially refutable. Second, one is allowed to define a partial order on the formulas in the knowledge base. A preferred subtheory, Brewka's analogue of a THEORIST maximal scenario, is a consistent subset of the knowledge base constructed by starting with formulas with the highest priority (as defined by the partial order) and progressively adding as many formulas of lower priority levels as can be consistently added. As with THEORIST, a knowledge base can have several preferred subtheories. Brewka shows in [3] that a knowledge base of this kind can be revised by simply adding the new formula and augmenting the partial order to incorporate any ordering relationships that might exist between this formula and the existing elements of the knowledge base. Also, if this framework is augmented to include THEORIST-style constraints, and a partial order is defined on the set containing both the formulas representing hypotheses and formulas representing constraints, then contraction is shown to be a simple case of adding a constraint to the knowledge base and augmenting the partial order. The improvements achieved by Brewka's belief change framework over earlier ones are twofold. First, the belief change operator is simple and totally incremental. Second, earlier information is not thrown away, but is retained in an elegant fashion. Nebel [23] establishes a restricted form of equivalence between nonmonotonic inference and belief change along similar lines.

In the case that the new belief (either a new hypotheses, as in revision, or a new constraint, as in contraction) always has a higher priority, under the partial ordering, than all existing beliefs, Brewka's framework turns out to be very similar to ours. As the following example shows, his framework avoids the problem of spurious beliefs in most cases.

Example: Let the initial knowledge base consist of the set $\{a, a \rightarrow b\}$ of hypotheses with no ordering relationship being defined on the hypotheses. In order to contract b from this knowledge base, we add the constraint $\neg b$, written as $< \neg b >$ to the knowledge base, together with the ordering relations $< \neg b > \geq a$ and $< \neg b > \geq a \rightarrow b$. We get two maximal scenarios , one containing a and the other containing $a \rightarrow b$. Further revision of the knowledge base with a results in the addition of this hypotheses at a higher priority level than all existing elements (hypotheses or constraints) of the knowledge base. There is only one maximal scenario at this point, consisting of a and its logical consequences. The spurious belief b does not reappear, as with the operators defined by Nebel [22] or by Alchourrón, Gärdenfors and Makinson [1]. □

The similarity of Brewka's framework to ours is not surprising, given that we like Brewka, use nonmonotonic theories which can generate possibly many different consistent sets of beliefs, to represent a belief state. Like Brewka, our approach is incremental, and information is never thrown away. Our choice of nonmonotonic formalism is very similar too, given the results in [5] relating PJ-default logic to THEORIST with constraints. However, since Brewka does not explicitly account for belief constraint propagation, his formalization is not entirely free of the problem of spurious beliefs, as the following example shows.

Example: Consider an initial knowledge base containing only one hypotheses and no constraints $\{a\}$. Let us now contract $a \vee b$ from this knowledge base. This entails the addition of the constraint $< \neg(a \vee b) >$ to the knowledge base, and augmenting the partial order such that the new constraint has higher priority than all existing elements of the knowledge base. If one were to revise the knowledge base with b, there would be one maximal scenario containing both a and b. Notice, however, that the new evidence obtained since retracting $a \vee b$ from the knowledge base does not warrant renewed belief in a. The problem arises because the presence of b at a higher priority level disables the constraint $< \neg(a \vee b) >$.□

Brewka does not explicitly consider disbelief propagation in his account of belief change based on the preferred subtheories framework. We improve upon Brewka's work by explicitly accounting for disbelief propagation. Necessary disbelief constraints are treated as a set of formulas to be explicitly disbelieved. We update this set at every belief change step, by retaining as many constraints, or parts of constraints, as are compatible with more recent constraints. Thus in the previous example, we would update this theory to account for revision with b by removing b from the set of necessary disbelief constraints, but leaving a intact. Brewka's framework does not achieve full disbelief propagation because it uses syntactic units (the constraints) which are enabled or disabled as whole units and not in terms of the individual components. In fact, his framework avoids the problem of spurious beliefs only if the only constraints permitted are atomic constraints.

We differ further from Brewka in that we factor out the use of priorities on beliefs entirely from the belief change process. Whereas Brewka's framework would only generate those maximal scenarios which respect the existing orderings on the beliefs, our framework would generate all maximal scenarios which satisfy the relevant belief constraints. Our framework would coincide with Brewka's, in this respect, if the only ordering relations were those generated by belief change steps.

Chapter 6

Conclusion

6.1 Contributions

- The primary contribution of this study is the identification of a set of requirements for belief change operators from the viewpoint of their iterated application:
  - *The Principle of Irrelevance of Choice*,
  - *The Persistence Postulate*
- A threefold categorization of beliefs in a belief state is introduced, as well as a precise characterization of belief migration across these classes is provided.
  - *Necessary Beliefs*,
  - *Necessary Disbeliefs* and
  - *Tentative Beliefs*

- The problem of spurious beliefs is identified, and is used to motivate the need both for maintaining multiple theories in the same belief state, as well as for propagating belief constraints.

- An alternative framework for iterated belief revision is proposed. Some of its advantages are:

  - A compact representation is obtained for the multiple possible outcomes of a belief change operation.

  - Belief constraints need not be recorded separately. They are clearly identifiable components of the PJ-default theory.

  - By retaining all outcomes at any given point during the process of belief change, we permit the use of different prioritizations of the beliefs at different times to actually select one theory out of the many that may potentially constitute a belief state.

- Our framework can be viewed as a formalization of the process through which a default theory evolves into a more accurate representation of the world. Belief change drives this process, and in abstract terms, this involves demoting facts known to be true to the status of defaults.

- By modelling the belief state as a *belief base*, finite set of sentences, we give the first step towards bridging the gap between theory and practice.

- By showing how a nonmonotonic formalism such as default logic can be crucial to knowledge revision, this study further explicates the close relationship between these related, yet separate, areas of inquiry.

6.2 Directions for Future Research

Some of the possible venues for further research are:

- Study how our rationale for iterated belief change can be applied to other nonmonotonic formalisms. This should be done not only at the theoretical level but it would also be interesting to adapt it to practical implementations (for example, THEORIST).

- As we already mentioned, our approach constitutes a first step towards bridging the gap between theory and practice. It would also be interesting to relate it to empirical studies on belief revision on human populations.

- Study if there is any counterpart in model–theoretic approaches to belief revision where our ideas on iterated belief change can be applied. It would also be interesting to investigate the relationship between our framework and the work done by [2]. He uses conditional logics for revision and the semantics is based on structures consisting of a set of possible worlds W and a binary accessibility relation R over W.

- Study how to apply our ideas in the context of dynamic worlds, (e.g. reasoning about action, planning), in the sense of Katsuno and Mendelzon [18].

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] Craig Boutilier. Sequences of revisions: On the semantics of nested conditionals. Technical Report 23, University of British Columbia, 1992.

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

[4] Johan deKleer. An assumption based truth maintenance system. *Artificial Intelligence*, 28:127–162, 1986.

[5] 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.

[6] Jon Doyle. A truth maintenance system. *Artificial Intelligence*, 12(2):231–272, 1979.

[7] 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.

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

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

[10] Peter Gärdenfors and Hans Rott. Belief revision. Technical Report 30, Zentrum Philosophie und Wissenschaftstheorie, Universität Konstanz. 1992.

[11] Aditya K. Ghose and Randy Goebel. Default reasoning as belief change: A rationale. In *Proc. of the 2nd Pacific Rim Int'l Conf. on AI*, 1992.

[12] Aditya K. Ghose, Pablo O. Hadjinian, Abdul Sattar, Jia-H. You, and Randy Goebel. Iterated belief chage. Technical Report 93-05, University of Alberta, Dept. of Computing Science, Edmonton, Alberta, 1993. Also published at Griffith University, School of Computing and Information Technology, Research Report No. 62.

[13] Matthew L. Ginsberg. *Readings in Nonmonotonic Reasoning*. Morgan Kaufmann, Inc., Los Altos, California, 1987.

[14] Sven O. Hansson. New operators for theory change. *Theoria*, 55:114–132, 1989.

[15] Sven O. Hansson. *Belief base dynamics.* PhD thesis, Uppsala University, 1991.

[16] W. L. Harper. Rational belief change, popper functions and counterfactuals. *Synthese*, (30):221–262, 1975.

[17] 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.

[18] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proc. of the Second Int'l Conf. on Principles of Knowledge Representation and Reasoning*, 1991.

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

[20] D. Makinson. How to give it up: a survey of some formal aspects of the logic of theory change. *Synthese*, 62(4):347–363, 1985.

[21] 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.

[22] 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.

[23] 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.

[24] David Poole, Randy Goebel, and Romas Aleliunas. Theorist: a logical reasoning system for defaults and diagnosis. In N.J. Cercone and G. McCalla, editors, *The Knowledge Frontier: Essays in the Representation of Knowledge*, pages 331–352. Springer Verlag, 1987.

[25] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1 and 2):81–132, 1980.

[26] Abdul Sattar. *Efficiency Aspects of Hypothetical Reasoning.* PhD thesis, University of Alberta, 1991.

Appendix A

Proofs

Theorem 1: Let $E_{(W,D)}$ be the set of extensions of a $\Gamma-$PJ-default theory (W, D). We then have that: $E_{(W,D)} = \{Cn(W \cup \delta_i) \mid \delta_i \in (Cons_{(W,D)} \downarrow \Gamma)\}$.
Proof: First, we show how to compute the extensions of a $\Gamma-$PJ-default theory.
It is easy to see from the definition of PJ-extensions that for a $\Gamma-$PJ-default theory:

$$\forall i \quad E_{J_i} \;=\; Th(E_{T_i} \cup \Gamma)$$

Thus, we can reformulate the following definition for $\Gamma-$PJ-extensions and note that it is strictly equivalent the definition of general PJ-extensions under the condition that for any two defaults: $\frac{:\beta_i \wedge \Gamma_i}{\beta_i} \in D$, $\frac{:\beta_j \wedge \Gamma_j}{\beta_j} \in D$, and $\Gamma_i = \Gamma_j = \Gamma$.

Let (W, D) be a closed $\Gamma-$PJ-default theory,
$E_0 = Cn(W)$

$\vdots$

$E_i = Cn(E_{i-1} \cup \{\beta_i\})$, *s.t.* $\frac{:\beta_i \wedge \Gamma}{\beta_i} \in D$ *and* $E_i \cup (\beta_i \wedge \Gamma)$ is satisfiable.

E is a Γ−PJ-default extension of (W, D) iff $E = \bigcup_{i=0}^{\infty} E_i$. It is easy to see that every extension of the Γ−PJ-default theory can be expressed as $E = Cn(W \cup \delta_i)$ where, $\forall \alpha \in \delta_i$, $\exists \frac{\beta}{\alpha} \in D$.

Finally, we show that $\forall i \;\; \delta_i \in (Cons_{(W,D)} \downarrow \Gamma)$.

As $\delta_i = (\bigcup_{j=0}^{k} \alpha_j)$ and $\forall j \alpha_j$ is the consequent of an *active default*, then $\delta_i \subseteq Cons_{(W,D)}$ and also $(\bigcup_{j=0}^{k} \alpha_j) \cup \neg\Gamma)$ is consistent.

We assume $\exists \chi \in (Cons_{(W,D)} \downarrow \Gamma)$ *s.t.* $\delta_i \subset \chi$, (*i.e.* $\exists x \in \chi$ *and* $x \notin \delta_i$). This also means that $(\delta_i \cup \{x\} \cup \neg\Gamma)$ is consistent. But, when computing extensions we add as many consequents of defaults as is consistently possible. So, $\nexists x \in \chi$ as supposed. Having reached a contradiction, we conclude that $\delta_i \in (Cons_{(W,D)} \downarrow \Gamma)$.□

Theorem 2:

2. $\mathbf{E(((W,D)_A^-) \subseteq E(W,D)}$.

Proof: Assume $\exists e_i \in E(W, D) \mid \forall e'_j \in E((W, D)_A^-)$, $e_i \not\subseteq e'_j$.

Let $(W', D') = (W, D)_A^-$. From Theorem 1 we have that:

$$E_{(W,D)} = \{Cn(W \cup \delta_i) \mid \delta_i \in (Cons_{(W,D)} \downarrow BC_{disbelief_{old}})\}$$
$$E_{((W,D)_A^-} = \{Cn(W' \cup \gamma_j) \mid \gamma_j \in (Cons_{(W',D')} \downarrow BC_{disbelief_{new}})\}.$$

To prove that our assumption is wrong, we show that $W = W'$ *and* $D = D'$. We analyze two cases:

1. A is a full belief: After the contraction operation,

(a) $W' \subset W$. As $A \in Cn(W)$, $\exists S \subseteq W \mid S \vdash A$.

(b) We need to show that $\forall \gamma_j$, $\exists e_i \mid \gamma_j \subseteq e_i$.

We know that $\forall x \in \gamma_j$, either:

- $x \in (W \setminus W')$. Then, $\forall e_i \in E(W,D)$, $x \in e_i$, or

- $\exists \frac{:\beta}{x} \in D'$. Then, $\frac{:\beta}{x} \in \Delta_{(W',D')}$ *and as* $BC_{disbelief_{old}} \subset BC_{disbelief_{new}}$, we have that $\frac{:\beta}{x} \in \delta_{(W,D)}$. Hence, $\exists e_i \mid x \in e_i$.

2. A is a tentative belief: After the contraction operation,

(a) $W' = W$, *as* $A \notin Cn(W)$.

(b) $\Delta_{(W',D')} \subset \Delta_{(W,D)}$, then $\delta_{(W',D')} \subset \delta_{(W,D)}$. Some defaults become passive, i.e. their consequents do not belong to any extension. And, it is clear that $\forall \gamma_j$, $\gamma_j \subseteq \delta_i$.

Finally, from (1) and (2) and having reached a contradiction, we conclude that no matter if A is a full belief or a tentative belief: $\forall e'_j \in E((W,D)^-_A)$ $\exists e_i \in E(W,D) \mid e'_j \subseteq e_i$.□

3. If $\forall \mathbf{e} \in \mathbf{E}(\mathbf{W},\mathbf{D})$ $\mathbf{e} \not\models \mathbf{A}$, then $\mathbf{E}((\mathbf{W},\mathbf{D})^-_{\mathbf{A}} = \mathbf{E}(\mathbf{W},\mathbf{D})$.

Proof: The proof is immediate. From hypothesis, $\forall e$ $e \not\vdash A$. Also, by definition, every extension contains only full or tentative beliefs. Hence, as the contraction operation consists only on adding A to $BC_{disbelief}$, we conclude that the set of extension $E((W,D)^-_A)$ are the same.□

4. If $\not\models \mathbf{A}$, then $\forall \mathbf{e} : (\mathbf{e} \in \mathbf{E}(\mathbf{W},\mathbf{D})^-_{\mathbf{A}}) \rightarrow (\mathbf{e} \not\models \mathbf{A})$.

Proof: The proof is trivial. When computing the new set of belief constraints we know that $-A \in BC_{new}$. In particular, $-A \in BC_{disbelief_{new}}$. Hence, by definition $\forall e \in E((W,D)^-_A)\ A \notin e.\Box$

5. If $\exists \in \mathbf{E(W,D)}$ s.t. $\mathbf{A \in e}$, then $\mathbf{E(W,D) \subseteq E(((W,D)^-_A)^+_A)}$.
Proof: $E(W,D) \subseteq E(((W,D)^-_A)^+_A)$.
Let $(W',D') = (W,D)^-_A$ *and* $(W'',D'') = (W',D')^+_A$.
We assume $\exists e_i \in E(W,D) \ \mid \ \forall e_j \in (W'',D''),\ \ e_i \not\subseteq e_j$. We analyze two cases:

(i) If $A \in W$, after the contraction operation $(W,D)^-_A$, $\forall bc_k \in (W \setminus W')$, bc_k becomes an active or passive *"default"*. Further, after the *expansion* operation $(W',D')^+_A$, A becomes again a full belief, that is $A \in Cn(W'')$, and as $-A$ ceases to be a *disbelief constraint*. Then, every bc_k will reappear in every extension as the reason for them to become defaults, (i.e., *contracting* A), has already disappeared. Indeed, $E(W,D) = E(W'',D'')$.

(ii) If A were a *tentative* belief in (W,D), the defaults in D that become *passive* as a result of performing $(W,D)^-_A$, become *active* defaults again, as $A \in W''$. In this case there is a gain of information as A becomes a *full* belief in (W'',D''). It is clear that $E(W,D) \subset E(W'',D'')$.
Hence, from (i) *and* (ii) we conclude that there is no extension $e_i \in E(W,D)$ as supposed.$\Box$

6. If $\models \mathbf{A \leftrightarrow B}$ then $\mathbf{E((W,D)^-_A) = E((W,D)^-_B)}$.
Proof: Let $(W',D') = (W,D)^-_A$ and $(W'',D'') = (W,D)^-_B$.
We show that $W = W'$ *and* $D = D'$. We analyze two cases:

1. A is a full belief. Then, $A \in Cn(W)$.

   Assume $W \neq W'$. Then, the sets of conflicting beliefs are different, i.e. $(W \setminus W') \neq (W \setminus W'')$.

   But, this is not possible because: $(W \setminus W') \models A \leftrightarrow (W \setminus W') \models B$, *as* $A \leftrightarrow B$ from hypothesis. The same argument applies for $(W \setminus W'')$. Then, contrary to our previous assumption, $W' = W''$, and so does the sets of conflicting beliefs. Hence, $D' = D''$.

2. A is a tentative belief. It means that $A \notin Cn(W)$.

   From hypothesis $A \leftrightarrow B$, then $B \notin Cn(W)$. So $W' = W'' = W$.

   As no new defaults are added and from hypothesis $A \leftrightarrow B$, then the set of new disbelief constraints is going to be the same for both $(W,D)^-_A$ *and* $(W,D)^-_B$.

   Hence, we also have that $D' = D''$.

So, from (1) and (2), and applying Theorem 1 we conclude that $E((W,D)^-_A = E((W,D)^-_B$.□

7. If $\forall e : (e \in E((W,D)^-_{A \wedge B})) \rightarrow (e \not\models A)$ then $E((W,D)^-_{A \wedge B}) = E((W,D)^-_A)$.

Proof Let $(W',D') = (W,D)^-_{A \wedge B}$ and $(W'',D'') = (W,D)^-_A$.

We analyze three cases:

1. $W \models (A \wedge B)$. The idea is to prove that $W' = W''$. Then, using Theorem 1 we show that they have the same set of extensions.

(a) Assume $W' \neq W''$ Then, $BC'_{new} \neq BC''_{new}$.

From hypothesis, $\forall e_i, e_i \not\models A$ so, when computing the new set of belief constraints BC'_{new}, we have that,

$\forall bc_j \in BC_{belief_{old}}, \mid bc_j \models A, \ (bc_j \cup Y_k)$ *is incompatible*, where $k < j$ and Y_k is the new set of belief constraints computed so far.

Then, $\forall x \in BC'_{new}, \ x \in BC''_{new},$, i.e. $BC'_{new} \subseteq BC''_{new}$.
When computing BC''_{new}, by definition, we know that $\forall e'_j, \ e'_j \not\models A$. Applying the same argument as before it is easy to see that $BC''_{new} \subseteq BC'_{new}$.
Having reached a contradiction, we conclude that $BC'_{new} = BC''_{new}$.
Hence, $W' = W''$.

(b) From Theorem 1 we have that:

$$E_{(W',D')} = \{Cn(W' \cup \delta_i) \mid \delta_i \in (Cons_{(W',D')} \downarrow \Gamma')\}$$
$$E_{(W'',D'')} = \{Cn(W'' \cup \gamma_j) \mid \gamma_j \in (Cons_{(W'',D'')} \downarrow \Gamma'')\}.$$

To prove that $E_{(W',D')} = E_{((W'',D'')}$, we first show that $Cons_{(W',D')} = Cons_{(W'',D'')}$.

- $\forall \alpha \in Cons_{(W',D')}, \ (W \cup \{\alpha\}) \not\models A, \ then \ \alpha \in Cons_{(W'',D'')}$.
- $\forall \beta \in (Cons_{(W'',D'')}, \ (W \cup \{\beta\}) \not\models A, \ then \ (W \cup \{\beta\}) \not\models (A \wedge B). \ So, \ \alpha \in (Cons_{(W',D')}$.

Using the result obtained in (a), $(BC'_{new} = BC''_{new})$, we can infer that $(\Gamma' \setminus \Gamma'') = \{A \wedge B\}$ *and* $(\Gamma'' \setminus \Gamma') = \{A\}$.
So, $\forall \delta_i, \ \delta_i \not\models A, \ then \ \delta_i \in (Cons_{(W'',D'')} \downarrow \Gamma'')\}$. Also, $\forall \gamma_j, \ \gamma_j \not\models$

A, *then* $\gamma_j \not\models (A \wedge B)$. *Hence*, $\gamma_j \in (Cons_{(W',D')} \downarrow \Gamma')\}$.

From (a) and (b), we conclude that $E_{(W',D')} = E_{(W'',D'')}$.

2. Both, A and B are tentative beliefs.

   As $W \not\models (A \wedge B)$ we have that after the contraction operation $W' = W'' = W$ and there are no new defaults. The only difference between the two Γ−PJ-default theories after contraction is given by the set of necessary disbeliefs. So, using $1.(b)$ from above, we conclude that $E_{(W',D')} = E_{(W'',D'')}$.

3. $W \models B$ and A is a tentative belief.

   As above, $W \not\models (A \wedge B)$. So, again, using $1.(b)$ we can conclude that $E_{(W',D')} = E_{(W'',D'')}$. □

Theorem 4: For a Γ-PJ-default theory (W, D), $E((W, D)^-_A) = K \downarrow A$ if $W = K$ where K is a deductively closed theory, $D = \{\frac{:\top}{\top}\}$ and given an empty constraint prioritization relation.

Proof: Let $E = \{e_i \mid e_i \in (K \downarrow A)\}$ and

$E' = E(W', D') = \{Cn(W' \cup \delta_i) \mid \delta_i \in (Cons(W', D') \downarrow \Gamma')\}$, *where* $\Gamma' = \{A\}$.

Note that $W' = \cap(W \downarrow A)$ because $D = \{\frac{:\top}{\top}\}$ *and* $BC = \{\{\}, \{\}\}$, from hypothesis.

Before proving that $E = E'$, we to show that:

- $\bigcap_{S \in (Cn(W) \downarrow A)} S = Cn(\bigcap_{S' \in (W \downarrow A)} S')$.

$\Rightarrow$) $\forall z \in X \mid X \subseteq \bigcap_{S \in (Cn(W) \downarrow A)} S$, either $z \in S'$, or $S' \models z$. Then, $\forall S'\ z \in (\cap S')$. So, we can conclude that $z \in Cn(\bigcap_{S' \in (W \downarrow A)} S')$.

$\Leftarrow$) *If* $a \in Cn(\bigcap_{S' \in (W \downarrow A)} S')$, as $S' \subseteq W$, $a \in W$. So, $a \in Cn(W)$. Also, as $S' \not\models A$ *then* $\forall S : S \in (Cn(W) \downarrow A)$, $a \in S$.

Hence, we can conclude that $\forall a$, $a \in \bigcap_{S \in (Cn(W) \downarrow A)} S$.

Finally, we show that $E = E'$. We do this by studying the full and tentative beliefs of each of the extensions.

1. $\forall \phi \in \cap e'_j$, $\phi \in Cn(W') = Cn(\cap(W \downarrow A)) = \cap(Cn(W) \downarrow A) = \cap(K \downarrow A)$.

   So, E *and* E' have the same set of full beliefs.

2. $\exists e_i \mid \phi \notin e_i$.

   So, $\phi \in (K \setminus \cap(K \downarrow A)) = Cn(W) \setminus \cap(Cn(W) \downarrow A) = Cn(W \setminus \cap(W \downarrow A)) = Cn(W \setminus W')$.

   So, E *and* E' also have the same set of tentative beliefs.

Hence, from (1) and (2) we conclude that $E = E'$.□

Theorem 5: Let (W, D) be a Γ-PJ-default theory with $W = B$, where B is a finite belief base, and $D = \{\frac{:\top}{\top}\}$. Then $Cn(B \simeq A) = Cn((\bigvee E((W, D)^-_A)) \wedge (B \vee \neg A))$ given an empty constraint prioritization relation.

Proof: The original definition of Nebel's *base contraction* is given by, $(B \simeq A) = (\bigvee_{C \in (B \downarrow A)} C) \wedge (B \vee \neg A)$.

Using Theorem 4 we know that $e = (B \downarrow A) = \{e_1, \ldots, e_n\}$ where $e_i \in E((W, D)^-_A)$, $W = B$ and $D = \{\frac{:\top}{\top}\}$. So, replacing the first part of the conjunction above by $\bigvee_{i=1}^{n} e_i$, we have that:

$Cn(B \simeq A) = Cn((\bigvee_{i=1}^{n} e_i) \wedge (B \vee \neg A))$.□