Pseudo-Contractions in Belief Revision

Yuri David Santos

Advisor: Prof. Renata Wassermann, PhD

Co-Advisor: Prof. Márcio Moretto Ribeiro, PhD

São Paulo, May 2016

Pseudo-Contractions in
Belief Revision

This version of the dissertation contains corrections and amendments suggested by the Judging Committee for the defense of the original version of the work carried out on 23/02/2016. A copy of the original version is available at the Instituto de Matemática e Estatística, Universidade de São Paulo.

Judging Committee:

- Prof. Renata Wassermann, PhD (advisor) - IME-USP
- Prof. Marcelo Finger, PhD - IME-USP
- Prof. Eduardo Fermé, PhD - Universidade da Madeira

To Seila and Rosa

Acknowledgments

Abstract

Belief Revision addresses the problem of how to change epistemic states, usually represented in the literature by sets of logical sentences. Solid theoretical results were consolidated with the AGM paradigm, which deals with theories (logically closed sets of sentences). After that, the theory was extended to belief bases, that is, arbitrary sets of sentences. Besides all this theoretical framework, AI researchers face serious difficulties when trying to implement belief revision systems. One of the major complications is the closure required by AGM theory, which cannot be easily computed. Even belief bases, which do not require closure, seem to be improper for practical purposes, since their changes are usually very rigid (syntax dependent).

Some operations, known as pseudo-contractions, are in the middle ground between belief set change and belief base change. In the present work we have proposed a new pseudo-contraction operation, studied its properties and characterized it. We have also found connections between this operator and some other pseudo-contractions.

Keywords: belief revision, pseudo-contractions.

Resumo

Revisão de Crenças aborda o problema de como alterar estados epistêmicos, normalmente representados na literatura como conjuntos de sentenças lógicas. Resultados teóricos sólidos foram consolidados com o paradigma AGM, que lida com teorias (conjuntos de sentenças logicamente fechados). Depois disso, a teoria foi estendida para bases de crenças, isto é, conjuntos arbitrários de sentenças. Apesar de todo esse arcabouço teórico, pesquisadores de IA enfrentam sérias dificuldades ao tentar implementar sistemas de revisão de crenças. Uma das maiores complicações é o fecho exigido pela teoria AGM, que não pode ser facilmente computado. Mesmo bases de crenças, que não exigem fechamento, parecem ser impróprias para fins práticos, pois suas alterações são geralmente muito rígidas (dependentes de sintaxe).

Algumas operações, conhecidas como pseudo-contrações, estão no meio do caminho entre mudanças para conjuntos de crenças e mudanças para bases de crenças. Nesse trabalho, propomos uma nova operação de pseudo-contração, estudamos suas propriedades e a caracterizamos. Também encontramos conexões entre esse operador e algumas outras pseudo-contrações.

Palavras-chave: revisão de crenças, pseudo-contrações.

Contents

Chapter 1

Introduction

Belief revision (also known as *Belief Dynamics*, *Belief Change* or *Theory Change*) is a field of study that addresses the problem of how to change belief states (or epistemic states). Rational agents, be they humans or machines, certainly have some representation of their knowledge or belief system (the problem of how to represent knowledge and belief systems is studied by a field in the intersection of computer science and philosophy, *Knowledge Representation*). How do these agents change these representations when they are faced with new information? This is the question that Belief Revision wants to answer. As Gärdenfors (one of the founders of the field) has claimed [Gär88, preface], there is little use in knowing how to represent knowledge if we do not know how to change these representations.

The initial hallmark of this research area is the paper by Alchourrón, Gärdenfors and Makinson [AGM85], in which they propose to represent belief states by sets of logical propositions closed under some consequence operator, the so-called *belief sets*. Three main operations can be applied on these belief sets: expansion, contraction and revision. *Expansion* is the simple addition of a new belief (a logical sentence), followed by the closure of the resulting set. *Contraction* is the removal of a belief and *revision* is the addition of a new belief, as in expansion, but with the requirement that the resulting set be logically consistent. Contraction and revision operations cannot be uniquely defined, but have been constrained by a set of *rationality postulates* instead.

In further studies (such as [Han92a]), a generalization of the AGM theory was proposed wherein epistemic states were represented by arbitrary sets of logical sentences, not necessarily closed, called *belief bases* (sometimes called just *bases*). From a computational standpoint, belief bases are very useful since they do not demand logical closure, which is hard to compute, if possible at all. It turns out that bases have a theoretical advantage also, for they are more expressive than belief sets. A belief set can be obtained from the closure of a belief base, whereas many belief bases can "represent" the same belief set.

One of the contraction postulates, *success*, requires the contracted element not to be entailed by the resulting epistemic state. Although belief bases are not necessarily closed, the success postulate still checks the contracted element against its closure, to see if it is implied. So, it is as if there are two classes of beliefs in the base: the represented beliefs (also called *explicit beliefs*) and the derived beliefs (*implicit beliefs*).

Belief bases are not free of problems, though. Since traditional contractions respect the *inclusion* postulate (a requirement that no new element is added in the operation), base contraction can be very destructive. Without adding new elements, it is not possible to "weaken" any formula, which could be enough to perform the contraction in consonance with success. This means that more than what is necessary will be removed, considering that we want *minimal informational change*[1]. This problem can also be viewed as the problem of *syntactic dependence*, i.e., for two bases representing the same belief set, the codification of the base (the individual formulas it contains) can lead to different contractions. This phenomenon does not exist in the context of belief sets, they represent

---

[1]For a thorough discussion on this principle, see [Rot00]. Harman also considers this principle in his book [Har86], for instance, in the discussions that start at pages 46 and 59.

a belief state in the *knowledge level* [New82], ignoring syntactic variations.

A first goal of this work is to investigate contraction operations that lie between belief base contractions and belief set contractions, trying to find some balance between syntactic independence and computational efficiency. We have described some of the properties of these operations and characterizations for them (*representation theorems*). In order to accomplish that, we have given up the inclusion postulate, giving rise to what is called *pseudo-contraction*. Secondly, we want to investigate the relations between our proposal and other pseudo-contractions.

The rest of this dissertation is organized as follows. Chapter 2 introduces the subject of logic with material found in the literature, defining classical propositional logic and some notational conventions. Chapter 3 lays out the theoretical background on belief revision and pseudo-contractions. Chapters 4, 5 and 6 represent our main contributions. In the first, we describe a new pseudo-contraction, explaining its usefulness through examples, proving some formal properties and relationships with other pseudo-contractions. In the second, we put forth a discussion about minimality, resource-bounded reasoners and postulates. In the third, we show two concrete examples where our theory can be applied. Chapter 7 briefly summarizes some related work, and finally, in Chapter 8 we conclude, analyzing possibilities for future work.

Chapter 2

Logic

Logic is the basis for Belief Revision and for most studies in the field of Knowledge Representation (and for a lot of other things, of course). For this reason, we are going to talk about logic in this chapter[1].

First we just introduce the subject of logic, in Section 2.1. Next, we present the classical propositional logic in Section 2.2. In Section 2.3, we define and talk a little about consequence relations and some of their properties. Then, in Section 2.4 we take account of basic notational conventions in use along this text.

2.1 Why Logic?

Belief Revision and Knowledge Representation are deeply related to Artificial Intelligence, and most of their applications have probably something to do with intelligent systems, "programs that think" in some (limited) way. With this in mind, we must quote this wise paragraph by Brachman and Levesque [BL04]:

> *Before any system aspiring to intelligence can even begin to reason, learn, plan, or explain its behavior, it must be able to formulate the ideas involved. You will not be able to learn something about the world around you, for example, if it is beyond you to even express what that thing is. So we need to start with a* language *of some sort, in terms of which knowledge can be formulated.*

So, we need a language in terms of which some knowledge can be expressed and then, bearing these representations, some reasoning can be done in order to learn new things. As machines cannot read natural language (very well) yet, to make this process easier, it is better to employ a formal language, that is, a language that is clearly and mathematically defined. Logics are good candidates for this kind of language.

In the next section we describe classical propositional logic in detail.

2.2 Classical Propositional Logic

Among a plethora of logics, we chose the classical propositional logic to describe here in detail because it is *perhaps* the most known, natural and commonly used logic today. It also serves as a backbone for many other logics. Besides that, it is the logic we use in most examples in this dissertation.

Still inspired by [BL04], we will define precisely the language of classical propositional logic[2] through two aspects: its syntax and semantics.

---

[1] This chapter was included only in order to make this dissertation self-contained for researchers of computer science and correlated areas. Those who already have a good background in logic can skip straight to Section 2.4.

[2] A better and deeper presentation of this logic can be found in [Eps90]. Classical propositional logic can be understood as a fragment of first order logic (FOL), which is carefully introduced in [BL04].

In this section we will call the language of classical propositional logic $\mathfrak{L}$. A formal language like $\mathfrak{L}$ can be considered to be a (possibly infinite) set of sentences.

2.2.1   Syntax

Each sentence of $\mathfrak{L}$ is a sequence of symbols. The set of allowed symbols together with rules specifying how they can be arranged to form a well-formed sentence is what we call syntax.

There are two classes of symbols in propositional logic: the logical and the nonlogical ones. The logical symbols are: *not* ($\neg$), *and* ($\wedge$) and *or* ($\vee$). The nonlogical symbols are an infinite set $\mathcal{P}$ of *propositional symbols* or *propositional atoms* (or just *atoms*), and are usually represented by lowercase latin letters ($p, q, r, ...$). We have also parentheses as ancillary symbols.

Now for the rules. Any atom (member of $\mathcal{P}$) is a (well-formed) sentence of $\mathfrak{L}$. If α and β are (well-formed) sentences of $\mathfrak{L}$, the following are (well-formed) sentences of $\mathfrak{L}$ as well:

- $\neg(\alpha)$

- $(\alpha \wedge \beta)$

- $(\alpha \vee \beta)$

Nothing more than these are sentences of $\mathfrak{L}$. So, by those rules, the atoms p and q are both well-formed sentences of $\mathfrak{L}$. Then, $(p \wedge q)$ is a sentence as well, and therefore $((p \wedge q) \vee p)$ and $\neg((p \wedge q) \vee p)$ are sentences too, but $(\neg \wedge p)$ and $(p \neg q)$ are not.

We also use the following abbreviations:

- $(\alpha \rightarrow \beta)$ for $(\neg \alpha \vee \beta)$

- $(\alpha \leftrightarrow \beta)$ for $((\alpha \rightarrow \beta) \wedge (\beta \rightarrow \alpha))$

- $\bot$ for $(\alpha \wedge \neg \alpha)$

- $\top$ for $(\alpha \vee \neg \alpha)$

2.2.2   Semantics

We just defined what the members (the sentences) of $\mathfrak{L}$ are. Now we need to have some standard way to interpret these sentences, a way to specify what they mean, what they say about the world.

It is not possible just to look at a sentence of $\mathfrak{L}$ and tell what it means in an absolute way. This is due to the fact that the atoms do not have a fixed meaning, their meaning are application-dependent. Each atom is a *proposition* about the world. It states that the world is one way, not the other. For example, we can say that p means that *The planet Mercury is round*. As any proposition, p can have two possible values: it can be *true*, or it can be *false*.

So, after defining what the atoms mean, the semantics allows us to understand unambiguously the meaning of any sentence of $\mathfrak{L}$, and it does so through *interpretations*. An interpretation of classical propositional logic is a *valuation*, a function $\mathcal{M} : \mathcal{P} \rightarrow \{true, false\}$. We can say an atomic proposition p is satisfied by an interpretation if it is *true* in that interpretation, that is, if $\mathcal{M}(p) = true$. In that case we write $\vDash_{\mathcal{M}} p$. In the opposite case, we write $\nvDash_{\mathcal{M}} p$. The same kind of notation applies to arbitrary sentences of $\mathfrak{L}$. If $\vDash_{\mathcal{M}} \alpha$, then the sentence α is satisfied by $\mathcal{M}$. But what does that mean? We can define what sentences are satisfied by an interpretation in a recursive fashion from the base case (single atoms), as we have done for the syntax. Let α and β be sentences. Then:

- $\vDash_{\mathcal{M}} \top$

- $\nvDash_{\mathcal{M}} \bot$

- $\vDash_{\mathcal{M}} \neg \alpha$ if and only if $\nvDash_{\mathcal{M}} \alpha$

- $\vDash_{\mathcal{M}} \alpha \wedge \beta$ if and only if $\vDash_{\mathcal{M}} \alpha$ and $\vDash_{\mathcal{M}} \beta$

- $\vDash_{\mathcal{M}} \alpha \vee \beta$ if and only if $\vDash_{\mathcal{M}} \alpha$ or $\vDash_{\mathcal{M}} \beta$

As promised, using this semantics we can now understand the meaning of any sentence of $\mathfrak{L}$ provided we know what the atoms mean. For example, consider the sentence $p \to q$, where p means "the planet is before Venus" and q means "the planet is Mercury". If we say an interpretation $\mathcal{M}$ satisfies this sentence, $\vDash_{\mathcal{M}} p \to q$, which by syntax rules is an abbreviation for $\vDash_{\mathcal{M}} \neg p \vee q$, then, using the rules above, we get that $\vDash_{\mathcal{M}} \neg p$ or $\vDash_{\mathcal{M}} q$, therefore $\nvDash_{\mathcal{M}} p$ or $\vDash_{\mathcal{M}} q$. So, we come to the conclusion that, in this interpretation, either the planet is Mercury (sentence q is *true* in $\mathcal{M}$), or it is not before Venus (sentence p is *false* in $\mathcal{M}$). So, the meaning of the entire sentence is "The planet is Mercury or it is not before Venus".[3]

If a sentence α is satisfied in all interpretations, then we say it is *valid* and write $\vDash \alpha$, otherwise it is *invalid*, and we write $\nvDash \alpha$. If it is satisfied in some interpretation we say it is *satisfiable*, otherwise we say it is *unsatisfiable*. If a sentence α is satisfied in all interpretations that satisfy all sentences in the set A, we say that A *entails* α or that α is a *logical consequence* of A, and write $A \vDash \alpha$. If an interpretation $\mathcal{M}$ satisfies a sentence α, we say that $\mathcal{M}$ is a *model* of α, and if it satisfies all sentences of a set A, we say that it is a model of A. Deciding whether a formula α is satisfiable is known as the *propositional satisfiability problem* (usually abbreviated as *SAT*).

2.2.3 Natural Deduction

Now we know what formulas are part of the language of propositional logic and what they mean, but we still need to have a method of *proof*, that is, a formal method which can decide, given a set of formulas A and a formula α, whether $A \models \alpha$. If we can find a proof within our method from A to α we write $A \vdash \alpha$. It is desirable that this method be *sound* (if $A \vdash \alpha$, then $A \models \alpha$) and *complete* (if $A \models \alpha$, then $A \vdash \alpha$). The method of proof for classical propositional logic known as *natural deduction* is both sound and complete, although the problem it tries to solve is computationally intractable in the general case[4]. In what follows, we are going to give a brief description of it.

The natural deduction method simply consists of taking a set of premises and repeatedly applying one among a set of rules over them to produce new sentences (which are added to the set of premises and over which the rules can also be applied), until the desired conclusion is obtained.

Now we quickly show the names of the rules[5] followed by its descriptions (inputs and outputs):

- $\wedge i$: if you already have the formulas α and β, you can use this rule to obtain $\alpha \wedge \beta$.

- $\wedge e_1$: if you have $\alpha \wedge \beta$, use this rule to obtain α.

- $\wedge e_2$: from $\alpha \wedge \beta$ obtain β.

- $\vee i_1$: from α obtain $\alpha \vee \beta$.

- $\vee i_2$: from β obtain $\alpha \vee \beta$.

- $\vee e$: if you can demonstrate δ from α and from β separately, and you already have $\alpha \vee \beta$, you can use this rule to obtain δ.

- $\to i$: assume α, if you can prove β from this and the other sentences, you can obtain $\alpha \to \beta$.

- $\to e$: if you have α and $\alpha \to \beta$ obtain β (this rule is also known by its latin name *modus ponens*).

---

[3]We can digress a bit (to a philosophical outlook) and wonder whether two sentences that always evaluate to the same value (*true* or *false*) in all models can be said to have the same meaning.

[4] An intractable problem is a problem that can be solved in theory but its solution takes too long to be calculated in practice, e.g., if it has exponential complexity in the size of the input.

[5]The names are pronounced as a concatenation of the name of the connective followed by *introduction/elimination*, e.g., $\wedge i$ is read *and introduction* and $\wedge e$ is *and elimination*.

- ¬i: obtain ¬α if you can prove $\bot$ from the premises by assuming α (also known by latin name *reductio ad absurdum*, reduction to absurdity).

- ¬e: if you have α and ¬α you can obtain $\bot$.

- $\bot e$: if you have $\bot$ you can obtain α.

- ¬¬i: if you have α you can obtain ¬¬α.

- ¬¬e: if you have ¬¬α you can obtain α.

As an example of a natural deduction procedure, let us derive the rule known as *modus tollens*, i.e., $p \rightarrow q, \neg q \vdash \neg p$.

1. $p \rightarrow q$ (premise)
2. $\neg q$ (premise)
$- - - - - - - - - - --$
3. p (assumption)
4. q (from 1 and 3, applying the rule $\rightarrow e$)
5. $q \wedge \neg q$ (abbreviated as $\bot$, from 2 and 4, applying the rule $\wedge i$)
$- - - - - - - - - --$ (hereafter the assumption p is no longer valid)
6. $\neg p$ (from 3 and 5, applying the rule $\neg i$)

So, since the method is sound and complete, $A \models \alpha$ if and only if $A \vdash \alpha$, i.e., whenever A entails α we can use this method to deduce α from A, and if some α is deduced from A using this method, we can assure that α is a logical consequence of A. The intractability arises from the difficulty in finding the sequence of rules we need to apply and over what input sentences.

2.3   Consequence Relations

A consequence relation is a function $C : 2^{\mathfrak{L}} \rightarrow 2^{\mathfrak{L}}$, that maps sets of sentences (of the language) into sets of sentences. Given a sentence α and a set of sentences A, we say that $\alpha \in C(A)$ if and only if $A \vdash \alpha$, that is, α is provable/deductible from A. The consequence relation of classical propositional logic (where $\vdash$ means provability by natural deduction, for instance) will be denoted by Cn throughout this dissertation (if we use Cn for designating another consequence relation than the classical one, it will be clearly specified). A (logically) closed set is a set K where $K = C(K)$, for some consequence relation C, usually Cn (unless otherwise specified).

We list below some important properties a consequence relation can have.

(Monotonicity) If $A \subseteq B$ then $C(A) \subseteq C(B)$

(Inclusion) $A \subseteq C(A)$

(Idempotence) $C(A) = C(C(A))$

(Supraclassicality) $Cn(A) \subseteq C(A)$

(Compactness) If $\alpha \in C(A)$, then $\alpha \in C(A')$ for some finite $A' \subseteq A$.

(Deduction) $\beta \in C(A \cup \{\alpha\})$ if and only if $\alpha \rightarrow \beta \in C(A)$.

Basically, *monotonicity* is saying that if we expand our set of premises, all the old conclusions will be kept, they should not disappear. *Inclusion* is a requirement that all premises are also conclusions, the reasoning process should not eliminate any premise. *Idempotence* says that the consequence relation computes everything at once, all the conclusions that can possibly be drawn will be drawn

in the first application of the function. A second application over the result of the first one will not yield new conclusions.

Supraclassicality says that the consequence relation will infer at least the classical consequences, that is, the conclusions that one could infer using first order logic. By *compactness*, if some conclusion can be drawn from a set of sentences, then it can be drawn from a finite subset of it.

We say a logic is *Tarskian* if its consequence relation satisfies monotonicity, inclusion and idempotence.

Additionaly, we define the following property:

(Subclassicality) $C(A) \subseteq Cn(A)$

This property just constrains the consequence relation to generate less consequences than classical propositional logic consequence (Cn).

It is worth mentioning that Cn satisfies all these properties (it is Tarskian and satisfies compactness, deduction, supraclassicality and subclassicality).

As a last remark, we should note that to compute the Cn closure of a set of sentences B is a difficult task, since deciding whether a single sentence α is a consequence of B is intractable, finding all the consequences is even harder.

2.4 Notational Conventions

In this section we settle some notational conventions.

We assume a logic is defined as a tuple $\langle \mathfrak{L}, C \rangle$, where $\mathfrak{L}$ is the language and C is a consequence relation. As already mentioned, we usually denote by Cn the consequence relation of classical propositional logic. A closed set is closed by Cn unless otherwise specified. An arbitrary set of sentences will be called a belief base, and a closed one is a belief set or a theory.

Subsets of $\mathfrak{L}$, that is, sets of logical sentences, will be denoted by uppercase latin letters ($A, B, C,...$), usually $A, B, A', B',...$ for arbitrary sets and $K, K',...$ for closed sets. Elements of $\mathfrak{L}$, i.e., logical sentences, are denoted by lowercase greek letters ($\alpha, \beta, \delta,...$) and propositional atoms, by lowercase latin letters ($p, q, r,...$).

The notation $A \vdash \alpha$ (α is provable from A) will be used interchangeably with its equivalent $\alpha \in C(A)$ (α is in the set of consequences of A), and in this case C is the consequence relation associated to the proof method denoted by $\vdash$. We write $A \nvdash \alpha$ if and only if $\alpha \notin C(A)$. Here again, we assume the consequence relation associated to $\vdash$ is Cn unless otherwise mentioned.

Chapter 3

Belief Revision

In this chapter we are going to summarize a fraction of the theory of Belief Revision that will be fundamental for the understanding of this work. First, in Section 3.1, we introduce the theory known as the *AGM Paradigm*, that has begun with the seminal paper by Alchourrón, Gärdenfors and Makinson [AGM85]. Subsequently, in Section 3.2, we present some theoretical background for *Base Change*, a generalization of the AGM theory to arbitrary sets of sentences. To finish, in Section 3.3, we talk about pseudo-contractions, operations similar to contractions but not meeting the inclusion postulate. This last section will be the starting point for our work.

3.1 AGM Paradigm

In this section we will shortly describe the most influential theory of belief revision, that has begun with Alchourrón, Gärdenfors and Makinson (henceforth *AGM*) [AGM85].

AGM have represented epistemic states partially[1] through belief sets, which are sets of sentences closed by some consequence relation (which we will call Cn). So, $K \subseteq \mathfrak{L}$ is a belief set if and only if $K = Cn(K)$. The underlying logic is $\langle \mathfrak{L}, Cn \rangle$, where Cn (which here in the AGM paradigm is not necessarily the consequence of classical propositional logic) is supposed to be Tarskian and to satisfy supraclassicality, compactness and deduction, and the language $\mathfrak{L}$ is supposed to be closed under the classical Boolean connectives[2].

3.1.1 Belief Change Operations

Three epistemic change operations were defined for belief sets: expansion, contraction and revision. These operations are supposed to take place as the agent's response when faced with new information.

Expansion (denoted by $+$) is the simple addition of a sentence to the set, followed by its closure:

$$K + \alpha = Cn(K \cup \{\alpha\})$$

The other two operations are not so simple, and have not been uniquely defined, but they have just been constrained by the rationality postulates. Contraction (denoted by $-$) is the complete removal of a belief from the belief set. Since the belief set must always be closed, the remaining beliefs must not imply the removed one (otherwise it would be present in the closure). In order to comply with this requirement, some other beliefs may be deleted together with the contracted sentence.

Revision (denoted by $*$) is similar to expansion, it is the addition of a new belief, but with the further requirement that the resulting belief set be consistent. To accomplish that, one can

---

[1]Partially because belief states are fully determined by the set of sentences (the belief set) together with its epistemic ordering.

[2]That means: if $\alpha, \beta \in \mathfrak{L}$, then $\neg\alpha \in \mathfrak{L}$, $\alpha \wedge \beta \in \mathfrak{L}$ and $\alpha \vee \beta \in \mathfrak{L}$.

first contract by the negation of the target sentence, then expand by it. This relationship between contraction and revision is known as the *Levi identity* [Gär88], and the inverse relationship is the *Harper identity* [Gär88].

(Levi Identity) $K * \alpha = (K - \neg\alpha) + \alpha$

(Harper Identity) $K - \alpha = (K * \neg\alpha) \cap K$

These relationships come in handy because we can now focus on contractions, after all, revisions can (usually) be obtained from them.

Notice that it is very important to avoid an inconsistent state in the AGM paradigm, as in this occasion we would lose all distinctions in the belief set. Any inconsistency lead to one and the same belief set, the whole language $\mathfrak{L}^3$.

3.1.2   AGM Postulates

As stated earlier, contraction and revision have not been completely defined. Instead, some rationality postulates have been elaborated in order to limit the space of possible contractions (and revisions), intending to allow just the operations considered "rational". The following are the six basic AGM postulates for contraction [AGM85]:

(Closure) $K - \alpha = Cn(K - \alpha)$

(Success) If $\alpha \notin Cn(\emptyset)$, then $\alpha \notin K - \alpha$

(Inclusion) $K - \alpha \subseteq K$

(Vacuity) If $\alpha \notin K$, then $K - \alpha = K$

(Recovery) $K \subseteq (K - \alpha) + \alpha$

(Extensionality) If $Cn(\alpha) = Cn(\beta)$, then $K - \alpha = K - \beta$

Closure requires the result of a contraction to be a belief set (a closed set of beliefs). The *success* postulate is satisfied if the contracted sentence is not implied by the resulting set, unless it is a tautology. *Inclusion* forbids us to add anything new in the operation. *Vacuity* requires that, if the contracted belief is not in the original belief set, nothing is done. *Recovery*, the most controversial one[4], says that if we add the contracted belief again after the contraction, the beliefs contained in the original belief set must be recovered. Finally, *extensionality* is just saying that equivalent formulas must have the same effect when contracted.

Revisions obtained from contractions respecting these postulates (using the Levi identity) also satisfy a set of postulates (the AGM revision postulates), which we will not mention here.

Other two postulates, known as the supplementary AGM postulates (for contraction), have been proposed: *conjunctive inclusion* and *conjunctive overlap*.

(Conjunctive Inclusion) If $\alpha \notin K - \alpha \wedge \beta$, then $K - \alpha \wedge \beta \subseteq K - \alpha$

(Conjunctive Overlap) $K - \alpha \cap K - \beta \subseteq K - \alpha \wedge \beta$

Conjuctive inclusion says that if when contracting by $\alpha \wedge \beta$ one has to remove α, then all sentences removed by the contraction of α are also to be removed in the contraction by $\alpha \wedge \beta$. Conjunctive overlap assures that a belief that is in $K - \alpha$ and in $K - \beta$ is also in $K - \alpha \wedge \beta$.

---

[3]This happens because the logic used is supraclassical, so it inherits from classical propositional logic the $\bot$e derivation rule (*falsum elimination*, from natural deduction): $\alpha \wedge \neg\alpha \vdash \beta$ for any α, β.

[4]See [Mak87] for clarification.

3.1.3   Partial Meet Contraction

Besides a set of postulates, AGM have proposed a construction for a contraction operation. Before defining it, we will introduce some basic concepts out of which it depends. The first one is the notion of *remainder set.*

Definition 3.1 *[AM82] Let $B \subseteq \mathfrak{L}$ and $\alpha \in \mathfrak{L}$. The remainder set $B\perp\alpha$ is such that $X \in B\perp\alpha$ if and only if:*

- $X \subseteq B$

- $X \nvdash \alpha$

- *For all sets Y, if $X \subset Y \subseteq B$, then $Y \vdash \alpha$*

Intuitively, a remainder set of a set B by a sentence α is the set of maximal subsets of B that do not imply α.

Next, we define *selection function.* A selection function chooses some elements of the remainder set. At least one of the elements must be picked, unless the remainder set is empty.

Definition 3.2 *[AGM85] A function γ is a selection function for the set B if and only if:*

- *If $B\perp\alpha \neq \emptyset$ then $\emptyset \neq \gamma(B\perp\alpha) \subseteq B\perp\alpha$*

- *Otherwise, $\gamma(B\perp\alpha) = \{B\}$*

Now we are in position to define the contraction proposed by AGM, *partial meet contraction.*

Definition 3.3 *[AGM85] Let γ be a selection function for a set of sentences B. The partial meet contraction of B by a sentence α is given by $B - \alpha = \bigcap \gamma(B\perp\alpha)$.*

If the selection function chooses exactly one element of the remainder set, the resulting contraction is called a *maxichoice contraction.* This type of contraction is maximal, in the sense that we cannot add back any removed element of the original set to it and still comply with success. If γ selects all the elements, it gives rise to *full meet contraction.* Full meet contractions are very conservative in a sense, since they just take the elements that are never "involved" in the derivation of the contracted sentence α, they are present in all the maximal subsets not implying α. The in-between cases are just called partial meet contractions.

A maxichoice selection function has very strange effects either in contractions or in revisions generated by them (via Levi identity). The following lemma and corollary ilustrate the unexpected behaviour.

Lemma 3.1 *[AM82] If $\alpha \in K$ and $K - \alpha$ is defined by means of a maxichoice contraction operation, then for any formula β, either $\alpha \vee \beta \in K - \alpha$ or $\alpha \vee \neg\beta \in K - \alpha$.*

Corollary 3.2 *[AM82] If a revision operation is defined from a maxichoice contraction by means of the Levi identity, then, for any α such that $\neg\alpha \in K$, $K * \alpha$ will be maximal, i.e., for every formula β, either $\beta \in K * \alpha$ or $\neg\beta \in K * \alpha$.*

Even meeting all the six AGM postulates, that should guarantee at least some sort of minimality, full meet contractions are too destructive, as can be seen in the following two results.

Lemma 3.3 *[AM82] If $\alpha \in K$ and $K - \alpha$ is defined by means of a full meet contraction operation, then $\beta \in K - \alpha$ if and only if $\beta \in K \cap Cn(\neg\alpha)$.*

Corollary 3.4 *[AM82] If a revision operation is defined from full meet contraction by means of the Levi identity, then, for any α such that $\neg\alpha \in K$, $K * \alpha = Cn(\alpha)$.*

Furthermore, we can define a partial meet operation based on a transitive relation between beliefs.

Definition 3.4 *[AGM85] A selection function γ for K is said to be transitively relational if and only if there is a transitive relation $\leq$ over 2^K such that the following identity holds:*

$$\gamma(K \perp \alpha) = \{X \in K \perp \alpha | X' \leq X \text{ for all } X' \in K \perp \alpha\}$$

Definition 3.5 *[AGM85] A partial meet function is transitively relational if and only if it can be determined by some transitively relational selection function.*

3.1.4   Representation Theorems

A *representation theorem* is a characterization of a mathematical construction. It sets up a list of postulates or axioms and ensures the exchangeability between an operation meeting these axioms and an operation conforming to the mathematical construction. The following representation theorem is one of the most important results of belief revision. It establishes an equivalence between the set of AGM postulates and partial meet contraction.

Theorem 3.5 *[AGM85] An operator $-$ is a partial meet contraction for a belief set K if and only if for all sentences α the operation $K - \alpha$ satisfies closure, success, inclusion, vacuity, recovery and extensionality.*

Another representation theorem, this time for contractions satisfying all basic plus the two suplementary postulates, can be found in the same paper.

Theorem 3.6 *[AGM85] An operator $-$ is a transitively relational partial meet contraction for a belief set K if and only if for all sentences α the operation $K - \alpha$ satisfies closure, success, inclusion, vacuity, recovery, extensionality, conjunctive inclusion and conjunctive overlap.*

3.2   Base Change

Base change is belief revision applied to arbitrary sets, not necessarily closed, the so-called belief bases.

This generalization is useful for computational reasons, as any system which represents knowledge will do it through a finite representation.

A belief base B always has an associated belief set $K = Cn(B)$. For that reason, belief bases are also more expressive than belief sets, since we can have different belief bases "generating" the same belief set. This link creates a distinction among the beliefs an agent holds. There are the *explicit beliefs*, which are the ones in the base, and the *implicit beliefs*, which are derived within the closure[5].

3.2.1   Partial Meet Contraction for Belief Bases

The construction of Definition 3.3 can also be applied to arbitrary sets of sentences (belief bases), in such a way that it satisfies the following set of postulates:

(success) If $\alpha \notin Cn(\emptyset)$, then $\alpha \notin Cn(B - \alpha)$

(inclusion) $B - \alpha \subseteq B$

---

[5]It is important to have a clear philosophical stance on what it means to be an explicit and an implicit belief. For instance, we can regard the explicit beliefs as the beliefs in fact belived by the agent, and the implicit beliefs as mere consequences of those real beliefs, or we can regard both explicit and implicit beliefs as equally important, the separation being just a technical issue.

(relevance) If $\beta \in B \setminus (B - \alpha)$, then there is a B' such that $B - \alpha \subseteq B' \subseteq B$, $\alpha \notin Cn(B')$, but $\alpha \in Cn(B' \cup \{\beta\})$

(uniformity) If for all $B' \subseteq B$, $\alpha \in Cn(B')$ if and only if $\beta \in Cn(B')$, then $B - \alpha = B - \beta$

The postulate of *relevance* says that, if something is removed, it is helping to imply the contracted element. *Uniformity* is similar to extensionality but stronger.

Furthermore, we have the following representation theorem:

Theorem 3.7 *[Han92b] An operator $-$ is a partial meet contraction for a belief base B if and only if for all sentences α the operation $B - \alpha$ satisfies success, inclusion, relevance and uniformity.*

When an operation satisfying all the postulates of Theorem 3.7 is applied over a belief set, it will also satisfy the postulates of Theorem 3.5 [Han99, p. 122-123][6].

3.2.2   Kernel Contraction

A generalization of partial meet contraction, developed in the context of belief bases, is *kernel contraction*, which was defined by Hansson[7]. Instead of picking the intersection of maximal consistent sets, the idea behind kernel contraction is to take the minimal sets implying the element to be contracted and to remove at least one element of each of those sets. Before defining this operation formally, we need to define *kernel sets* and *incision functions*.

Definition 3.6 *[Han94] Let $B \subseteq \mathfrak{L}$ and $\alpha \in \mathfrak{L}$. Then the kernel operation over the set B and the sentence α is the set $B \perp\!\!\!\perp \alpha$, such that $X \in B \perp\!\!\!\perp \alpha$ if and only if:*

- $X \subseteq B$

- $X \vdash \alpha$, *and*

- *If $Y \subset X$, then $Y \nvdash \alpha$.*

The elements of the kernel set are the *α-kernels*. So, in contrast with the notion of remainder sets, an α-kernel of B is a minimal subset of B which implies α.

An incision function is a function that selects at least one element of each α-kernel.

Definition 3.7 *[Han94] An incision function σ for B is a function such that for all α:*

- $\sigma(B \perp\!\!\!\perp \alpha) \subseteq \bigcup(B \perp\!\!\!\perp \alpha)$

- *If $\emptyset \neq X \in B \perp\!\!\!\perp \alpha$, then $X \cap \sigma(B \perp\!\!\!\perp \alpha) \neq \emptyset$*

Now, we are in position to define kernel contraction.

Definition 3.8 *[Han94] Let σ be an incision function for B. The kernel contraction $\approx_\sigma$ for B is defined as follows:*

$$B \approx_\sigma \alpha = B \setminus \sigma(B \perp\!\!\!\perp \alpha)$$

Kernel contraction satisfies the weaker postulate of *core-retainment* (proposed in [Han91]) instead of *relevance*.

(core-retainment) If $\beta \in B \setminus (B - \alpha)$, then there is a B' such that $B' \subseteq B$, $\alpha \notin Cn(B')$, but $\alpha \in Cn(B' \cup \{\beta\})$

Hansson has also proven the representation theorem for kernel contractions.

---

[6]This holds for classical logics, for other cases see [RWFA13].
[7]A generalization of *safe contraction* [AM85].

Theorem 3.8 *[Han94] An operator $\approx_\sigma$ is a kernel contraction for a belief base B if and only if for all sentences α the operation $B \approx_\sigma \alpha$ satisfies success, inclusion, core-retainment and uniformity.*

A desirable property for contractions that kernel contraction unfortunately does not have is relative closure [Han91].

(relative closure) $B \cap Cn(B - \alpha) \subseteq B - \alpha$

This property is a consequence of the postulate of relevance [Han99, p. 71], which kernel contraction does not satisfy. Due to the lack of this property, kernel contraction violates the principle of minimality of belief change [Han99, p. 90]. For instance, consider the logically independent sentences p and q, and let $A = \{p, p \lor q, p \leftrightarrow q\}$. The kernel contraction $A - (p \land q) = \{p\}$ is possible, whereas partial meet contraction cannot have this outcome. As observed by Hansson, it is not sensible to give up $p \lor q$, since p was kept. Nevertheless, a special type of kernel contraction called *smooth kernel contraction* [Han94] satisfies relative closure.

3.2.3   Base-Generated Operators

A base contraction operator gives rise to a theory contraction operator on the base's corresponding belief set [Neb89, Fuh91].

Definition 3.9 *Let $B \subseteq \mathfrak{L}$, $K = Cn(B)$ and $-$ be a contraction operator for B. We say that $\div$ is a contraction operator for K generated from $-$ if for all $\alpha \in \mathfrak{L}$ we have that $K \div \alpha = Cn(B - \alpha)$.*

Any base-generated contraction operation does not satisfy recovery [Han99, p. 307].

3.3   Pseudo-Contractions

Partial meet contraction over belief bases may be too destructive. Consider, for instance, the belief base $B = \{p \land q\}$. If we (partial meet) contract p, the result would be $B - p = \emptyset$. This outcome can be justified if we regard p and q as beliefs obtained together, then, if we give up one, we give up the other. But the outcome $B - p = \{q\}$ is also (and, by intuition, possibly even more) conceivable[8]. This result cannot be obtained due to inclusion.

Hansson has proposed a weakening of the inclusion postulate, called *logical inclusion* [Han89].

(Logical Inclusion) $Cn(B - \alpha) \subseteq Cn(B)$

This author has suggested to call operations satisfying success and logical inclusion *pseudo-contractions* [Han93]. Naturally, all contractions (operations satisfying success and inclusion) are also pseudo-contractions.

3.3.1   Nebel's Pseudo-Contraction

Nebel has proposed a pseudo-contraction for bases (which we will abbreviate as NPC) that generates a contraction which satisfies all the six AGM postulates [Neb89].

Definition 3.10 *Let $\bigwedge B$ be the conjunction of all elements of B. Nebel's pseudo-contraction for the set B is the operator $-$ such that for all sentences α:*

$$B - \alpha = \left\{ \begin{array}{ll} B & \text{if } \alpha \in Cn(\emptyset) \\ \bigcap \gamma(B \perp \alpha) \cup \{\alpha \rightarrow \bigwedge B\} & \text{otherwise} \end{array} \right.$$

---

[8] And here we reveal why it is essential to have a clear philosophical stance on what a belief base is. If we regard each explicit belief as a piece of information obtained at once and independently of other beliefs, then the sentence $p \land q$ of the base was obtained in this form, p and q together, and maybe we should not remove one without removing the other. In any case, we can regard the beliefs in the base as arbitrary codifications of the agent's beliefs, and as such it is only their consequences/content that matter (and not the syntactic form).

In [RW08], Ribeiro and Wassermann have already noted that there is no other intuition behind Nebel's operation than maintaining recovery, a postulate which has been deemed as polemic (as we have already mentioned). These authors, in the same paper, have proposed some reformulations of Nebel's construction. The first (here called NPC$'$) is identic to the above mentioned one, but it adds $\{\alpha \rightarrow \beta | \beta \in B\}$ instead. The second one (NPC$''$) adds $\{\alpha \rightarrow \beta | \beta \in B \setminus \bigcap \gamma(B \perp \alpha)\}$.

3.3.2   General Partial-Meet Pseudo-Contraction (GPMPC)

Although the belief set operation generated from Nebel's pseudo-contraction satisfies all the AGM postulates, it adds unnecessary information to the base. As shown in [RW08], it suffices to add $\{\alpha \rightarrow \bigwedge B'\}$, where $B' = B \setminus \bigcap \gamma(B \perp \alpha)$. Intending to avoid useless additions and yet maintaining some "degree" of recovery, they have proposed a new contraction operator, based on the notion of an *extension* of a selection function.

Definition 3.11 *[RW08] Let γ be a selection function for a set B and let B^* contain B. An extension of γ to B^* is a selection function γ^* such that for every $Y \in \gamma^*(B^* \perp \alpha)$ there is an $X \in \gamma(B \perp \alpha)$ such that $X \subseteq Y$.*

Then, they define the *general partial meet pseudo-contraction*.

Definition 3.12 *[RW08] Let B be a finite belief base, $\alpha \in \mathfrak{L}$ and γ a selection function for B. The general partial meet pseudo-contraction $B - \alpha$ is given by:*

$$B - \alpha = \begin{cases} B & \text{if } \alpha \in Cn(\emptyset) \\ \bigcap \gamma^*(B^* \perp \alpha) & \text{otherwise} \end{cases}$$

where $B^* = B \cup Cn^*(B \setminus \bigcap \gamma(B \perp \alpha))$, γ^* is an extension of γ to B^* and Cn^* is a consequence relation.

In order to axiomatize some notion of minimality in the additions to the base, these authors propose the *core-addition* postulate [RW08], which the last construction does not always satisfy (it depends on the Cn^* used). In the same paper, following the idea of Hansson's logical inclusion, they have also proposed logical relevance, a weakening of relevance[9], and logical recovery, a weakening of recovery.

(Core-Addition) If $\beta \in (B - \alpha) \setminus B$, then there is a $\beta' \in B \setminus (B - \alpha)$ and a $B' \subseteq B - \alpha$ such that $\alpha \rightarrow \beta' \notin Cn(B')$ but $\alpha \rightarrow \beta' \in Cn(B' \cup \{\beta\})$.

(Logical Relevance) If $\beta \in B \setminus (B - \alpha)$, then there is a B' such that $B - \alpha \subseteq B' \subseteq Cn(B)$, $\alpha \notin Cn(B')$, but $\alpha \in Cn(B' \cup \{\beta\})$

(Logical Recovery) $B \subseteq Cn(B - \alpha + \alpha)$

Some properties of this pseudo-contraction can be found in the paper [RW08], but we reproduce the main results here. The next two propositions concern properties of the GPMPC operation itself.

Proposition 3.9 *[RW08] GPMPC satisfies success, logical inclusion, logical vacuity, extensionality and core-retainment.*

Proposition 3.10 *[RW08] If Cn^* satisfies subclassicality, then GPMPC satisfies logical relevance.*

The next two propositions are properties of the operation on belief sets generated by the GPMPC operation.

---

[9]We first thought that a base contraction operation satisfying logical relevance would yield a base-generated operation satisfying relevance, but this is not the case. This would imply a base-generated contraction satisfying recovery, what is impossible. The same argument applies to logical recovery.

Proposition 3.11 *[RW08] The operation on belief sets generated from GPMPC satisfies closure, inclusion, vacuity, success and extensionality.*

In the paper, the authors propose to use $Cn^*(X) = \{\alpha \to \beta | \beta \in X\}$ as a means to obtain a new version of Nebel's construction (NGPMPC, Nebel's GPMPC). This construction is equivalent to NPC$''$.

3.4 Conclusion

So, in this chapter we briefly presented the AGM paradigm, the dominant approach to belief revision. It is a solid framework, with different facets, two of them being the rationality postulates and the partial meet construction, which we exposed here. It is easy to see that only highly idealized agents could implement belief revision in the way the AGM paradigm demands, due to two reasons. The first of them is that the epistemic state is represented by belief sets, very large mathematical entities. If we are dealing with classical logic, these sets contain infinitely many sentences. The second reason is that in order to compute the remainder sets of these infinite entities one must do an enormous amount of work (recall from Chapter 2 the SAT problem, whose instances cannot be solved in polynomial time yet[10]).

After presenting the AGM paradigm, we introduced the topic of base change, the theory of belief change for belief bases. Belief bases are more expressive than belief sets and also more easily representable, since they are not necessarily infinite. Nevertheless, their change tends to be quite destructive, deleting more information than necessary, partly due to the inclusion postulate.

Finally, we presented pseudo-contractions, that in the cases seen so far are also base change operations, but allowing the addition or weakening of formulas, which contradict the inclusion postulate. Nebel's pseudo-contraction, one of the first constructions of this type to be proposed, lacks intuition beyond attaining the recovery postulate, which no base-generated partial meet contraction has. It adds more information than necessary to the base. As a way to fix this, GPMPC was proposed. We start our research trying to better understand this construction, by studying a simplified operation based on it.

---

[10] And never will, if the most accepted conjecture that $P \neq NP$ is correct.

Chapter 4

A New Pseudo-Contraction

In this work we want to further explore the possibility of working with belief bases with logical inclusion, allowing for some syntax independence without having to resort to belief sets. In Section 4.1 we present our proposal, its definition and some basic properties satisfied by it. Section 4.2 has some practical examples of applications to motivate our work. Section 4.3 connects our operation with other pseudo-contraction found in the literature.

Most results and discussion in this chapter are our contribution (unless otherwise mentioned), and part of them (mostly from Sections 4.1 and 4.2) was already published [SRW15].

4.1 Definition and Basic Properties

The direct application of partial meet contraction over closed belief sets and over belief bases creates problems of practical (computational infeasibility) and theoretical (syntax dependence) nature, respectively. One of the aims of this study is to assess the effects of doing the traditional partial meet contraction on belief bases closed by a consequence operation that is between the classical consequence operator and the identity (i.e., the base itself). Hence, we will assume that this operator (here called Cn^*) is Tarskian.

We will study the properties of the application of the partial meet contraction over a set closed under Cn^*, i.e., the operator defined as:

Definition 4.1 *Let B be a set of sentences, Cn^* a consequence relation and γ a selection function for $Cn^*(B)$. The operator $-_*$ is such that, for all sentences α:*

$$B -_* \alpha = \bigcap \gamma(Cn^*(B) \perp \alpha)$$

Notice that $B -_* \alpha = \bigcap \gamma(Cn^*(B) \perp \alpha) = Cn^*(B) -_\gamma \alpha$, where $-_\gamma$ is partial meet contraction. Since $-_\gamma$ satisfies the postulates of success, inclusion, relevance and uniformity, it follows directly (details in the proof just ahead) that $-_*$ satisfies success and the following "starred" versions of inclusion, relevance and uniformity:

(inclusion*) $B - \alpha \subseteq Cn^*(B)$

(relevance*) If $\beta \in Cn^*(B) \setminus (B - \alpha)$, then there is a B' such that $B - \alpha \subseteq B' \subseteq Cn^*(B)$, $\alpha \notin Cn(B')$, but $\alpha \in Cn(B' \cup \{\beta\})$

(uniformity*) If for all $B' \subseteq Cn^*(B)$, $\alpha \in Cn(B')$ if and only if $\beta \in Cn(B')$, then $B - \alpha = B - \beta$

For several applications it is important that the construction satisfies the original success postulate, and not only a starred version of it:

(success*) If $\alpha \notin Cn^*(\emptyset)$, then $\alpha \notin Cn^*(B - \alpha)$

We want that the sentence to be contracted ceases to be logically (classically) implied by the resulting set after the contraction. In this case, the role of the consequence relation Cn^* is just to give a degree of syntactic independence to the operation.

As our purpose here is to make the contraction on a set closed by a Cn^* that does not generate as many consequences as the classic Cn, it is desirable that Cn^* satisfies the property of subclassicality (see Section 2.3).

Clearly, if $Cn^*(A) = A$ (identity), we have that $-_*$ is the usual operation of partial meet contraction on bases. Similarly, for all Tarskian Cn^* that also satisfies subclassicality, applying $-_*$ to belief sets (i.e., $K = Cn^*(K) = Cn(K)$) yields the usual AGM partial meet contraction on belief sets.

Observation 4.1 *If Cn^* satisfies*

- *subclassicality, then an operation satisfying inclusion* also satisfies logical inclusion.*

- *inclusion, then an operation satisfying uniformity* also satisfies uniformity.*

- *subclassicality and inclusion, then an operation satisfying relevance* also satisfies logical relevance.*

The next corollary follows:

Corollary 4.1 *If Cn^* satisfies inclusion and subclassicality, then an operation that satisfies inclusion*, relevance* and uniformity* also satisfies logical inclusion, logical relevance and uniformity.*

With a proof that is very similar to that of the representation theorem for partial meet contraction on bases (which can be found in [Han99]), we can prove the following representation theorem:

Theorem 4.2 *Provided that Cn^* satisfies inclusion, idempotence and subclassicality, an operation is a $-_*$ operator if and only if it satisfies success, inclusion*, relevance* and uniformity*.*

To prove this representation theorem, first, we will reproduce here (with minor adjustments) two properties that will be needed in the proofs.

Observation 4.2 *[AM81](Upper Bound Property) If $X \subseteq A$, and $\alpha \notin Cn(X)$, then there is some X' such that $X \subseteq X' \in A\bot\alpha$.*

Observation 4.3 *[Han99, p. 39, observation 1.39] The following two conditions are equivalent:*

- *$A\bot\beta_1 = A\bot\beta_2$*

- *For all subsets D of A : $\beta_1 \notin Cn(D)$ if and only if $\beta_2 \notin Cn(D)$.*

Now we are ready for the proof.

Proof of Theorem 4.2:
Construction-to-postulates: We know that $A -_* \alpha = \bigcap \gamma(Cn^*(A)\bot\alpha) = Cn^*(A) -_\gamma \alpha$, where $-_\gamma$ is the partial meet contraction. We also know that $-_\gamma$ satisfies success, inclusion, relevance and uniformity. So, we have:

- If $\alpha \notin Cn(\emptyset)$, then $\alpha \notin Cn(Cn^*(A) -_\gamma \alpha)$

- $Cn^*(A) -_\gamma \alpha \subseteq Cn^*(A)$

- If $\beta \in Cn^*(A) \setminus (Cn^*(A) -_\gamma \alpha)$, then there is a B' such that $Cn^*(A) -_\gamma \alpha \subseteq B' \subseteq Cn^*(A)$, $\alpha \notin Cn(B')$, but $\alpha \in Cn(B' \cup \{\beta\})$.

- If for all $B' \subseteq Cn^*(A)$, $\alpha \in Cn(B')$ if and only if $\beta \in Cn(B')$, then $Cn^*(A) -_\gamma \alpha = Cn^*(A) -_\gamma \beta$.

Since $Cn^*(A) -_\gamma \alpha = A -_* \alpha$, we are done.

Postulates-to-construction:

This part is is almost trivially obtained from the proof of the representation theorem for partial meet contraction for bases, which can be found in [Han99, p. 129].

Let $-_*$ be an operation for A that satisfies success, inclusion*, relevance* and uniformity*. From Cn^* satisfying inclusion and subclassicality and corollary 4.1 we conclude that $-_*$ also satisfies logical relevance and uniformity. Let γ be a function such that:

- If $Cn^*(A)\perp\alpha = \emptyset$, then $\gamma(Cn^*(A)\perp\alpha) = \{Cn^*(A)\}$.

- Otherwise $\gamma(Cn^*(A)\perp\alpha) = \{X \in Cn^*(A)\perp\alpha \,|\, A -_* \alpha \subseteq X\}$

We need to show that (1) γ is a well-defined function, (2) γ is a selection function and (3) $\bigcap\gamma(Cn^*(A)\perp\alpha) = A -_* \alpha$ for all α.

Part 1: For γ to be a well-defined function, for all α and β, if $Cn^*(A)\perp\alpha = Cn^*(A)\perp\beta$, we must have $\bigcap\gamma(Cn^*(A)\perp\alpha) = \bigcap\gamma(Cn^*(A)\perp\beta)$. Suppose that $Cn^*(A)\perp\alpha = Cn^*(A)\perp\beta$. It follows from observation 4.3 that any subset of $Cn^*(A)$ implies α if and only if it implies β. By uniformity (of partial meet), $Cn^*(A) -_* \alpha = Cn^*(A) -_* \beta$. By the definition of γ we have $\gamma(Cn^*(Cn^*(A))\perp\alpha) = \gamma(Cn^*(Cn^*(A))\perp\beta)$. By idempotence of Cn^*, the result follows.

Part 2: For γ to be a selection function it remains to be proven that if $Cn^*(A)\perp\alpha$ is not empty, then $\gamma(Cn^*(A)\perp\alpha)$ is not empty as well. Then, assuming $Cn^*(A)\perp\alpha \neq \emptyset$, we know that there is at least one $X \in Cn^*(A)\perp\alpha$, and we must show that at least one of these X contains $A -_* \alpha$. Since $Cn^*(A)\perp\alpha$ is not empty, $\alpha \notin Cn(\emptyset)$, and by success, $\alpha \notin Cn(A -_* \alpha)$. By inclusion*, $A -_* \alpha \subseteq Cn^*(A)$, then, by the subclassicality of Cn^* and the upper bound property (observation 4.2), there is an A' such that $A -_* \alpha \subseteq A'$ and $A' \in Cn^*(A)\perp\alpha$. By the construction of γ, $\gamma(Cn^*(A)\perp\alpha)$ is non-empty.

Part 3: Case 1, $\alpha \in Cn(\emptyset)$. Then, by logical relevance, since there is no A' such that $\alpha \notin Cn(A')$, no element is in $A \setminus A -_* \alpha$, then, using inclusion*, $A \subseteq A -_* \alpha \subseteq Cn^*(A)$. We know that $Cn^*(A)\perp\alpha = \emptyset$, then $\bigcap\gamma(Cn^*(A)\perp\alpha) = Cn^*(A)$. We need to show that $Cn^*(A) \subseteq A -_* \alpha$. By relevance*, we know that $Cn^*(A) \setminus A -_* \alpha = \emptyset$, then $Cn^*(A) \subseteq A -_* \alpha$.

Case 2, $\alpha \notin Cn(\emptyset)$. $Cn^*(A)\perp\alpha$ is non-empty and by part 2, $\gamma(Cn^*(A)\perp\alpha)$ is non-empty as well. Since $A -_* \alpha$ is a subset of all elements of $\gamma(Cn^*(A)\perp\alpha)$, $A -_* \alpha \subseteq \bigcap\gamma(Cn^*(A)\perp\alpha)$. We need to show that $\bigcap\gamma(Cn^*(A)\perp\alpha) \subseteq A -_* \alpha$.

Take $\varepsilon \notin A -_* \alpha$. If $\varepsilon \notin Cn^*(A)$, obviously $\varepsilon \notin \bigcap\gamma(Cn^*(A)\perp\alpha)$. If $\varepsilon \in Cn^*(A) \setminus A -_* \alpha$, then by relevance* there is an A' such that $A -_* \alpha \subseteq A' \subseteq Cn^*(A)$, $\alpha \notin Cn(A')$ but $\alpha \in Cn(A' \cup \{\varepsilon\})$. It follows from the upper bound property that there is an A'' such that $A \subseteq A''$ and $A'' \in Cn^*(A)\perp\alpha$. From $A \subseteq A''$, $\alpha \in Cn(A' \cup \varepsilon)$ and $\varepsilon \in A''$ we conclude that $\alpha \in Cn(A'')$, so we must have $\varepsilon \notin A''$. By our definition of γ, $A'' \in \gamma(Cn^*(A)\perp\alpha)$, and since $\varepsilon \notin A''$, we conclude that $\varepsilon \notin \bigcap\gamma(Cn^*(A)\perp\alpha)$, concluding the proof. ∎

From this theorem and the previous corollary, it also follows:

Corollary 4.3 *If Cn^* satisfies inclusion, idempotence and subclassicality, then $-_*$ satisfies success, logical inclusion, logical relevance and uniformity.*

It is interesting that we have here a set of postulates that are independent from Cn^*. Nonetheless, these postulates do not characterize the operation, and are in general weaker than the postulates with *. Hansson's logical inclusion postulate is quite reasonable for base operations, as it brings syntactic independence, although with inclusion* we already have a degree of independence, and with better preservation of the original set (since Cn^* is subclassical), and, depending on the chosen Cn^*, we avoid the complexity problem we have with the closure of Cn.

Another desirable property in rational contraction operations is relative closure (cf. Section 3.2). This property is a consequence of the postulate of relevance, which $-_*$ does not satisfy. Nevertheless, relative closure is satisfied, given the condition that Cn^* satisfies inclusion.

Proposition 4.4 *If Cn^* satisfies inclusion, the $-_*$ operator satisfies relative closure.*

Proof: We know that $A -_* \alpha = Cn^*(A) -_\gamma \alpha$, where $-_\gamma$ is the partial meet contraction. Since partial meet satisfies relative closure [Han99, p. 71], $Cn^*(A) \cap Cn(Cn^*(A) -_\gamma \alpha) \subseteq Cn^*(A) -_\gamma \alpha$ is valid. From this we have $Cn^*(A) \cap Cn(A -_* \alpha) \subseteq A -_* \alpha$. By the inclusion property of Cn^* and set theory we get $A \cap Cn(A -_* \alpha) \subseteq Cn^*(A) \cap Cn(A -_* \alpha)$ and hence $A \cap Cn(A -_* \alpha) \subseteq A -_* \alpha$. ∎

As seen in Section 3.2, kernel contraction is an alternative construction for contraction, which is characterized by the same postulates as partial meet contraction on bases, except for relevance, which is weakened to core-retainment.

Kernel contraction may have erratic behaviour due to its non-satisfaction of relevance, as mentioned in Chapter 3. Would the weakening of relevance to relevance* or logical relevance be enough so as to make these behaviours show up in the $-_*$ operator? Kernel contraction does not satisfy any of these last two postulates. Furthermore, Hansson had already noticed that the lack of relative closure also contributes to these unnecessary removals in contraction. The operator $-_*$, as shown above, satisfies relative closure.

A property of our pseudo-contraction is enforced closure*.

(enforced closure*) $B - \alpha = Cn^*(B - \alpha)$

Proposition 4.5 *If Cn^* is Tarskian and satisfies subclassicality, an operator that satisfies inclusion* *and relevance* *also satisfies enforced closure*.*

Proof: Since Cn^* is Tarskian, by inclusion, $A - \alpha \subseteq Cn^*(A - \alpha)$. We want to show that $Cn^*(A - \alpha) \subseteq A - \alpha$. Suppose by contradiction that $\beta \in Cn^*(A - \alpha) \setminus (A - \alpha)$. From inclusion*, monotonicity and idempotence of Cn^* we obtain $\beta \in Cn^*(A) \setminus A - \alpha$. Relevance* guarantees that there is an A' such that $A - \alpha \subseteq A' \subseteq Cn^*(A)$, $\alpha \notin Cn(A')$ but $\alpha \in Cn(A' \cup \{\beta\})$. By subclassicality of Cn^* and $\beta \in Cn^*(A - \alpha)$ we have $\beta \in Cn(A - \alpha)$. By $A - \alpha \subseteq A'$ and by the inclusion property of Cn, $\beta \in Cn(A')$. So, we have $Cn(A') = Cn(A' \cup \{\beta\})$, which is a contradiction. ∎

Whenever $A \subset Cn^*(A)$, this postulate will imply that vacuity is not satisfied, which is an essential postulate from the point of view of rational contractions (that respect the principle of minimal change). It has as effect that the belief base will always end up closed by Cn^* after the contraction, even though the original base was not closed.

Nevertheless, our construction does satisfy a weaker form of vacuity:

(vacuity*) If $\alpha \notin Cn(B)$, then $B - \alpha = Cn^*(B)$

Proposition 4.6 *If Cn^* satisfies subclassicality, an operator that satisfies inclusion* *and relevance* *also satisfies vacuity*.*

Proof: Assume $\alpha \notin Cn(B)$.

By inclusion* we already have $B - \alpha \subseteq Cn^*(B)$. To finish the proof it is sufficient to prove that $Cn^*(B) \setminus (B - \alpha) = \emptyset$.

By relevance*, if $\beta \in Cn^*(B) \setminus (B - \alpha)$, then there is a B' such that $B - \alpha \subseteq B' \subseteq Cn^*(B)$ and $\alpha \in Cn(B' \cup \{\beta\})$. From this and subclassicality of Cn^* we get $B' \cup \{\beta\} \subseteq Cn^*(B) \subseteq Cn(B)$ and then by monotonicity and idempotence of Cn we have $Cn(B' \cup \{\beta\}) \subseteq Cn(B)$. Since $\alpha \notin Cn(B)$ by assumption, we cannot have such α. ∎

Corollary 4.7 *If Cn^* is Tarskian and satisfies subclassicality, then the operator $-_*$ satisfies enforced closure* and vacuity*.*

A simple way to restore vacuity is to redefine the $-_*$ operator in the following manner:

Definition 4.2 *Let B be a set of sentences, Cn^* a function from sets of sentences to sets of sentences and γ a selection function for $Cn^*(B)$. The operator $-'_*$ is such that, for all sentence α:*

$$B -'_* \alpha = \left\{ \begin{array}{ll} B & \textit{if } \alpha \notin Cn(B) \\ \bigcap \gamma(Cn^*(B) \bot \alpha) & \textit{otherwise} \end{array} \right.$$

Observation 4.4 *The $-'_*$ operator satisfies success, inclusion*, uniformity* and vacuity. If Cn^* satisfies inclusion and subclassicality, $-'_*$ also satisfies logical inclusion, uniformity and relative closure.*

The proof of the observation above is not given, but can be trivially obtained from Theorem 4.2, Corollary 4.1 and Proposition 4.4. Notice that relevance* and logical relevance were lost and although we have attained vacuity, we have only partly gotten rid of enforced closure* (just when $\alpha \notin Cn(B)$).

If on one hand logical inclusion seems to make more sense than inclusion for base contractions, by allowing some syntactic independence, effects such as enforced closure* illustrate the need to refrain from careless additions of sentences in the contraction. Here we should recall the postulate of core-addition [RW08], seen in Section 3.3.

Any operator satisfying inclusion will satisfy this postulate trivially. If we break $\{\alpha \to \bigwedge B\}$ into the set of sentences $\{\alpha \to \beta \mid \beta \in B\}$, Nebel's pseudo-contraction will not satisfy core-addition. Clearly the $-_*$ operator does not satisfy it also (neither does $-'_*$), and it does not satisfy vacuity as well. In the effort to fix these two problems, the operations of general partial-meet pseudo-contraction and Δ-partial-meet pseudo-contraction, proposed in [RW08], seem to be viable solutions.

4.2   Examples and Applications

In this Section we are going to show some concrete examples of Cn^* functions that can be useful in the solution of practical problems. The first example we mention is the Cleopatra example, adapted from [Han99].

Example 4.1 *Consider a language with three propositional letters, p, q and r and a belief base $B = \{p \wedge q\}$, where p stands for* Cleopatra had a son *and q,* Cleopatra had a daughter. *If we want to contract by p, applying a partial meet contraction produces $B - p = \emptyset$. This is not always the expected result, because the loss of faith in the belief that Cleopatra had a son also made us lose faith in the belief that she had a daughter.*

With the classic partial meet construction for bases we would have $B \bot p = \{\emptyset\}$, so the selection function needs to choose $\{\emptyset\}$, causing the overall contraction process to produce $\emptyset$ as final result.

On the other hand, if we take B to represent the belief set $K = Cn(B)$, then K contains both p and q and the belief that Cleopatra had a daughter (q) may survive the contraction, i.e., we may have $q \in K - p$. But then we would also have $p \vee r$, $r \to p$ and many other irrelevant formulas in the resulting set, since it is closed under Cn.

Let us consider an intermediate consequence operator:

$Cn_1^*(A) = \{\alpha \mid \alpha \in A$ or for any formulas β, δ, $\alpha \wedge \beta \in A$ or $\beta \wedge \alpha \in A$ or $\beta \wedge \alpha \wedge \delta \in A\}$

This operator is "breaking the conjunctions into conjuncts". We can use Cn_1^* with the $-_*$ operator to solve the problem of the preceding example. In this case, we have $Cn_1^*(B) = \{p \wedge q, p, q\}$ and $Cn_1^*(B) \bot p = \{\{q\}\}$, hence the selection function would choose the whole remainder set, $\{\{q\}\}$, and, accordingly, $B -_* p = \{q\}$.

Although the usefulness of this consequence operation is dubious, its use already brings better results than the typical base contraction in some cases, as in the former example.

Example 4.2 *Suppose I believe that the town of Juazeiro do Norte is located in the state of Pernambuco $(j \to p)$ and that the state of Pernambuco is located in Brazil $(p \to b)$. Speaking with a colleague, I found that this town is not located in his state (Pernambuco), that is, I contract $j \to p$ from my base. The outcome is $B - (j \to p) = \{p \to b\}$. So, I no longer know whether Juazeiro do Norte is located in Brazil.*

In this example, as well as in the previous one, one can blame the poor codification of the belief base for the problems. The knowledge that Juazeiro do Norte is located in Brazil, if obvious, perhaps would be individually justified, and so it would deserve to be explicitly in the base. At this point the syntactic independence dilemma reappears. In some cases we want to have it, but without having to generate infinitely many derivative sentences with little utility.

However, when working with ontologies, for instance, it is possible that the user does not want to make explicit every possible relationship, trusting the transitivity of some properties (i.e., he would be more concerned with his ontology on the knowledge level than on the syntactic level). One may also want a knowledge base with little redundancy.

Again, neither the belief base nor the belief set approach would give us the desired result.

Returning to the foregoing example, we could use a Cn_2^* that adds to the base the result of one application of the transitivity of $\to$:

$$Cn_2^*(A) = A \cup \{\alpha_1 \to \alpha_2 | \text{ for some } \beta, \alpha_1 \to \beta, \beta \to \alpha_2 \in A\}.$$

In that case, we would have $Cn_2^*(B) = B \cup \{j \to b\}$, which results in $Cn_2^*(B) \perp (j \to p) = \{\{p \to b, j \to b\}\}$. As in the last example, the selection function must choose the only member of the remainder set, therefore $B -_* (j \to p) = \{p \to b, j \to b\}$. It is interesting to note that we could also use here a $Cn_{2'}^*(B)$ being the set of all Horn consequences of B.

The following example was adapted from [Han93].

Example 4.3 *Suppose I believe, for good and independent reasons, that Andy is son of the mayor (a) and Bob is son of the mayor (b). Then I hear the mayor say: "I certainly have nothing against our youth studying abroad. My only son did it for three years". I then have to retract $a \wedge b$ from my base $B = \{a, b\}$. But it is reasonable to retain a belief that either Andy or Bob is the son of the mayor, i.e., the result of the contraction should be $\{a \vee b\}$.*

The remainder set for the operation above is $B \perp (a \wedge b) = \{\{a\}, \{b\}\}$. So, the resulting partial meet contraction is either $\{a\}$, $\{b\}$ or $\emptyset$, the first two being odd since we do not seem to have reasons to prefer a over b or vice-versa.

In the same paper where he presented the example above, Hansson has done an extensive study of partial meet contraction on disjunctively closed bases. If we define $Cn_3^*(A)$ as the disjunctive closure of A, as defined by Hansson, that is, $Cn_3^*(A)$ is the set of sentences that are either elements of A or disjunctions of elements of A, we can manage to get the desired result.

$$Cn_3^*(A) = A \cup \{\bigvee \alpha_i | \alpha_i \in A\}$$

We have $Cn_3^*(B) = \{a, b, a \vee b\}$. Then, the remainder set is $Cn_3^*(B) \perp (a \wedge b) = \{\{a, a \vee b\}, \{b, a \vee b\}\}$, and so the selection function may choose both sets, producing the expected result in the lack of evidence for a or b: $B -_* (a \wedge b) = \{a \vee b\}$.

Consider now the case where the language has three propositional letters (a, b, and c). If we take the belief set $K = Cn(B)$, we have that K contains $a \vee c$ and $b \vee c$. It is not hard to see that there are two remainder sets containing these two formulas, namely $\{a, a \vee b, a \vee c, b \vee c\}$ and $\{b, a \vee b, a \vee c, b \vee c\}$, and hence, they may survive contraction, even if the original set did not mention c. Again, this shows the inconvenience of using classical closure.

4.3   Connections with GPMPC

We now show the connections between general partial-meet pseudo-contraction and our new pseudo-contraction $-_*$ (that was conceived as a simplification of the former). We should keep in mind that we turned back to this more sophisticated construction as a way to prevent a problem (or feature, if you will) of the $-_*$ operator, which we mentioned in the end of Section 4.1: its violation of minimality, which manifests itself in the form of enforced closure* and the failure to satisfy vacuity (two properties which, in turn, are closely related) and core-addition. The construction of GPMPC makes it easier to reach core-addition, but we will postpone the discussion about this postulate to chapter 5.

First we define a two-place selection function.

Definition 4.3 *[Han99, p. 105] A two-place selection function γ is a function which maps sets of sentences (B) into suitable selection functions (γ_B for B).*

If the context is clear, we will omit the subscript of γ_B, writing just γ. Now, we can redefine the operator from Ribeiro and Wassermann as:

Definition 4.4 *Let γ be a two-place selection function. The new general partial-meet pseudo-contraction operator $-$ is an operator that, for all sets B and all sentences α:*

$$B - \alpha = \bigcap \gamma(Cn^*_{\alpha,\gamma}(B) \bot \alpha)$$

where $Cn^*_{\alpha,\gamma}(B) = B \cup f_\alpha(B \setminus \bigcap \gamma(B \bot \alpha))$, and $f_\alpha : 2^{\mathfrak{L}} \to 2^{\mathfrak{L}}$.

The intuition behind the function f_α comes from [RW08]. Their construction was designed as a pseudo-contraction, which will expand the original base with some consequences of the elements that would usually be removed by a partial meet contraction (represented by the set $B \setminus \bigcap \gamma(B \bot \alpha)$). These consequences are drawn according to a consequence relation, intended to be weaker than the classical Cn, here denoted by this f_α. The idea is that some formulas that would be removed by partial meet contraction could just be weakened, and this weakening is done by extracting some consequences with f_α and making the traditional partial meet operation afterwards.

Comparing to Ribeiro and Wassermann's operator, we eliminated the exception clause for the case when α is a tautology, because it was redundant in the case when $f_\alpha(\emptyset) = \emptyset$ (which is a sensible requirement on f_α). The set B^* was replaced by $Cn^*_{\alpha,\gamma}$ so as to make explicit that we are dealing with a "weak closure" of the original set (as in our $-_*$ operation) and that this consequence relation depends both on the contracted sentence and on the selection function. Note that since γ is a two-place selection function, the two appearances of γ in the definition above are not referring to the same function. The first is actually $\gamma_{Cn^*_{\alpha,\gamma}(B)}$, a selection function for $Cn^*_{\alpha,\gamma}(B)$, and the second is γ_B, a selection function for B[1].

We replaced the extension of a selection function γ^* from the original definition by the more general concept of a two-place selection function. We can then verify in what situations the relationship between $\gamma_{Cn^*_{\alpha,\gamma}(B)}$ and γ_B satisfies a property that makes the former equivalent to an extension of the latter (recall definition 3.11).

We cannot regard this (redefined) general partial-meet pseudo-contraction as a special case of the $-_*$ operator, because its consequence relation depends on both the belief base and the contracted element. Still, we want to investigate the conditions under which $Cn^*_{\alpha,\gamma}$ satisfies inclusion, monotonicity, idempotence and subclassicality, so we could apply some of the results showed before to that operation as well. However, before working with this new definition, we will present some further properties of GPMPC (as it was originally defined), besides the ones already listed in Section 3.3.

---

[1]This definition of $\gamma_{Cn^*_{\alpha,\gamma}(B)}$ may seem circular but it is not, since $Cn^*_{\alpha,\gamma}(B)$ depends only on γ_B, not on γ itself. In fact, we should have written $Cn^*_{\alpha,\gamma_B}(B)$ but the definition is already too complicated, so we are going to leave it as it is for the sake of simplicity.

4.3.1   Properties of GPMPC

Basic properties of GPMPC are enumerated in Section 3.3. In [RW08] they claim the following observation:

Observation 4.5 *[RW08] Let γ be a selection function for B and let B^* contain B. If γ^* is an extension of γ to B^*, then for any $X \in \gamma(B \bot \alpha)$ there is a $Y \in \gamma^*(B^* \bot \alpha)$ such that $X \subseteq Y$.*

Nonetheless, if one wants this property, it should also be part of the definition of an extension of a selection function, because the following proposition holds true:

Proposition 4.8 *If γ^* is defined according to definition 3.11, then observation 4.5 is not valid.*

Proof: We show a counterexample. Let $A = \{p, q, r\}$. So $A \bot (p \wedge q) = \{\{p, r\}, \{q, r\}\}$. Let $A^* = A$, so $A^* \bot (p \wedge q) = A \bot (p \wedge q)$. If $\gamma(A \bot (p \wedge q)) = A \bot (p \wedge q)$, then for $\gamma^*(A^* \bot (p \wedge q)) = \{p, r\}$, γ^* satisfies the definition but not the observation. ∎

The following lemma is an intermediary result we are going to use to prove another property of GPMPC.

Lemma 4.9 *Let $A^* = A \cup X$, γ be a selection function for A and γ^* be an extension of γ to A^* (according to definition 3.11). If $\bigcap \gamma^*(A^* \bot \alpha) \subseteq A$, then $\bigcap \gamma^*(A^* \bot \alpha) = \bigcap \gamma(A \bot \alpha)$.*

Proof: If $\varepsilon \in \bigcap \gamma(A \bot \alpha)$, then $\varepsilon \in X$ for all $X \in \gamma(A \bot \alpha)$. Since, by definition of γ^*, for all $Y \in \gamma^*(A^* \bot \alpha)$, there is an $X \in \gamma(A \bot \alpha)$ such that $X \subseteq Y$, we conclude that $\varepsilon \in Y$ for all $Y \in \gamma^*(A^* \bot \alpha)$, therefore $\bigcap \gamma(A \bot \alpha) \subseteq \bigcap \gamma^*(A^* \bot \alpha)$.

For the other direction of the proof, let us assume $\bigcap \gamma^*(A^* \bot \alpha) \not\subseteq \bigcap \gamma(A \bot \alpha)$. So, there must be an $\varepsilon \in \bigcap \gamma^*(A^* \bot \alpha)$ such that $\varepsilon \notin \bigcap \gamma(A \bot \alpha)$. Thus, $\varepsilon \in Y$ for all $Y \in \bigcap \gamma^*(A^* \bot \alpha)$, but $\varepsilon \notin X$ for some $X \in \gamma(A \bot \alpha)$, accordingly, for all those Y's and some of those X's, $\varepsilon \in Y \setminus X$. Since $X \in A \bot \alpha$, for any $\varepsilon \in A \setminus X$ we have $X \cup \{\varepsilon\} \vdash \alpha$, and since $Y \nvdash \alpha$, it must be the case that $\varepsilon \in A^* \setminus A$. But we know that $\bigcap \gamma^*(A^* \bot \alpha) \subseteq A$, therefore this ε cannot exist and so we come to a contradiction. ∎

As an immediate result we have the next corollary:

Corollary 4.10 *Let $A^* = A$, γ be a selection function for A and γ^* be an extension of γ to A^*. Then $\bigcap \gamma^*(A^* \bot \alpha) = \bigcap \gamma(A \bot \alpha)$.*

This corollary is basically saying that any extension of a selection function to the same set will always produce the same partial meet contraction (notice that we have not even used observation 4.5).

Now we can easily prove the seemingly obvious property that follows:

Proposition 4.11 *If a GPMPC operation satisfies inclusion, then it satisfies relevance.*

This proposition is not as obvious as it seems. Respecting inclusion does not mean that B^* is the same as B. The inclusion property refers to the overall result of the operation, not to the consequence relation that generates B^*. So it could be the case that an operation added elements to B^* but they were lost in the contraction afterwards. Likewise, we do not know much about the behaviour of γ^* in relation to γ, for instance, even if $B^* = B$ it could be the case that the resulting contractions were different (actually, now we know that this is impossible, thanks to the previous corollary).

Proof of Proposition 4.11: Let $-$ be a GPMPC and $-_\gamma$ be a partial meet contraction. If α is a tautology, $B - \alpha = B = B -_\gamma \alpha$, so let us consider the main case when $\alpha \notin Cn(\emptyset)$.

Since inclusion holds, $\bigcap \gamma^*(B^* \bot \alpha) \subseteq B$, and thus, by lemma 4.9, $\bigcap \gamma^*(B^* \bot \alpha) = \bigcap \gamma(B \bot \alpha)$. We know relevance holds for partial meet contraction, so it must hold for GPMPC in this case as well.   ∎

Now we prove some additional properties of GPMPC.

Proposition 4.12 *GPMPC does not satisfy uniformity (in general).*

Proof: We give a counterexample.

$$A = \{p, q, r\} \qquad\qquad A^* = A \cup \{p \vee x\}$$
$$A \bot p = \{\{q, r\}\} \qquad\qquad A^* \bot p = \{\{q, r, p \vee x\}\}$$
$$A \bot (p \vee x) = \{\{q, r\}\} \quad A^* \bot (p \vee x) = \{\{q, r\}\}$$

For all $A' \subseteq A$, $A' \vdash p$ if and only if $A' \vdash p \vee x$, but $\bigcap \gamma^*(A^* \bot p) = A - p \neq A - (p \vee x) = \bigcap \gamma^*(A^* \bot (p \vee x))$.   ∎

Proposition 4.13 *GPMPC satisfies relative closure.*

Proof: In the case when the contracted element α is a tautology the result follows trivially (and, again, in this case GPMPC behaves exactly like partial meet contraction, which satisfies relative closure). So let us examine the main case, where α is not a tautology.

Let $\beta \notin A - \alpha$. We want to show that $\beta \notin A \cap Cn(A - \alpha)$. If $\beta \notin A$ the result follows trivially, so let us consider the case where $\beta \in A$. We want $\beta \notin Cn(A - \alpha)$. Suppose to the contrary that $\beta \in Cn(A - \alpha)$. Then $\bigcap \gamma^*(A^* \bot \alpha) \vdash \beta$, so by compactness of Cn there is a subset B of $\bigcap \gamma^*(A^* \bot \alpha)$ such that $B \vdash \beta$, and thus all elements of $\gamma^*(A^* \bot \alpha)$ contain B. Since $\beta \notin \bigcap \gamma^*(A^* \bot \alpha)$, then there is some $Y \in \gamma^*(A^* \bot \alpha)$ such that $\beta \notin Y$. By the definition of an extension of a selection function, since γ^* is an extension of γ, there is some $X \in \gamma(A \bot \alpha)$ such that $X \subseteq Y$ and, of course, $\beta \notin X$. Since X is a maximal subset of A not implying α, we cannot add any element of A to it without implying α. $\beta \in A \setminus X$, so $X \cup \{\beta\} \vdash \alpha$. Since Y contains B that implies β and also contains X, Y implies α, which is a contradiction.   ∎

Proposition 4.14 *GPMPC does not satisfy core-addition (in general).*

Proof: Let $A = \{p \wedge q\}$, $A \bot p = \emptyset$. Let $A^* = \{p \wedge q, q \vee x\}$, so $A^* \bot p = \{q \vee x\}$. The sentence $q \vee x$ was legally added to A^* but $\{q \vee x\} \cup \{p\} \nvdash p \wedge q$.   ∎

4.3.2   Relating the Operations

We depart from definition 4.4 and try to connect this operation with our previously defined $-_*$ operation. Unfortunately, we cannot say that the New GPMPC (from definition 4.4) is a special case of our $-_*$ pseudo-contraction (as we initially thought), because the consequence relation of the New GPMPC depends on both the belief base and the input sentence, whereas the consequence relation of our pseudo-contraction disregards the input sentence. In any case, we want to know when the consequence relation $Cn^*_{\alpha, \gamma}$ has inclusion, idempotence, monotonicity and subclassicality, so we can verify which of the results from Section 4.1 apply to this construction.

Now we prove that $Cn^*_{\alpha, \gamma}$ is Tarskian and subclassical under some conditions.

Proposition 4.15 $Cn^*_{\alpha, \gamma}$ *satisfies inclusion.*

Proof: Follows directly from the definition of $Cn^*_{\alpha, \gamma}$.   ∎

The following condition, which is divided into two independent parts, is necessary for the next proofs. The plausibility of this condition will be discussed just ahead.

Definition 4.5 *Let γ be a two-place selection function. We say that γ satisfies* Selection Linkage *(*part 1 *and* part 2*) if and only if the following conditions are (respectively) met. For all A, A', if $A \subseteq A'$, then, for all α:*

  1. and for all $X \in \gamma(A \bot \alpha)$ there is an $X' \in \gamma(A' \bot \alpha)$ such that $X \subseteq X'$.

  2. and for all $X' \in \gamma(A' \bot \alpha)$ there is an $X \in \gamma(A \bot \alpha)$ such that $X \subseteq X'$.

Proposition 4.16 *If f_α is monotonic and γ satisfies part 1 of Selection Linkage, then $Cn^*_{\alpha,\gamma}$ satisfies monotonicity.*

Proof: Let $A' = A \cup B$ and $A \cap B = \emptyset$. If $B = \emptyset$ then $A' = A$ and so $Cn^*_{\alpha,\gamma}(A') = Cn^*_{\alpha,\gamma}(A)$. So let $B \neq \emptyset$ and then $A \subset A'$. Therefore, $Cn^*_{\alpha,\gamma}(A) \subseteq Cn^*_{\alpha,\gamma}(A')$ if and only if $A \cup f_\alpha(A \setminus \bigcap \gamma(A \bot \alpha)) \subseteq A \cup f_\alpha(A' \setminus \bigcap \gamma(A' \bot \alpha)) \cup B$. The latter is true if (but not only if) $f_\alpha(A \setminus \bigcap \gamma(A \bot \alpha)) \subseteq f_\alpha(A' \setminus \bigcap \gamma(A' \bot \alpha))$. Since f_α is monotonic, we just need $A \setminus \bigcap \gamma(A \bot \alpha) \subseteq A' \setminus \bigcap \gamma(A' \bot \alpha)$, which is true if $\bigcap \gamma(A' \bot \alpha) \setminus \bigcap \gamma(A \bot \alpha) \subseteq B$, by set theory. So, we want that, if it is the case that for all $X' \in \gamma(A' \bot \alpha)$ and some $X \in \gamma(A \bot \alpha)$ we have $\beta \in X'$ and $\beta \notin X$, then $\beta \in B$.

Suppose $\beta \notin B$, $\beta \in X'$ for all $X' \in \gamma(A' \bot \alpha)$ and $\beta \notin X$ for some $X \in \gamma(A \bot \alpha)$. By part 1 of Selection Linkage, we know that for all $Z \in \gamma(A \bot \alpha)$ there is a $Z' \in \gamma(A' \bot \alpha)$ such that $Z \subseteq Z'$. If $Z = Z'$ and $\beta \notin Z$, then $\beta \notin Z'$. Otherwise $Z \subset Z'$, and if $\beta \notin Z$, since $\beta \notin B$, $\beta \in A$. But since $Z \subset Z'$, $Z' \nvdash \alpha$ and $\beta \in Z'$, $Z \cup \{\beta\} \nvdash \alpha$, which is a contradiction since $Z \in A \bot \alpha$. ∎

The next condition is used in the proof of idempotence. It seems a bit puzzling at first, but the idea behind it is quite simple. We are going to discuss its intuition and plausibility just after the proofs.

Definition 4.6 *(Condition 1) For all $B \subseteq \mathfrak{L}$ and all $\alpha \in \mathfrak{L}$, $Cn^*_{\alpha,\gamma}$ and γ are such that $Cn^*_{\alpha,\gamma}(B) \setminus B \subseteq Cn^*_{\alpha,\gamma}(B) - \alpha \cup D$, where D is some set such that $f_\alpha(X) \setminus f_\alpha(X \setminus D) \subseteq B$ for all $X \subseteq Cn^*_{\alpha,\gamma}(B)$.*

Proposition 4.17 *Let f_α be monotonic, Condition 1 hold and γ be a two-place selection function such that for all $X' \in \gamma(Cn^*_{\alpha,\gamma}(B) \bot \alpha)$ there is an $X \in \gamma(B \bot \alpha)$ such that $X \subseteq X'$[2]. Then, $Cn^*_{\alpha,\gamma}$ satisfies idempotence.*

Proof: Let $B_0 = B$ and $B_{i+1} = Cn^*_{\alpha,\gamma}(B_i) = B_i \cup f_\alpha(B_i \setminus \bigcap \gamma(B_i \bot \alpha))$ for all $i \geq 0$. We want $Cn^*_{\alpha,\gamma}(B) = Cn^*_{\alpha,\gamma}(Cn^*_{\alpha,\gamma}(B))$, that is, $B_i = B_{i+1}$ for all $i \geq 1$, or, equivalently, $B_1 = B_2$. So we want $B_1 = B_0 \cup f_\alpha(B_0 \setminus \bigcap \gamma(B_0 \bot \alpha)) = B_1 \cup f_\alpha(B_1 \setminus \bigcap \gamma(B_1 \bot \alpha)) = B_2$, i.e., $f_\alpha(B_1 \setminus B_1 - \alpha) \subseteq B \cup f_\alpha(B \setminus B - \alpha)$, where $-$ is partial meet contraction. From the conditions on γ, we get $B - \alpha \subseteq B_1 - \alpha$.

Again, we want $f_\alpha(B_1 \setminus B_1 - \alpha) \subseteq B \cup f_\alpha(B \setminus B - \alpha)$. Since f_α is monotonic, $B_1 \setminus (B_1 - \alpha) \subseteq B \setminus (B - \alpha)$ would be enough, but it is not necessary since some elements of $f_\alpha(B_1 \setminus B_1 - \alpha)$ can be in B, so we just need $(B_1 \setminus (B_1 - \alpha)) \setminus D \subseteq B \setminus (B - \alpha)$, or $B_1 \setminus (B_1 - \alpha) \subseteq B \setminus (B - \alpha) \cup D$, where $f_\alpha(B_1 \setminus (B_1 - \alpha)) \setminus f_\alpha((B_1 \setminus (B_1 - \alpha)) \setminus D) \subseteq B$. We already have $B - \alpha \subseteq B_1 - \alpha$, so, by set theory, it only remains for us to prove that $B_1 \setminus B \subseteq B_1 - \alpha \cup D$, but this is granted by Condition 1. ∎

Before proceeding to the discussion about Selection Linkage and Condition 1, we will prove that $Cn^*_{\alpha,\gamma}$ can be subclassical.

Proposition 4.18 *If $f_\alpha(B \setminus (B - \alpha)) \subseteq Cn(B)$ or f_α satisfies subclassicality and monotonicity, then $Cn^*_{\alpha,\gamma}$ satisfies subclassicality.*

Proof: If $f_\alpha(B \setminus (B - \alpha)) \subseteq Cn(B)$, the result follows directly from the construction of $Cn^*_{\alpha,\gamma}$. In the other case, since $B \setminus (B - \alpha) \subseteq B$, by monotonicity and subclassicality of f_α we get $f_\alpha(B \setminus (B - \alpha)) \subseteq f_\alpha(B) \subseteq Cn(B)$. ∎

---

[2] The last condition mentioned could have been replaced by the stronger Selection Linkage part 2.

At this moment we have all sufficient conditions to make $Cn_{\alpha,\gamma}^*$ Tarskian and subclassical, which allow us to use most results from Section 4.1 on the New GPMPC.

Now, we are going to talk a bit about the two conditions defined earlier in order to prove the Tarskian properties of $Cn_{\alpha,\gamma}^*$. First, we will tackle Condition 1 (definition 4.6) via two examples.

We can slightly modify Nebel's pseudo-contraction (see definition 3.10), redefining it as a New GPMPC, where $Cn_{\alpha,\gamma}^*(B) = B \cup f_\alpha(B \setminus \bigcap \gamma(B \bot \alpha))$, and $f_\alpha(X) = \{\alpha \to \beta | \beta \in X\}$. In this case we have that the included sentences $\alpha \to \beta$ will never help to imply the removed sentence α, so all elements added, i.e., $Cn_{\alpha,\gamma}^*(B) \setminus B$, will survive the contraction (by relevance of partial meet contraction), and so they will be in $Cn_{\alpha,\gamma}^*(B) - \alpha$. Therefore, this construction will satisfy Condition 1, and the set D (referred to in Condition 1) is not even required in this case.

For constructions such as $f_\alpha(X) = \{\beta | \beta$ is a conjunct in some $\varepsilon \in X\}$, we have the property that $f_\alpha(f_\alpha(X)) = \emptyset$ for any X, since once the conjunctions are broken, no more consequences are produced (we are considering, for example, that $f_\alpha(\{p \wedge q\}) = \{p, q\}$, but $f_\alpha(\{p, q\}) = \emptyset$). So, the consequences, $Cn_{\alpha,\gamma}^*(B) \setminus B$, will constitute the set D, where $f_\alpha(X) = f_\alpha(X \setminus D)$ for any X, since the freshly created conjuncts cannot be broken any further, which means that the set D does not influence the outcomes of f_α in any way. Thus, in this example the construction satisfies Condition 1 as well.

Condition 1 establishes that all conclusions generated by $Cn_{\alpha,\gamma}^*$ must be in one of three possible categories. The first category is made up of sentences which survive the contraction by α. The second is the group of sentences which do not influence the results of f_α. And the third, the set of sentences which are already in B. The two examples above produce consequences in the first and second category, respectively. Although a pair of examples are scarcely sufficient to prove that Condition 1 is reasonable, it is illustrative of its plausibility, and other similar instances of f_α can be built which will also cause the consequences to be in the foregoing categories. Therefore, there are reasonable scenarios where $Cn_{\alpha,\gamma}^*$ is idempotent.

Now for the discussion on Selection Linkage. The next results show that Selection Linkage is a too strong assumption.

Proposition 4.19 *If Selection Linkage part 1 holds, then γ is full meet.*

Proof: For all sets such that $B \bot \alpha = \{B\}$ we must have $\gamma(B \bot \alpha) = \{B\}$. These sets do not imply α. So, for any set B' we will have $B' \bot \alpha = \{B_1, B_2, ..., B_n\}$, where each B_i is selected in $B_i \bot \alpha$, i.e., $B_i \in \gamma(B_i \bot \alpha)$. So, by Selection Linkage part 1, they all should be selected, and thus $\gamma(B \bot \alpha) = B \bot \alpha$ for any B. ∎

Proposition 4.20 *Selection Linkage part 2 is impossible.*

Proof: Take a set B such that $B \bot \alpha = \{B_1, B_2\}$. Some subset of $B \bot \alpha$ must be selected. If B_1 is selected, then, since $B_2 \subseteq B$ and $B_2 \bot \alpha = \{B_2\}$, we must have by Selection Linkage part 2 that $B_2 \subseteq B_1$, what is impossible since $\{B_1, B_2\} \subseteq B \bot \alpha$. Analogously, the same happens if B_2 is selected. ∎

Corollary 4.21 *Selection Linkage is impossible.*

So, only Selection Linkage part 1 is possible, and even that requires quite strong conditions (full meet selection function). Fortunately, to prove Proposition 4.17, we can use another condition (which is possible, at least) instead of Selection Linkage part 2, as shown in the proof.

These last results are interesting not just with regard to the New GPMPC but also with respect to the subject of (two-place) selection functions.

4.3.3   Concluding Remarks

The New GPMPC is not a particular case of the $-_*$ operator, because its consequence relation considers the input sentence in addition to the belief base. It is a particular case of a *two-place* $-_*$ operator, whose consequence relation depends not only on the belief base but also on the input sentence. The conditions for $Cn^*_{\alpha,\gamma}$ to be Tarskian are very strong, specially for monotonicity. Nevertheless, it turns out that the only result of Section 4.1 that needs monotonicity of Cn^* is Proposition 4.5. $Cn^*_{\alpha,\gamma}$ inclusion is granted and idempotence holds on reasonable conditions.

New GPMPC satisfies vacuity, as a result of $Cn^*_{\alpha,\gamma}(B) = B$ whenever $\alpha \notin Cn(B)$, assuming $f_\alpha(\emptyset) = \emptyset$. Comparing to GPMPC, the New GPMPC has gained uniformity (if $Cn^*_{\alpha,\gamma}$ satisfies inclusion).

Usually, in belief revision, change operators are defined only for a single set (the other case is called "iterated revision"). If the New GPMPC operator is defined only for a single set B (as is the case of GPMPC in [RW08]), we can define $Cn^*_{\alpha,\gamma}$ only for B as well. Since $Cn^*_{\alpha,\gamma}$ depends on γ_B, which is defined only for B, it is reasonable to define $Cn^*_{\alpha,\gamma}$ only for B too. So, if we consider the domain of f_α as being $2^{\mathfrak{L}}$, f_α can be a partial function undefined for any $X \subseteq \mathfrak{L}$, except for $B \setminus \bigcap \gamma(B \bot \alpha)$. Thus, $f_\alpha(B \setminus \bigcap \gamma(B \bot \alpha))$ can take any value, say $h_\alpha(B)$, without worrying if $B \setminus \bigcap \gamma(B \bot \alpha) = B' \setminus \bigcap \gamma(B' \bot \alpha)$ for some B' (in which case we could not assign arbitrary values for f_α with these arguments, the values should be the same since f_α is a function). So, $Cn^*_{\alpha,\gamma} = B \cup h_\alpha(B)$. Since $h_\alpha(B)$ can be anything we want, $Cn^*_{\alpha,\gamma}$ is any operator of consequence with inclusion. This is quite astonishing, for it converts the New GPMPC operator into the two-place $-_*$ operator with a Cn^* that respects inclusion. It is desirable that $Cn^*_{\alpha,\gamma}$ satisfies inclusion, so, in the most typical case, the New GPMPC is equivalent to the two-place $-_*$, and not just a specific case of it, as was thought (and stated) before. We can say this version of New GPMPC is "not global", by virtue of it being applicable only for the belief base B (which is, again, the most common case in belief revision).

The limitation that γ^* must be an extension of γ is not a problem too (in regard to the equivalence between the operations) because, as long as $f_\alpha(B \setminus \bigcap \gamma(B \bot \alpha))$ can be anything we desire, we can define γ^* as we wish and make a compatible γ.

The vacuity clause (when $\alpha \in Cn(\emptyset)$) is not necessary if $f_\alpha(\emptyset) = \emptyset$, so, even if it is kept, it will not affect the result.

If we make $Cn^*_{\alpha,\gamma}$ and f_α total functions, though, we do not have the equivalence anymore. But this is not really necessary since we are not making a global operator. Although, the independence between the functions and the set B can be useful for practical reasons. Moreover, this equivalence is more of a theoretical nature, and just reveals that the expressive power of both constructions is the same. By this we are claiming (without formal proof) that the following observation is valid:

Observation 4.6 *Provided that Cn^* and $Cn^*_{\alpha,\gamma}$ satisfy inclusion, a pseudo-contraction can be constructed as a two-place $-_*$ operation if and only if it can be constructed as a (not global) New GPMPC.*

Chapter 5

Minimality and the Postulates

This small chapter, a sequel of the work developed in the previous one, is a discussion concerning the postulates, mainly about how they could encode minimality, but also about success.

First, in Section 5.1, we talk about general principles behind belief contraction. Then, in Section 5.2, we discuss the core-addition postulate. Next, in Section 5.3, we discuss the success postulate.

5.1 Guiding Principles of Contractions

We have seen that a contraction can have various possible outcomes. What are the guiding principles which guarantee that one contraction is more rational than another? If we go through the AGM postulates for contraction (closure, success, inclusion, vacuity, recovery and extensionality) we can acknowledge some of these principles.

Maybe success is the key one. If we contract a non-tautological belief α and it remains there after contraction, we can barely call this a contraction of α[1].

Closure is somewhat "technical", and it is useful mostly because we are dealing with ideal agents in AGM paradigm. However, as noted by Hansson in [Han99, p. 69], a weaker version of closure, relative closure, is important to preserve the coherence of contractions. It seems illogical to remove a consequence of the contracted set if this consequence was already in the original set.

Extensionality tries to bring sanity to the behaviour of contractions across different situations. For example, when we contract two formulas that say the same thing via different syntactic representations, the result should be the same (in other words, syntactic independence is desirable). However, by the same argument, if I had two belief sets B and B' and I contracted equivalent formulas, should not I had the same effect over the "relevant parts" of B and B' (supposing this relevant part of both sets is the same)? Extensionality does not avoid this, since contractions are usually defined just for one belief set.

The other three postulates have the same underlying principle: minimality of changes. If we are going to change our beliefs due to the removal of α, this change must be just the minimal acceptable change (i.e., a minimal one that respects the other postulates).

Vacuity requires minimality of change (no change, indeed) when the contracted α is not in the belief set. Inclusion requires minimality with respect to additions: nothing is added during contraction. Recovery tries to guarantee minimality in removals. So little is removed that if we add α back we have again all the knowledge we had before. It is as though just the information that comes with α has been removed.

Since we are working with pseudo-contractions, we replace inclusion by logical inclusion, what makes sense for belief bases. What we now want is to strengthen this minimality of additions, and we can do this with the postulate of core-addition.

---

[1]In any event, success is not above criticism (even in this chapter we present some issues about it).

5.2   Core-Addition

The core-addition postulate (Section 3.3.2) has some problems. Even if adding $\{\alpha_1, \alpha_2, ..., \alpha_n\}$ to a contraction causes it to violate core-addition, it is possible that adding only the sentence $\bigwedge_{i=1}^{n} \alpha_i$ does not. The problem resides in the fact that some conjuncts in $\bigwedge_{i=1}^{n} \alpha_i$ may not be necessary to imply some removed element, while the entire sentence is. So, the postulate is sensible to syntactic variations. We must be very emphatic in the minimalism demanded by core-addition, since we are already allowing additions (contrary to AGM, base change postulates and minimality principle in general), they must really be very small, minimal ideally. With this ambition, we propose two adjustments in the postulate.

We do not want that an added element (say β) only helps in deriving some removed element (say β') for *some* subset B' of $B - \alpha$. We want it to be mandatory to the derivation of this β' inside $B - \alpha$. Which means that there is some $\beta' \in B \setminus (B - \alpha)$ for which $\beta' \notin Cn(B - \alpha \cup \{\alpha\} \setminus \{\beta\})$, that is, β' is not implied by the contraction expanded by α without β, but it is implied with the help of β, or $\beta' \in Cn(B - \alpha \cup \{\alpha\})$. In other words, β' would not be recovered without β. This leads us to the following modification of core-addition:

(Core-Addition′) If $\beta \in (B - \alpha) \setminus B$, then there is some $\beta' \in B \setminus (B - \alpha)$ such that $\beta' \notin Cn(B - \alpha \cup \{\alpha\} \setminus \{\beta\})$ but $\beta' \in Cn(B - \alpha \cup \{\alpha\})$

Furthermore, if β was added because it is necessary to imply some removed β', no formula ε weaker than β could be used to replace β with this purpose. By weaker we mean a formula which is implied by β but does not imply it, $\varepsilon \in Cn(\beta)$, $\beta \notin Cn(\varepsilon)$. As shown in observation 5.1 ahead, this is the same as $Cn(\varepsilon) \subset Cn(\beta)$, and in a way says that "ε contains less information than β". So we can modify core-addition in order to demand that the added elements be as weak as possible:

(Core-Addition″) If $\beta \in (B - \alpha) \setminus B$, then there is some $\beta' \in B \setminus (B - \alpha)$ and some $B' \subseteq (B - \alpha)$ such that for all ε such that $Cn(\varepsilon) \subset Cn(\beta)$, we have that $\beta' \notin Cn(B' \cup \{\alpha, \varepsilon\})$ and $\beta' \in Cn(B' \cup \{\alpha, \beta\})$.

So, for example, if $a \wedge b$ was included, then a is not enough to derive some β' removed in contraction (in presence of the contracted α and the rest of the base).

Combining both modifications, we get the following:

(Core-Addition‴) If $\beta \in (B - \alpha) \setminus B$, then there is some $\beta' \in B \setminus (B - \alpha)$ such that and for all ε such that $Cn(\varepsilon) \subset Cn(\beta)$ we have that $\beta' \notin Cn(B - \alpha \cup \{\alpha, \varepsilon\} \setminus \{\beta\})$ but $\beta' \in Cn(B - \alpha \cup \{\alpha\})$.

We can easily see that the satisfaction of core-addition‴ implies the satisfaction of core-addition′ and core-addition″, and that the satisfaction of each of those two implies the satisfaction of the original postulate. The implications do not hold in the opposite direction.

In the way Nebel's Pseudo-Contraction is formulated (see Section 3.3), it does not even satisfy the original core-addition (in the vacuity case: if the contracted element is not in the base, a sentence is added anyway), so it does not satisfy any of the postulates proposed here. Nebel's GPMPC (NGPMPC) satisfies core-addition, but does not satisfy core-addition′ nor core-addition″.

The Section "Minimal Additions" in [RW08] starts with an argument in favor of core-addition, namely, that it is possible to have logical recovery without adding all sentences that NGPMPC adds. They gave an example. Consider the base $\{p \wedge \neg r, p \wedge \neg r \wedge q\}$ and the contraction by p. Both formulas are given up, but we only need $\alpha \rightarrow p \wedge \neg r \wedge q$ to maintain logical recovery. It turns out that the original core-addition is so weak that it is not enough to prevent the problem of this example. Clearly, a contraction respecting core-addition′ will not accept the addition of $\alpha \rightarrow p \wedge \neg r$ if $\alpha \rightarrow p \wedge \neg r \wedge q$ has already been added in this example.

A problem (or feature, if you will) of core-addition″ can be illustrated with the following example. Consider the contraction $\{p \wedge q\} - p$. Core-addition″ allows the result $p \rightarrow q$, but forbids q (which, intuitively, seems very plausible).

Some versions of core-addition can be viewed as a strengthening of logical inclusion, but some can not. The results follow.

Proposition 5.1 *An operation satisfying core-addition does not necessarily satisfy logical inclusion.*

Proof: We can show a counterexample.

Take the base $B = \{p \wedge q\}$ and the contraction $B - q = \{q \wedge x\}$. This operation satisfy core-addition but does not satisfy logical inclusion. ∎

Proposition 5.2 *An operation satisfying core-addition' does not necessarily satisfy logical inclusion.*

Proof: The proof is identical to the previous one. ∎

Proposition 5.3 *For a logical language closed under the boolean operators, any operation satisfying core-additon'' will also satisfy logical inclusion.*

Proof: Consider $\beta \in (B - \alpha) \setminus B$. Then, there is some $\beta' \in B \setminus (B - \alpha)$ and some $B' \subseteq (B - \alpha)$ such that for all ε such that $Cn(\varepsilon) \subset Cn(\beta)$, we have that $\beta' \notin Cn(B' \cup \{\alpha, \varepsilon\})$ and $\beta' \in Cn(B' \cup \{\alpha, \beta\})$. (We will refer to all these conditions on β as "the conditions".) We want to show that $\beta \in Cn(B)$.

Take $\varepsilon = \alpha \to \beta'$. Since $\beta' \in B$, then $\varepsilon \in Cn(B)$. This ε is such that $\beta' \in Cn(B' \cup \{\alpha, \varepsilon\})$. If there is an ε weaker than that satisfying $\beta' \in Cn(B' \cup \{\alpha, \varepsilon\})$, then we can take it instead. We can repeat this step until we find the weakest ε' satisfying the above conditions, and it holds that $\varepsilon' \in Cn(B)$.

Suppose there is a ε'' for which the above conditions hold, but $Cn(\varepsilon'') \not\subset Cn(\varepsilon')$ (otherwise we should have chosen ε'' instead of ε') and $Cn(\varepsilon') \not\subset Cn(\varepsilon'')$ (otherwise we just pick ε'). We are supposing there is no ε weaker than ε'' for which the conditions hold. This would be the only way ε'' could satisfy the conditions without being inside $Cn(B)$. But, if both ε' and ε'' are weakest formulas satisfying the conditions (in the sense that there is no formula weaker than them), we can take their disjuction, which is weaker than them, and it would satisfy the conditions. This is a contradiction since we assumed ε' and ε'' had no weaker formulas satisfying the conditions. ∎

Corollary 5.4 *For classical propositional logic, any operation satisfying core-additon''' will also satisfy logical inclusion.*

Observation 5.1 $\alpha \in Cn(\beta)$ *and* $\beta \notin Cn(\alpha)$ *if and only if* $Cn(\alpha) \subset Cn(\beta)$.

Proof: If $\{\alpha\} \subset Cn(\beta)$, by monotony and idempotence of Cn, $Cn(\alpha) \subset Cn(\beta)$. If $Cn(\beta) \subseteq Cn(\alpha)$, by inclusion, $\beta \in Cn(\alpha)$, so we conclude that $Cn(\beta) \not\subseteq Cn(\alpha)$, i.e., $Cn(\alpha) \subset Cn(\beta)$.

For the other direction, by inclusion, $\alpha \in Cn(\alpha)$, and since $Cn(\alpha) \subset Cn(\beta)$, $\alpha \in Cn(\beta)$. If $\beta \in Cn(\alpha)$, by monotony and idempotence, $Cn(\beta) \subseteq Cn(\alpha)$, contradicting $Cn(\beta) \not\subseteq Cn(\alpha)$. ∎

5.3   Success

The success postulate is reasonable in the context of perfect reasoners, as thought by AGM, but we can cast doubt upon it for *resource-bounded reasoners* (that is, reasoners that may not be able to realize all the consequences of its beliefs). We are going to show some examples.

Example 5.1 *I believe some sentence α. For some reason, I note that α is not necessarily true, I doubt it, so I contract α from my set of beliefs. Still, my knowledge base may have beliefs that together imply α, without my noticing, and afterwards I may derive α again.*

The above example seems perfectly sane for a human reasoner, for instance.

Example 5.2 *We have a theory which comprises α. We contract α. This does not mean we belief α is false (in that case we would revise the base with $\neg\alpha$), we just do not know about α. We do not have to seek and remove anything from our theory that imply α. Maybe it is hidden somewhere, we are just unsure of it, so we remove it.*

Of course, the above example makes no sense for perfect reasoners, since the closure would readily add the contracted belief.

As remarked by Hansson [Han99, p. 64], it is difficult to find an example of *pure* contraction[2]. These two examples have something to do with this. If instead of just contracting some sentence by doubting it we had revised our base by the negation of it, then the total elimination of the sentence and the beliefs that imply it would be mandatory (otherwise we end up with an inconsistent base). But revision, as defined by the Levi Identity, if contraction was like the above examples, would be inconsistent. Thus, perhaps we should allow inconsistencies, or perhaps revision should be defined in another way (independent of contraction, maybe). If we are thinking about resource-bounded reasoners, allowing inconsistencies makes sense, since humans, for example, do not have an inconsistency-free set of beliefs.

It is strange that, in many cases, the examples used to justify the reasonableness of AGM postulates are intuitions about human beings and our reasoning, after all, we are not perfect reasoners.

Maybe it would be interesting to have a *naive contraction* operation, that just removes α (much like an opposite of expansion), and another *purge* operation, that removes what is needed to avoid α being forced to be true (the contractions as we know them).

We can have "degrees of success". Borrowing a lot from the philosophical outlook of [Was00][3], one can devise a stepwise consequence operation, which Wassermann calls Inf. We can compute success over Inf^* (all steps of Inf) when dealing with perfect reasoners, and over some Inf^n for resource-bounded reasoners. The bigger the n, the smarter the agent. Thus, the naive contraction could use 0-success ($\alpha \notin Inf^0(B) = B$) and purge could use ∞-success ($\alpha \notin Inf^*(B)$, where $Inf^*(B)$ approaches $Cn(B)$).

The use of two different types of logical consequence (one for closure and one for success) is one of the fundamental differences of this work. The consideration of success under weaker logics was made before by Wassermann [Was03], where just the "relevant" beliefs were taken into account when contracting by some formula. The agent then could continue having this belief implied, but she can not see it, so she believes that contraction was successful.

---

[2]Although, in computer science, removal from logical databases can be viewed as pure contractions.

[3]In which we see a belief revision approach for resource-bounded agents.

Chapter 6

Application Examples

In this chapter we are going to present two application examples where our Cn^* closure and our pseudo-contraction can be used as a theoretical grounding. The first, in Section 6.1, demonstrates how our weak closure is implicitly present in popular inference engines for ontologies. The second, in Section 6.2, shows how our pseudo-contraction can be applied to change of probabilistic bases.

6.1 Pseudo-contractions in Ontology Engineering

Ontologies[1] are a particularly useful realization of knowledge representation (and of the millennial philosophical idea of ontology, which dates back to Parmenides in Ancient Greece). They are used to explicitly specify some body of knowledge, designating a set of concepts and their properties, as well as relationships between those concepts. Ontologies have a wide variety of applications today, many of them related to the semantic web.

Those ontologies are for the most part internally represented through a logical language, which is usually a description logic[2]. One of the most basic description logics is $\mathcal{ALC}$. Sentences of $\mathcal{ALC}$ can be translated into sentences of first order logic[3].

Ontology engineering (the process of creating and maintaining ontologies) is often done with the aid of ontology editing software, such as Protégé[4]. With the assistance of this kind of software, the user can specify an ontology and then query it, using some query language, like SPARQL [PS06, PAG06]. Commonly, the ontology is recorded as a belief base, with just the information explicitly provided by the user. So, the user can only use SPARQL to query this "belief base". In order to obtain some of the logical consequences of her ontology, the user is required to employ some automated reasoner, which will transform her dataset into something which resembles our weak closure. Some widely known reasoners used for this purpose are HermiT[5], Pellet[6] and FaCT++[7]. We are going to examine some aspects of the inference done by HermiT, that seems to outperform its contestants in most cases [SMH08] (although the inference done by the other reasoning engines are probably very similar, the main differences lie in the implementation details and performance). Since we have not found any formal specification of the inference done by HermiT, we will define this inference based on observations of the behaviour of this reasoner, and assume what is more plausible.

---

[1]Since our intention here is just to depict the usage of our theory, detailed information on ontologies is beyond the scope of this dissertation. For an introduction on ontologies see [Gru93, Gru95, NM01].

[2]A good introduction on description logics and the language $\mathcal{ALC}$ can be found in [LMP08, BHS08].

[3]A translation method is in [BHS08, p. 144].

[4]http://protege.stanford.edu/

[5]http://www.hermit-reasoner.com/

[6]https://github.com/complexible/pellet

[7]http://owl.man.ac.uk/factplusplus/

6.1.1   The Logic $\mathcal{ALC}$

Before proceeding, we reproduce here some definitions about the language $\mathcal{ALC}$. First, we define its syntax:

Definition 6.1 *[BHS08] Let N_C be a set of concept names and N_R be a set of role names. The set of $\mathcal{ALC}$-concept descriptions is the smallest set such that*

- *$\top$,$\bot$, and every concept name $A \in N_C$ is an $\mathcal{ALC}$-concept description,*

- *if C and D are $\mathcal{ALC}$-concept descriptions and $r \in N_R$, then $C \sqcap D$, $C \sqcup D$, $\neg C$, $\forall r.C$, and $\exists r.C$ are $\mathcal{ALC}$-concept descriptions.*

The syntax of $\mathcal{ALC}$ accepts as valid concept descriptions: $\top$ (*top*) and $\bot$ (*bottom*), which are the concepts containing all and no members of the domain, respectively; concept names; intersection and union of two simpler concepts; negation of a concept; universal and existential quantification of a role with a concept (for example, $\forall child.Female$ is a concept whose members only have female children).

The following definition shows the semantics of $\mathcal{ALC}$.

Definition 6.2 *[BHS08] An interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ consists of a non-empty set $\Delta^{\mathcal{I}}$, called the domain of $\mathcal{I}$, and a function $\cdot^{\mathcal{I}}$ that maps every $\mathcal{ALC}$-concept to a subset of $\Delta^{\mathcal{I}}$, and every role name to a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ such that, for all $\mathcal{ALC}$-concepts C, D and the role names r,*

- *$\top^{\mathcal{I}} = \Delta^{\mathcal{I}}$, $\bot^{\mathcal{I}} = \emptyset$,*

- *$(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$, $(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cup D^{\mathcal{I}}$, $\neg C^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$,*

- *$\exists r.C = \{x \in \Delta^{\mathcal{I}} |$ There is some $y \in \Delta^{\mathcal{I}}$ with $(x, y) \in r^{\mathcal{I}}$ and $y \in C^{\mathcal{I}}\}$,*

- *$\forall r.C = \{x \in \Delta^{\mathcal{I}} |$ For all $y \in \Delta^{\mathcal{I}}$, if $(x, y) \in r^{\mathcal{I}}$, then $y \in C^{\mathcal{I}}\}$*

In other words, the interpretation contains a domain, and a function which maps every concept into a subset of the domain and roles into binary relations between elements of the domain. The interpretation of intersection/union between concepts is straightforward: it is given by the intersection/union between the subsets of the domain which represent these concepts. The negation of a concept is interpreted as the elements of the domain which are not in the interpretation of this concept.

Now, we define axioms of concept inclusion and TBox.

Definition 6.3 *[BHS08] A general concept inclusion (GCI) is of the form $C \sqsubseteq D$, where C, D are $\mathcal{ALC}$-concepts. A finite set of GCIs is called a TBox. An interpretation $\mathcal{I}$ is a model of a GCI $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$; $\mathcal{I}$ is a model of a TBox $\mathcal{T}$ if it is a model of every GCI in $\mathcal{T}$.*

A TBox is a set of concept inclusions, which are, in turn, sentences asserting that all elements of a concept are also elements of another.

Next, we define assertional axioms and ABox.

Definition 6.4 *[BHS08] An assertional axiom is of the form $x : C$ or $(x, y) : r$, where C is an $\mathcal{ALC}$-concept, r is a role name, and x and y are individual names. A finite set of assertional axioms is called an ABox. An interpretation $\mathcal{I}$ is a model of an assertional axiom $x : C$ if $x^{\mathcal{I}} \in C^{\mathcal{I}}$, and $\mathcal{I}$ is a model of an assertional axiom $(x, y) : r$ if $(x^{\mathcal{I}}, y^{\mathcal{I}}) \in r^{\mathcal{I}}$; $\mathcal{I}$ is a model of an ABox $\mathcal{A}$ if it is a model of every axiom in $\mathcal{A}$.*

An ABox is a set of assertional axioms, which are sentences saying that an individual is a member of a concept or that two individuals are related through a certain role.

Finally, we define knowledge base.

Definition 6.5 *[BHS08] A knowledge base (KB) is a pair $(\mathcal{T}, \mathcal{A})$, where $\mathcal{T}$ is a TBox and $\mathcal{A}$ is an ABox. An interpretation is a model of a KB if it is a model of its TBox and of its ABox.*

We say that $KB \models \alpha$ if and only if every model of KB is also a model of α, where α is an element of a TBox or of an ABox.

6.1.2   Analysing the Consequence Relation

Based on the observed behaviour of HermiT inside the Protégé environment, and making some plausible assumptions, we presume that its consequence relation Cn_H, given an $\mathcal{ALC}$ ontology (or knowledge base) KB, generates an inferred ontology $Cn_H(KB)$ that is the smallest set for which the following properties hold:

- If an axiom is in KB, it is in $Cn_H(KB)$.

- Applying it twice is the same as applying it once, that is, $Cn_H(KB) = Cn_H(Cn_H(KB))$

- For all pairs of user defined concepts C, D, if $KB \models C \sqsubseteq D$, then $C \sqsubseteq D \in Cn_H(KB)$

- For all user defined individuals a and concepts C, if $KB \models a : C$, then $a : C \in Cn_H(KB)$.

- For all user defined individuals x, y and roles r, if $KB \models (x, y) : r$, then $(x, y) : r \in Cn_H(KB)$.

The first two observations warrant that Cn_H satisfies inclusion and idempotence, respectively. It is easy to see that if one of the three last rules are used to generate a consequence over KB, this same consequence would be generated for any KB' such that $KB \subseteq KB'$, since the $\mathcal{ALC}$ consequence relation $Cn_{\mathcal{ALC}}$[8] is monotonic. Thus, Cn_H is also monotonic.

Because all the consequences generated by the five rules above are also consequences of $Cn_{\mathcal{ALC}}$, we can say that Cn_H satisfies the following:

$$\text{For all } A, Cn_H(A) \subseteq Cn_{\mathcal{ALC}}(A)$$

Since $Cn_{\mathcal{ALC}}$ is subclassical[9], Cn_H is too. So we have a Cn_H that satisfies inclusion, idempotence, monotonicity and subclassicality, what makes it a perfect example of our weak closure Cn^* described in the previous chapters. Therefore, all those properties hold for Cn_H as well, and a pseudo-contraction based on it would have the aforementioned properties.

The importance of this result lies in the fact that we gave some theoretical background for an operation that is already being used a lot in practice.

6.2   Pseudo-contractions in Probabilistic Belief Revision

In the work [dBFR$^+$16], we deal with consolidation (the operation of turning an inconsistent belief base into a consistent one) of probabilistic knowledge bases. (Under the AGM paradigm, one can think of consolidation simply as the contraction by *falsum*.) We work with a logical language composed by sentences of the form $P(\phi|\psi) \geq q$, called a probabilistic conditional, which means that "the probability of ϕ being true given that ψ is true is at least q", where ϕ and ψ are formulas of classical propositional logic and q is a real number in the interval $[0, 1]$.

In the AGM paradigm, contraction (and consequently, consolidation) is done via the removal of formulas. However, in probabilistic logic, another way of making a belief base consistent is possible, via the changing of probabilities. When probability bounds are changed, AGM postulate of inclusion is violated. When we decrease the value of the probability bound of a sentence, what we are really doing is a weakening of this formula.

---

[8] $\alpha \in Cn_{\mathcal{ALC}}(KB)$ if and only if $KB \models \alpha$, where $\models$ is defined as in the previous subsection.
[9] Considering the translation to FOL.

The weakening of formulas was made possible by what is known as pseudo-contractions, subject thoroughly discussed in this dissertation. As we already know, this weakening is just the addition of formulas implied by the original set of beliefs, which is done before contraction. The problem is that there is little clue on how to choose the consequences, problem aggravated in the case of consolidation by the fact that the consequences of an inconsistent set is the whole language. To avoid this situation, one possibility is to use a subclassical consequence instead of the classical one[10], operation described in Chapter 4 of this dissertation. This was the approach taken by the authors.

We propose an element-wise consequence relation, Cn_{ew}, given by the following expressions:

- For all $\Gamma \in \mathbb{K}$, $Cn_{ew}(\Gamma) = \bigcup\{Cn_{ew}(\{\alpha\})|\alpha \in \Gamma\}$

- For all $\alpha = P(\phi|\psi) \geq q \in \mathcal{L}$, $Cn_{ew}(\{\alpha\}) = \bigcup\{P(\phi|\psi) \geq q'|q' \in [0,1], q' \leq q\}$

where $\mathbb{K}$ is the set of all probabilistic knowledge bases and $\mathcal{L}$ is the language.

In that paper, a type of consolidation called maximal consolidation is studied. This consolidation is one wherein the contradiction is removed solely through the decreasing of some probabilities, in a way that the distance between the original array of probabilities and the new consistent one is minimal, according to some measure of distance. Interestingly enough, taking Cn_{ew} as Cn^*, maximal consolidations always respect inclusion*, but relevance* is violated. If the result of a maximal consolidation is closed under Cn^* afterwards, relevance* is recovered.

We go on with the goal of characterizing a maximal consolidation. We come to the conclusion that an operation is the Cn^* closure of a maximal consolidation if and only if it respects success, inclusion* and fullness* (again, taking Cn_{ew} as Cn^*).

(Fullness*) If $\beta \in Cn^*(B) \setminus (B - \alpha)$, then $\alpha \notin Cn(B - \alpha)$ and $\alpha \in Cn((B - \alpha) \cup \{\beta\})$

Notice that fullness* is a strengthening of relevance*. Its non-starred version is one of the postulates that characterize maxichoice contraction. Uniformity* was not required in the former characterization because consolidations are always contractions by $\perp$.

---

[10]In fact, the term subclassical is not adequate here, since their standard consequence relation Cn is the one from probabilistic logic.

Chapter 7

Related Work

In this chapter we will make short surveys of articles that are not necessarily a basis for our work, but similar to it in design or intention.

7.1 Degrees of Recovery and Inclusion

This subject has been covered in the paper [RW08], which was the primary source of inspiration for this master's dissertation. We have already described parts of it in Section 3.3.

The general idea of that paper is relaxing the postulate of inclusion to obtain some "degree" of recovery. Doing so, they define some pseudo-contractions, one of which is General Partial Meet Pseudo-Contraction (GPMPC) which was further (and deeply) explored here. As this pseudo-contraction is rather complicated, we have studied in this project the properties of a pseudo-contraction that seems like a simplified version of GPMPC.

7.2 Disjunctively Closed Bases

In the paper [Han93], Hansson develops one of the closest proposals to the one described here. The author takes the disjunctive closure of the belief base before making the contraction, and compares the results. The disjunctive closure adds all disjunctions of elements of the set to it (of course we must not add equivalent disjunctions if we want the set to keep finite).

Our approach is a generalization of Hansson's, we replace disjuctive closure by a general consequence operation.

7.3 Local Change

In [HW02], the authors propose the use of local change operators. When a base is revised or contracted by some belief, for instance, the base is not affected in its entirety, only parts of the base are *relevant* for the target belief. A local operator looks only at this relevant chunk of the base to make the changes.

In the course of their text, they define very general belief revision operators, not necessarily local, and prove some results and representation theorems for them. They have found properties of partial meet based on an inference operation C which is required only to be compact and monotonic.

The main difference from this general partial meet to our pseudo-contraction regards the success postulate. Whereas they define belief change operations for any logic whose consequence relation respects some properties, here we use a general consequence relation just for computing the closure. Thus, in our case, success is still verified against classical entailment.

7.4   Approximate Belief Revision

In [CPW01], the authors propose the use of approximate entailment as defined by [SC95]. One of the approximate entailment relations, denoted by $\vDash^1_S$, is unsound and complete, the other, $\vDash^3_S$, is sound but incomplete. Since we have that if $B \vDash^3_S \alpha$ then $B \vDash \alpha$ and that if $B \nvDash^1_S \alpha$ then $B \nvDash \alpha$, we can use these relations to approximate $\vDash$ and $\nvDash$, respectively. The context is a set $S \subseteq L$, where L is the set of atoms of B (also known as the language of B).

They propose a strategy to define the context set S, based on a notion of relevance between atoms. They define a sketch of an algorithm that improves the approximation iteratively. The extreme cases happen when $S = \emptyset$, and when $S = L$, in which case we have classical entailment.

They show that we can define consequence relations based on these entailments, and that they satisfy a set of properties that allow them to be used together with the axiomatizations of local change [HW02] to implement belief revision.

Our approach differs from this in the same way as it does from the local change approach, since this is a practical realization of local change theory of belief revision.

7.5   Infobase Pseudo-Contraction

A different approach on pseudo-contractions was taken by [MLH00] and further developed in [Mey01]. The authors replace the formulas that cannot be retained with weakened formulas ("replacing" actually means including new formulas in the base, so they are dealing with pseudo-contractions). They call their bases infobases (a belief base in which all beliefs have independent standing).

They use a pre-order on U (the set of all interpretations of $\mathfrak{L}$). The contraction for belief sets obtained from this pre-order satisfies the 8 AGM postulates (the six basic ones plus the two supplementary postulates). This pre-order is assembled taking the syntactic form of the infobase into account. The formulas retained in this theory contraction will be kept in the base contraction. The others will be weakened.

In those papers some examples can be found which show that in the knowledge level the results are pretty decent, but in the symbolic level the resulting base is made up of unnatural and redundant formulas. The authors make it clear that their purpose is to use the expressiveness of the infobase to improve belief set contraction.

They define base contraction from belief set contraction. Therefore, due to the closure of belief sets, we can conclude that their main concern is not of a computational nature. So, our works are similar in the sense they address pseudo-contractions, but have very different focuses, since our approach seems closer to the possibility of an actual implementation.

7.6   Horn Belief Revision

Works such as [Del08],[DW13] and [BMVW11] have addressed belief revision for Horn logic, which is a fragment of classical propositional logic.

In the AGM paradigm everything is done inside a classical logic framework (which means that the remainder sets are calculated according to a classical logic and success is verified according to it, also). Considering the general results of [HW02], everything can be done inside some particular logic (which can be very different from the classical ones). In these papers, everything is done inside Horn logic. In this case, the difference from these works to ours is that here we deal with two different consequence relations: one to compute the remainder sets and another to compute success.

Chapter 8

Conclusions and Future Work

Here we will make a small summary of what was done in this project and then list some of the possible issues future work can cover.

8.1 Conclusion

As already said, traditional belief revision is difficult to implement due to the infinite set of consequences produced by the deductive closure of any set of beliefs. Working with belief bases, sets not closed by logical consequence, is possible, however partial meet contraction applied on these sets has a destructive effect, usually eliminating much more information than necessary.

The idea of employing a consequence relation which produces less consequences than the classical one on belief bases to produce computable and yet reasonable belief base contractions was posed by Ribeiro and Wassermann [RW08]. The characterization of the operation proposed there (here abbreviated to GPMPC) is difficult, since the operation is very complicated. In order to further the understanding of GPMPC and hence make a small contribution to the study of the problem of implementation of belief revision, we studied the properties of a simplified operation, inspired by GPMPC. This new operation is simply partial meet contraction for bases closed under a weak consequence relation.

Some formal properties of the operation have been found and demonstrated, and the connections between it and GPMPC revealed, and also some further properties of GPMPC have been found. The possible applications of this simple operation have been explored, and some examples have been given, the main ones being related to ontology engineering and belief dynamics in probabilistic logic.

Moreover, we have put forward a discussion about the principles underlying belief revision, centered on pseudo-contractions. The postulate of core-addition, proposed in the paper mentioned [RW08], has been deeply analysed and some modifications of it have been proposed.

With this work we hope to have done a meaningful contribution to the study of implementation of belief revision and bounded reasoners.

8.2 Future Work

As the reader may have noticed, much of what was done here was inspired by some ideas from the paper from Ribeiro and Wassermann [RW08]. If I had more time, I would naturally continue investigating the other constructions proposed there, trying to find their representation theorems and connections between them and other pseudo-contractions, and to improve the intuitions and discussions about them and their parts.

Of course, revision operations have been deliberately ignored in this dissertation. Even though they can be defined via contraction and Levi identity, studying the operation proposed here by the standpoint of revisions can perhaps show new interesting facts. This also applies for consolidations and other less famous belief revision operations.

At the moment, a study of the computational complexity of various classes of Cn^* functions is lacking. As pointed out by professor Marcelo Finger, it would be interesting to see how much of the classical consequences we could obtain in polynomial time.

We still need a lot more comprehension on how to build useful, meaningful Cn^* functions. It is not too evident what should be kept and what could be removed from classical consequence in order to produce a significant subset of it.

We briefly mentioned a two-place pseudo-contraction, whose weak consequence relation also considers the input sentence as a parameter. This generalization of the pseudo-contraction studied here still must be examined.

One last point to be explored and which is not clear yet is the understanding of the implications of this theoretical background to the practical examples displayed here. We have showed how our operator relates, for example, to an inference operation that is being thoroughly used in ontology engineering, but the implications of our theory to that area are not completely clear yet.

Bibliography

[AGM85]   Carlos Alchourrón, Peter Gärdenfors and David Makinson. On the logic of theory change. *Journal of Symbolic Logic*, 50(2):510–530, 1985. 1, 9, 10, 11, 12

[AM81]   Carlos Alchourrón and David Makinson. *New studies in Deontic Logic*, chapter Hierarquies of regulation and their logic, pages 125–148. Reidel Publishing Company, 1981. 18

[AM82]   Carlos Alchourrón and David Makinson. On the logic of theory change: contraction functions and their associated revision functions. *Theoria*, 48:14–37, 1982. 11

[AM85]   Carlos E Alchourrón and David Makinson. On the logic of theory change: Safe contraction. *Studia logica*, 44(4):405–422, 1985. 13

[BHS08]   Franz Baader, Ian Horrocks and Ulrike Sattler. Description logics. *Handbook of Knowledge Representation*, 3:135–179, 2008. 33, 34, 35

[BL04]   Ronald J. Brachman and Hector J. Levesque. *Knowledge Representation and Reasoning.* Elsevier, 2004. 3

[BMVW11] Richard Booth, Thomas Andreas Meyer, Ivan José Varzinczak and Renata Wassermann. On the link between partial meet, kernel, and infra contraction and its application to horn logic. *Journal of Artificial Intelligence Research (JAIR)*, 42:31–53, 2011. 38

[CPW01]   Samir Chopra, Rohit Parikh and Renata Wassermann. Approximate belief revision. *Logic Journal of the IGPL*, 9(6):755–768, 2001. 38

[dBFR⁺16] Glauber de Bona, Marcelo Finger, Márcio M. Ribeiro, Yuri D. Santos and Renata Wassermann. Consolidating probabilistic knowledge bases via belief contraction. In *Proceedings of the 15th International Conference of Principles of Knowledge Representation and Reasoning*, 2016. 35

[Del08]   James P. Delgrande. Horn clause belief change: Contraction functions. In Gerhard Brewka and Jérôme Lang, editors, *Proceedings of the 11th International Conference of Principles of Knowledge Representation and Reasoning*, pages 156–165. AAAI Press, 2008. 38

[DW13]   James P. Delgrande and Renata Wassermann. Horn clause contraction functions. *Journal of Artificial Intelligence Research (JAIR)*, 48:475–511, 2013. 38

[Eps90]   Richard L. Epstein. *The semantic foundations of logic.* Springer, 1990. 3

[Fuh91]   André Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20:175–203, 1991. 14

[Gär88]   Peter Gärdenfors. *Knowledge in Flux - Modeling the Dynamics of Epistemic States.* MIT Press, 1988. 1, 10

[Gru93] Thomas R. Gruber. A translation approach to portable ontology specifications. *Knowledge acquisition*, 5(2):199–220, 1993. 33

[Gru95] Thomas R. Gruber. Toward principles for the design of ontologies used for knowledge sharing. *International journal of human-computer studies*, 43(5):907–928, 1995. 33

[Han89] Sven Ove Hansson. New operators for theory change. *Theoria*, 55(2):114–132, 1989. 14

[Han91] Sven Ove Hansson. Belief contraction without recovery. *Studia Logica*, 50(2):251–260, 1991. 13, 14

[Han92a] Sven Ove Hansson. A dyadic representation of belief. In Peter Gärdenfors, editor, *Belief Revision*, volume 29 of *Cambridge Tracts in Theoretical Computer Science*, pages 89–121. Cambridge University Press, 1992. 1

[Han92b] Sven Ove Hansson. Reversing the Levi identity. *Journal of Philosophical Logic*, 22:637–639, 1992. 13

[Han93] Sven Ove Hansson. Changes of disjunctively closed bases. *Journal of Logic, Language and Information*, 2(4):255–284, 1993. 14, 22, 37

[Han94] Sven Ove Hansson. Kernel contraction. *Journal of Symbolic Logic*, 1994. 13, 14

[Han99] Sven Ove Hansson. *A Textbook of Belief Dynamics.* Kluwer Academic, 1999. 13, 14, 18, 19, 20, 21, 23, 29, 32

[Har86] Gilbert Harman. *Change In View: Principles of Reasoning.* Bradford Books. MIT Press, 1986. 1

[HW02] Sven Ove Hansson and Renata Wassermann. Local change. *Studia Logica*, 70(1):49–76, 2002. 37, 38

[LMP08] Vladimir Lifschitz, Leora Morgenstern and David Plaisted. Knowledge representation and classical logic. *Handbook of Knowledge Representation*, 1:3–88, 2008. 33

[Mak87] David Makinson. On the status of the postulate of recovery in the logic of theory change. *Journal of Philosophical Logic*, 16:383–394, 1987. 10

[Mey01] Thomas Andreas Meyer. Basic infobase change. *Studia Logica*, 2001. 38

[MLH00] Thomas Andreas Meyer, Willem Adrian Labuschagne and Johannes Heidema. Infobase change: A first approximation. *Journal of Logic, Language and Information*, 9:353–377, 2000. 38

[Neb89] Bernhard Nebel. A knowledge level analysis of belief revision. In *Proceedings of the 1st International Conference of Principles of Knowledge Representation and Reasoning*, 1989. 14

[New82] Allen Newell. The knowledge level. *Artificial intelligence*, 18(1):87–127, 1982. 2

[NM01] Natalya F. Noy and Deborah L. McGuinness. Ontology development 101: A guide to creating your first ontology. http://www-ksl.stanford.edu/people/dlm/papers/ontology-tutorial-noy-mcguinness-abstract.html, 2001. 33

[PAG06] Jorge Pérez, Marcelo Arenas and Claudio Gutierrez. Semantics and complexity of SPARQL. In *International semantic web conference*, volume 4273, pages 30–43. Springer, 2006. 33

[PS06] Eric Prud and Andy Seaborne. SPARQL query language for RDF. http://www.w3.org/TR/rdf-sparql-query/, 2006. 33

[Rot00]   Hans Rott. Two dogmas of belief revision. *The Journal of Philosophy*, 97(9):503–522, 2000. 1

[RW08]   Márcio Moretto Ribeiro and Renata Wassermann. Degrees of recovery and inclusion in belief base dynamics. *12th International Workshop on Non-Monotonic Reasoning (NMR 2008)*, 2008. 15, 16, 21, 23, 24, 28, 30, 37, 39

[RWFA13]   Márcio Moretto Ribeiro, Renata Wassermann, Giorgos Flouris and Grigoris Antoniou. Minimal change: Relevance and recovery revisited. *Artificial Intelligence*, 201:59–80, 2013. 13

[SC95]   Marco Schaerf and Marco Cadoli. Tractable reasoning via approximation. *Artificial Intelligence*, 74(2):249–310, 1995. 38

[SMH08]   Rob Shearer, Boris Motik and Ian Horrocks. Hermit: A highly-efficient owl reasoner. In *OWLED*, volume 432, page 91, 2008. 33

[SRW15]   Yuri David Santos, Márcio Moretto Ribeiro and Renata Wassermann. Between belief bases and belief sets: Partial meet contraction. In *Proceedings of the International Workshop on Defeasible and Ampliative Reasoning (DARe-15)*, 2015. 17

[Was00]   Renata Wassermann. *Resource Bounded Belief Revision*. PhD Thesis, Universiteit van Amsterdam, 2000. 32

[Was03]   Renata Wassermann. Generalized change and the meaning of rationality postulates. *Studia Logica*, 73(2):299–319, 2003. 32