Kernel, Cumulative, and Safe Contractions

Alexander Mark Smith
ams1@andrew.cmu.edu

August 2009

Preface

The original idea for this thesis was to find a characterization theorem for the class of safe contractions over theories, a problem which has been open since Alchourrón and Makinson defined them and gave a construction-to-postulates result in 1985 [4]. After being unable to extend the technique Rott used in his result for regular and virtually connected safe contractions over theories [20], I attempted to reverse any of several steps Hansson used to redefine safe contractions in [14]. As explained in this thesis, this is impossible at each step.

I have been unable to solve the original problem and provide an axiomatization, but I do have several interesting partial results. Unfortunately, almost all of these are negative, indicating that the behavior of kernel contractions in general and safe contractions in particular is much more complicated than previously thought. In addition to my own work, I have also collected and organized all the results in the Alchourrón-Gärdenfors-Makinson tradition which relate to characterization.

It is my hope that this thesis will be useful to the next person who studies the only major class of AGM operators still lacking a characterization theorem, and that it inspires future work on the other subclasses of kernel contractions which have to this point gone unstudied.

Acknowledgements

Specifically related to my thesis, I must thank my advisor Horacio Arló-Costa for his advice during my research; in particular, for pointing me toward the example of a non-cumulative smooth incision function, and for providing the main idea behind the falsification of relevance. My second reader, Wilfried Sieg, has given me valuable, consistent encouragement, assistance, and advice for years, and was a very careful and thorough reader of this thesis. Without him, I would not have found that I enjoyed logic so much and started my study of philosophy.

Paul Pedersen and I had a few very enlightening conversations. His intuition about the postulate of conjunctive overlap resulted in my exploration of that idea, and his worry about revision equivalence caused me to put in the section on that topic. Conor Mayo-Wilson asked why only finite cycles were forbidden in the definition of a hierarchy, which is a useful point that needed clarification.

I would also like to thank some non-academics. Many friends, particularly Caitlin Osbahr, gave me much-needed support. Lastly, I would be completely remiss if I failed to mention the immense help I was lucky to receive from my parents, Mark and Amy Smith. Without these people, I would not have finished.

Contents

Chapter 1

Motivation

How does a believer change his beliefs? How can a theory (in the logical/mathematical sense) become a different theory? What guidance is there for how to update a database? These questions are the subject of the field variously called belief dynamics, theory change, or database updating, but most commonly referred to as belief revision.

The modern research program in belief revision was started following the landmark 1985 paper by Alchourrón, Gärdenfors, and Makinson [1] in which they introduced a type of belief change called partial meet contractions and gave a list of properties called the Gärdenfors postulates, then showed that the former notion was characterized by the latter. That is, every use of the partial meet contraction method to change beliefs has exactly the properties set out in the list of Gärdenfors postulates, and every change of belief which has those properties can be represented using the partial meet contraction method.

Since then, and in large part because of this cornerstone result, the vast majority of research in the field has focused on partial meet contraction or similar ideas, but there have been two other types of contraction which have received attention: the entrenchment contractions Gärdenfors and Makinson introduced in 1988 [8] and the safe contractions Alchourrón and Makinson introduced in another 1985 paper [4]. A characterization result for entrenchment contractions is also known, but many of the properties of safe contractions remain mysterious.

However, safe contractions themselves are by far the clearest philosophically and the most intuitive of the three types of changes. When using safe contraction, a believer needs only to have a preference structure called a "hierarchy" on his beliefs which indicates "comparative willingness to stop believing." For example, I may prefer dropping the belief that I am hungry over dropping the belief that I have a stomach, and this preference may be used in determining how my beliefs should change when I stop experiencing pain in my torso. The only property this preference structure must have is that it has no (finite) cycles;

that is, there is no (finite) list of beliefs that I hold for which I prefer losing each belief to the next one in the list and also prefer losing the last one to the first.

By contrast, partial meet contractions have a rather questionable philosophical status. Believers are required to make choices using "selection functions" amongst sets of their beliefs, which it seems strange to do—even unconsciously. Entrenchment contractions are not much better; they have the benefit, like safe contractions, of also only requiring preference structures (here called "entrenchment relations") on beliefs, but there are many more properties these preferences must adhere to. For one thing, preferences must follow logical implication. It is clear already that neither of these types of contractions describe how believers actually behave. At the same time, whether or not these are reasonable restrictions on how believers should behave is a subject of much debate.

Moreover, partial meet contractions and entrenchment contractions are difficult to iterate—there is no straightforward way for a believer to make another change to his beliefs after the first change using these methods. This is because the selection function generating a partial meet contraction and the entrenchment relation generating an entrenchment contraction must be related in certain precise ways to a believer's current beliefs, and they simply will fail to do so after those beliefs change. By contrast, the hierarchies behind safe contractions can continue to work for as many changes as desired.

For these reasons, the properties of safe contractions should be explored, and characterization results like the ones for partial meet and entrenchment contractions need to be found. The project of this thesis is to use the "kernel contraction" notion Hansson introduced in [14] to explore the formal properties of safe contractions and some other closely related types of contractions.

In the next chapter the existing literature on characterizing the properties of various types of changes in the Alchourrón-Gärdenfors-Makinson tradition is reviewed. Familiarity with modern logic but no background in belief revision is assumed. This is a technical thesis and its focus is not on the philosophical plausibility or motivation for each particular formal result, but periodic sections throughout that chapter titled "Philosophical notes" will offer pointers to relevant literature on philosophical problems relating to the formal constructions introduced.

After the literature is reviewed and the reader is acquainted with the current state of the field, I present the new results I have obtained. These, rather unfortunately, show that the situation is even more complex than previously thought. The thesis closes with a summary of the current state of affairs in AGM-style belief revision, and suggestions for further work, since many properties of safe contractions remain to be explored.

Chapter 2

Background

As noted, it will be useful to first review the current state of the literature—this will allow the thesis to be self-contained and readable by anyone with a basic understanding of logic. My own results do not make much sense stripped of context.

I shall give a short overview of the basics of the Alchourrón-Gärdenfors-Makinson (AGM) approach to belief revision which originated in [1], define and construct the three major types of contraction operations traditionally considered, and mention all of the important known results relating to their characterization and classification.

The majority of this presentation can be found in Sven Ove Hansson's book [16], which is the standard reference textbook in belief revision. Not all of the following is present there, though, and the results which do appear are scattered throughout the text (for example, the difference made explicit here between Rems and RemSets is left implicit and not brought out well). Even for the results which appear there, I will favor referencing the original papers where applicable and available.

2.1 Preliminaries

For the rest of this thesis, let $\mathcal{L}$ be a formal language with countably many propositional variables $p, q, r, \ldots$ and a complete set of logical connectives $\wedge$, $\vee$, $\rightarrow$, $\leftrightarrow$, $\neg$, and $\perp$. I will use $\alpha, \beta, \gamma, \ldots$ as metasyntactic variables ranging over the well-formed formulae in $\mathcal{L}$, and $A, B, C, \ldots$ will range over sets of these well-formed formulae. By abuse of notation, let $\mathcal{L}$ also be a classical propositional logic over the language $\mathcal{L}$.

2.1.1 Consequence operations and change operations

We are now in a position to define the following very useful notion, originally due to Tarski and treated quite well in Hansson's text.

Definition 2.1.1 ([16], p. 26)**.** A *consequence operation on* $\mathcal{L}$ is a function Cn : $\mathcal{P}(\mathcal{L}) \to \mathcal{P}(\mathcal{L})$ such that for all subsets A, B of $\mathcal{L}$

inclusion $A \subseteq \mathrm{Cn}(A)$,

monotony if $A \subseteq B$, then $\mathrm{Cn}(A) \subseteq \mathrm{Cn}(B)$,

iteration $\mathrm{Cn}(A) = \mathrm{Cn}(\mathrm{Cn}(A))$.

Moreover, we stipulate a few additional requirements. Namely, for all subsets A and sentences α and β:

supraclassicality if α is derivable from A in $\mathcal{L}$, then $\alpha \in \mathrm{Cn}(A)$,

deduction $\beta \in \mathrm{Cn}(A \cup \{\alpha\})$ if and only if $\alpha \to \beta \in \mathrm{Cn}(A)$,

compactness if $\alpha \in \mathrm{Cn}(A)$, then there is a finite $A' \subseteq A$ such that $\alpha \in \mathrm{Cn}(A')$.

We will also write $A \vdash \alpha$ for $\alpha \in \mathrm{Cn}(A)$.

Any operation on sets of well-formed formulae in $\mathcal{L}$ which satisfies the above six principles is a consequence operation. Using this concept we are able to abstract away from the particular features of one method of deduction and simultaneously handle many notions of consequence. Every result in this thesis is relative to a consequence operation, so at this time fix a particular operation, call it Cn.

We can now define how we will formally represent epistemic objects and states.

Definition 2.1.2. Let α be a sentence in $\mathcal{L}$; α is a *belief*. Let A be a set of sentences in $\mathcal{L}$; A is a *belief base* (or just a *set*). If A is closed, i.e. $A = \mathrm{Cn}(A)$, A is a *belief set* or *theory*.

The "belief set" and "belief base" terminology is somewhat unfortunate, because "belief set" implies closure while the more common term "set" covers both cases. Since the terms "set" and "theory" are familiar to logicians of all types, I will favor them.

In the AGM tradition, sentences are treated as the basic epistemic objects, and theories are treated as epistemic states. There are, of course, well-known objections to this characterization. Firstly, why should belief states merely contain beliefs and nothing else? It seems that an agent's dispositions toward his or her beliefs also form part of the belief state. This has intimate connections with iterated belief changes, and will be revisited in the philosophical sections of this chapter.

Even accepting that belief states are simply collections of beliefs themselves, the deductive closure requirement of theories presents its own problems It is clear no one actually believes all that closed belief sets require of us (this is the so-called "problem of logical omniscience," for example, no one believes all the mathematical truths), and therefore this representation of a belief state cannot be used in any descriptive analysis of what believers actually do. Furthermore, all theories have infinite cardinality, and therefore cannot be used directly in any computational implementation of belief change.

At the same time, closure is a natural condition to require: any statement which is derivable in one step from something we believe it intuitively seems that we must believe. The philosopher Isaac Levi has made interesting remarks in which he gets around this problem by interpreting belief sets as epistemic *commitments* rather than epistemic states strictly speaking, for more on this see [17] and [18]. Others working in the field, in particular Hansson, prefer thinking of our epistemic states as belief bases which are in general not closed. This approach allows one to privilege certain beliefs as "foundational" (i.e. actually present in the belief base) and leave the others as merely "derived." For more on this idea see Hansson's writings, in particular [11] and [15]. The debate on the status of logical closure, normatively speaking, remains quite open.

These philosophical issues are outside of the scope of this formal thesis, so I will report and provide results both for the case of arbitrary sets and the restricted case of logically closed theories.

Now that we have representations of belief states, we need to represent changes to those states.

Definition 2.1.3. Let $\cdot$ be a function $\cdot : \mathcal{P}(\mathcal{L}) \times \mathcal{L} \rightarrow \mathcal{P}(\mathcal{L})$ that takes a set and a sentence and returns a new set. $\cdot$ is a *(global) belief change operation*. Similarly, let A be a set and $\cdot : \mathcal{L} \rightarrow \mathcal{P}(\mathcal{L})$ a function, dependent on A, which takes a sentence and returns a new set. $\cdot$ is a *belief change operation for A*.

Since they are only functions of one argument that depend on a belief state, non-global operations of belief change should be written $\cdot_A(\alpha)$, but they are traditionally written $A \cdot \alpha$ and I will follow tradition and write all operations in this thesis in infix notation.

This distinction between binary operations of a global sort and unary operations for particular sets is a deeply important one and one that is too often overlooked in the AGM literature. To quote Hans Rott, even in Alchourrón, Gärdenfors, and Makinson's classic paper [1], they "sometimes use [belief change operations] as binary functions taking various belief sets as their first argument, but this is not in the spirit of what they actually do," ([21], p. 260). The problem is that sometimes it is impossible to describe changes of belief as a function with two inputs that produces reasonable outputs, but it is much easier to describe how a particular belief state can change as a function of one input. Again, this will be revisited later in the philosophical notes sections on iteration.

Traditionally, three kinds of belief change operations are considered: *Expansions* are the addition of a new belief to a belief state, *contractions* are the removal of a belief from a belief state, and *revisions* are a general form of change from one belief state to another.

Expansions are by far the most straightforward. For expansion, in the case of theories we will follow Levi, simply adding by brute force the new belief and closing under consequence.

Definition 2.1.4 ([17])**.** A belief change operation + is the *expansion* operation (for theories) if and only if for all theories A and sentences α we have that $A + \alpha = \mathrm{Cn}(A \cup \{\alpha\})$.

Notice that this actually is a two-place function, and provides a way to characterize adding a belief to any belief state at all.

Of course, this only makes sense if we are restricting our epistemic states to theories, as Levi does. In the more general case of (possibly non-closed) sets, we simply add the belief and do not close.

Definition 2.1.5. A belief change operation + is the *expansion* operation (for sets in general) if and only if for all sets A and sentences α we have that $A+\alpha = A\cup\{\alpha\}$.

The same symbol is used for both; when we are concerned only with theories we will use the closing definition, in the general case we will use the non-closing definition. This will be obvious from context.

Let us assume for a moment we understand the notion of contraction. The following says that this is all we need to understand in the current framework:

Definition 2.1.6 (Levi identity)**.** Let $\div$ a contraction operation be given. We may define the associated revision operation $* : \mathcal{P}(\mathcal{L}) \times \mathcal{L} \to \mathcal{P}(\mathcal{L})$ in the following way: for all sets A and sentences α, $A * \alpha = (A \div \neg\alpha) + \alpha$.

The Levi identity says that whenever we need to revise by a new belief α, we may first contract by its negation, getting rid of any possibly contradictory beliefs, and then expand by α. This is also controversial—many ask, "When do we ever purely lose a belief, without gaining a new one?" but this discussion is also outside of the scope of a technical thesis. The lack of examples of "pure" contractions, losses of belief without any new beliefs arising, is of central importance to the "revision equivalence" objection discussed in 2.4.2. For an introduction to the Levi identity in the AGM framework, see [3].

Using the Levi identity and the simple understanding of expansion, we can turn our attention entirely to contractions for the remainder of this thesis.

2.1.2 The Gärdenfors postulates

Since there are many different kinds of contractions, rather than constructing them now, it is good to first discuss their properties as a way of comparing various types. Thus we are led to a discussion of the *Gärdenfors postulates* for contractions. These originate in Peter Gärdenfors's [7].

Contractions are supposed to *remove* beliefs. This means two things: a contraction should not add any beliefs not already present in the belief state, and a contraction should actually remove the specified belief.

In all the following, let $\div$ be a belief change operation, A a set, and α and β sentences.

success If $\alpha \notin \mathrm{Cn}(\emptyset)$, then $\alpha \notin \mathrm{Cn}(A \div \alpha)$.

inclusion $A \div \alpha \subseteq A$.

These two postulates capture exactly the above two ideas. (Notice that we cannot ever really remove any of the tautologies, since they will be elements of $\mathrm{Cn}(A)$ for any A by the definition of Cn, so we must put the antecedent on the postulate of success for any change operation to satisfy it.) Thus, we make the following definition:

Definition 2.1.7. Let $\div$ be a belief change operation. $\div$ is a *contraction* if and only if it satisfies success and inclusion.

There are a few other properties it seems natural for contraction functions to have simply from a formal standpoint. Namely, in the case of theories the results of contractions should be closed, and equivalent sentences should be contracted equivalently.

closure If $A = \mathrm{Cn}(A)$, then $A \div \alpha = \mathrm{Cn}(A \div \alpha)$.

extensionality If $\alpha \leftrightarrow \beta \in \mathrm{Cn}(\emptyset)$, then $A \div \alpha = A \div \beta$.

Notice at this point almost no restrictions limiting the behavior of the contraction have been given; we can define an operation such that for all α, $A \div \alpha = \mathrm{Cn}(\emptyset)$ and it would satisfy the above postulates. Preventing such extreme behavior by requiring that a contraction actually preserve beliefs which are in some way "unrelated" is the role of vacuity and recovery.

vacuity If $\alpha \notin \mathrm{Cn}(A)$, then $A \div \alpha = A$.

recovery $A \subseteq (A \div \alpha) + \alpha$.

Vacuity requires that if one does not believe the proposition being contracted from, one should not change one's beliefs, which seems intuitive and forbids the extreme operation given above. However, as noted in Section 2.4.2, the five postulates given above (excluding recovery) are still too permissive. No restrictions other than closure have been placed on the action of the contraction for sentences in the theory.

Recovery requires that if one first contracts by a belief and then expands by it, all of one's original beliefs are again present. On first reading it may be plausible, but there have been many philosophical objections raised against it (see section 2.3 of [16] and [9] for an introductory discussion). As with all controversial parts of belief revision, the debate is centered on the status of the normative claim that believers should obey recovery. The related descriptive claim is easily falsified, and it must be noted that the postulates of belief revision were not defined from a descriptive standpoint.

However, even in a purely formal context, recovery is harder to remove than one might think. First, in the case of theories the much-weaker postulate of core retainment implies recovery, as we shall see in Section 2.5.1. Further, as reported in Section 2.4.2, Theorem 2.4.12 means that even when no other postulates are stipulated, recovery is essentially present for all contractions of theories. This is part of the motivation for Hansson's work on defining contractions over arbitrary sets.

These six postulates together are the *basic Gärdenfors postulates*.

Definition 2.1.8. Let $\div$ be a belief change operation. $\div$ satisfies the basic Gärdenfors postulates if and only if it satisfies success, inclusion, closure, extensionality, vacuity, and recovery.

After a moment's reflection, one realizes that the only "choices" one may make when contracting come when contracting conjunctions. Closure and success together require that when contracting disjunctive beliefs one contracts both disjuncts as well, and conditional beliefs are of course essentially disjunctive. However, when contracting a conjunction, one may contract either or both of its conjuncts. In [1], the following two postulates for conjunctions were introduced.

conjunctive overlap $(A \div \alpha) \cap (A \div \beta) \subseteq A \div (\alpha \wedge \beta)$.

conjunctive inclusion If $\alpha \notin \mathrm{Cn}(A \div (\alpha \wedge \beta))$, then $A \div (\alpha \wedge \beta) \subseteq A \div \alpha$.

Conjunctive overlap requires that anything that is retained when contracting by α and also when contracting by β is retained when contracting by their conjunction. Conjunctive inclusion requires that if α was *not* given up when contracting by $\alpha \wedge \beta$ (and therefore β alone was chosen for removal since one of them must be), then anything retained in the contraction of the conjunction is retained when contracting by α.

These two postulates are also controversial (in a normative context; as descriptive claims they are even more easily falsified than recovery is, being more complex), but they possess great formal interest, and again I will set discussions of their plausibility aside.

Definition 2.1.9. Let $\div$ be a belief change operation. $\div$ satisfies the *supplementary Gärdenfors postulates* if and only if it satisfies conjunctive overlap and conjunctive inclusion.

Now, with the notion of contraction clarified and some possible properties of contractions expressed, we can move to constructing particular classes of contraction operations.

2.2 Partial meet contractions

Partial meet contractions are by far the most well-known class of belief change operations. They are the subject of the classic AGM paper [1], which is where the following definitions and theorems originate. There are two central ideas: *remainder sets*, and *selection functions*.

When contracting α from A, *success* requires the result of the contraction to fail to imply α. Given *inclusion*, the result must also not include anything new. There are two immediate ideas about how these goals can be accomplished: pick a subset of A that does not imply α to be your result (and try to retain as much of A as possible), or remove an element out of each subset which *does* imply α, and return the result. The latter idea will be discussed in Section 2.5 on kernel contractions, the first idea is the basis for partial meet contractions.

Definition 2.2.1 ([2])**.** Given a set A and a sentence α, $A \perp \alpha$ is the set containing all inclusion-maximal subsets of A which do not imply α. It is called the *remainder set of A by α*, and its elements are the *α-remainders of A*.

The set of remainder sets of A will be written RemSets(A). (Notice $\mathrm{RemSets}(A) \subseteq \mathcal{P}(\mathcal{P}(A))$.)

Any of the elements of $A \perp \alpha$ are *a priori* reasonable candidates for the result of contracting A by α, and in fact in early work on belief revision *choice contractions* were studied which simply picked elements of $A \perp \alpha$ as their results. Another idea is to let $A \div \alpha$ be $\bigcap(A \perp \alpha)$, the elements which are preserved under every choice contraction. This is *(full) meet contraction*. Partial meet contraction is a common generalization of these ideas. The unifying structure is that of a selection function.

Definition 2.2.2 ([1])**.** Given a set A, a *selection function γ for A* is a function $\gamma : \mathrm{RemSets}(A) \to \mathcal{P}(\mathcal{P}(A))$ such that for all α

1. If $A \perp \alpha \neq \emptyset$, $\gamma(A \perp \alpha)$ is a nonempty subset of $A \perp \alpha$,

2. if $A \perp \alpha = \emptyset$, then $\gamma(A \perp \alpha) = \{A\}$.

A selection function chooses elements of the remainder set which are preferred to the others. Notice that it returns a set of equally preferred remainders, not simply a single preferred remainder.

Definition 2.2.3 ([1])**.** Let A be a set and γ a selection function for A. The *partial meet contraction on A generated by γ* is the operation $\sim_\gamma$ such that for all α we have $A \sim_\gamma \alpha = \bigcap \gamma(A \perp \alpha)$.

An operation $\div$ on A is a partial meet contraction for A if and only if there is a selection function γ such that $A \div \alpha = A \sim_\gamma \alpha$ for all α.

Notice that the study of partial meet contractions is simply the study of selection functions since there is a clear interdefinability between contractions and selections.

The following characterization result is the central result in the AGM tradition.

Theorem 2.2.4 ([1])**.** *$\div$ is a partial meet contraction for a theory A if and only if it satisfies the basic Gärdenfors postulates.*

To satisfy the supplementary Gärdenfors postulates, additional structure must be placed on the selection functions.

2.2.1 Relational selections and transitivity

This structure is obtained by generating the selection from a relation. It is well known from the theory of rational choice that only some choice functions are "rationalizable," in that they are coherent with an actor's preference relations. Expressing choices in terms of a relation places constraints on the choices, and similar constraints are at play here.

Definition 2.2.5 ([1])**.** A selection function γ for A is *relational* if and only if there is a relation $\sqsubseteq$ on $\mathcal{P}(A)$ which generates it. That is, if and only if there is a relation $\sqsubseteq$ such that if $A \perp \alpha$ is nonempty then γ selects the elements of $A \perp \alpha$ that are maximal under $\sqsubseteq$, i.e. $\gamma(A \perp \alpha) = \{B \in A \perp \alpha | (\forall C \in A \perp \alpha)(C \sqsubseteq B)\}$ for all α.

A partial meet contraction is relational if and only if it is generated by a relational selection function. If a relation has a property, we also say that the selection function and partial meet contraction generated by that relation have that property. For example, a transitive relation generates a transitively relational selection function and a transitively relational partial meet contraction.

Properties of relations which are studied in other branches of logic and mathematics can now be used to define subclasses of relational partial meet contractions. There are two properties we are particularly interested in:

transitivity If $A \sqsubseteq B$ and $B \sqsubseteq C$, then $A \sqsubseteq C$.

maximizing If $A \subset B$, then $A \sqsubset B$.

Both of these are well-known "rationality constraints" from the theory of rational choice, and are equally plausible here. Maximizing turns out to be implied by transitivity in the case of contractions on theories.

Lemma 2.2.6 ([12]). *Let A be a theory and $\sim_\gamma$ a partial meet contraction for A. Then $\sim_\gamma$ is transitively maximizingly relational if and only if $\sim_\gamma$ is transitively relational.*

Using this result, we arrive at the other major result of the original AGM paper.

Theorem 2.2.7 ([1]). *An operation $\div$ on a theory A is a transitively relational partial meet contraction for A if and only if it satisfies the basic and the supplementary Gärdenfors postulates.*

Now we have models of contractions that characterize the two main sets of postulates. This is the foundation of the "two-tiered" approach that is a hallmark of the AGM tradition, and the other types of contractions reviewed here will generally fit into one of these two tiers. The pattern of a general contraction based on arbitrary functions satisfying the basic postulates and a specialized contraction based on relations satisfying the supplementary postulates appears frequently.

However, all of this work has been done only for contractions on theories. As noted in an earlier section of this thesis, there are philosophical and computational objections to this restriction to theories, so it is important to see how partial meet contractions behave on arbitrary sets which may fail to be closed.

2.2.2 Partial meet contractions on arbitrary sets

Hansson studied the behavior of partial meet contractions on arbitrary sets in his 1992 [10], finding a characterization result which occupies the central place in the study of contractions on sets that the AGM results just reported occupy in the study of contractions on theories.

The postulates satisfied are of course slightly different. Uniformity is a strengthened version of extensionality, and relevance has connections to vacuity, closure, and recovery.

uniformity If for all $A' \subseteq A$ it is the case that $\alpha \in \mathrm{Cn}(A')$ if and only if $\beta \in \mathrm{Cn}(A')$, then $A \div \alpha = A \div \beta$.

relevance If $\beta \in A$ and $\beta \notin A \div \alpha$, then there is A' such that $A \div \alpha \subseteq A' \subseteq A$ and that $\alpha \notin \mathrm{Cn}(A')$ but $\alpha \in \mathrm{Cn}(A' \cup \{\beta\})$.

These postulates and the postulates defining contractions produce a characterization result.

Theorem 2.2.8 ([10]). *$\div$ is a partial meet contraction for A if and only if $\div$ satisfies success, inclusion, relevance, and uniformity.*

It is clear that the class of partial meet contractions on arbitrary sets is strictly larger than those on theories, simply because of that restriction. Therefore, these postulates imply but are not implied by the basic Gärdenfors postulates. What is less clear is how exactly this implication comes about. Here are some helpful lemmata.

Lemma 2.2.9 ([16]). *Let $\div$ be an operation on an arbitrary set A.*

1. *If $\div$ satisfies uniformity, then $\div$ satisfies extensionality.*
2. *If $\div$ satisfies relevance and inclusion, then $\div$ satisfies vacuity and closure.*
3. *If A is a theory and $\div$ satisfies relevance, then $\div$ satisfies recovery.*

Corollary. *Let $\div$ be an operation on a set A. If $\div$ satisfies success, inclusion, uniformity, and relevance, then $\div$ satisfies the basic Gärdenfors postulates except recovery. If in addition A is a theory, $\div$ also satisfies recovery.*

Note that the assumption in the corollary that A is closed is only used to obtain recovery. For the other direction, that assumption is needed in every case.

Lemma 2.2.10 ([16]). *Let $\div$ be an operation on a theory A.*

1. *If $\div$ satisfies extensionality and vacuity, then $\div$ satisfies uniformity.*
2. *If $\div$ satisfies closure, inclusion, vacuity, and recovery, then $\div$ satisfies relevance.*

Corollary. *Let $\div$ be an operation on a theory A. If $\div$ satisfies the basic Gärdenfors postulates, then $\div$ satisfies success, inclusion, uniformity, and relevance.*

Unfortunately, for arbitrary sets a characterization result relating to the supplementary postulates has not been found, though soundness results have and are reported below. Note that Lemma 2.2.6 is only valid in the case of theories, so the relevant operations here are the transitively *maximizingly* relational partial meet contractions for arbitrary sets.

Theorem 2.2.11 ([13]). *If $\div$ is a transitively maximizingly relational partial meet contraction for a set A, then $\div$ satisfies conjunctive overlap.*

Conjunctive inclusion has only been shown for finite and disjunctively closed sets. (Disjunctive closure is a weakening of logical closure, requiring a set to contain disjunctions of each pair of its non-disjunctive elements.)

Theorem 2.2.12 ([16], p. 152). *If $\div$ is a transitively maximizingly relational partial meet contraction for a finite and disjunctively closed set A, then $\div$ satisfies conjunctive inclusion.*

2.2.3 Philosophical notes

Before we turn our attention to other contraction types, it is worth remarking on two problems with the partial meet approach. The first is purely philosophical: the relations $\sqsubseteq$ which generate partial meet contractions are relations on $\mathcal{P}(A)$. A believing agent actually possessing preferences between subsets of his beliefs is strange even from a normative point of view, let alone a descriptive one.

Secondly, recall that in the definition of a selection function γ for A, we require that $\gamma(\emptyset) = \{A\}$. This implies that γ is *not* a selection function for any set different from A, in particular, by success it will not be a selection function for $A \div \alpha$ for any (non-tautologous) α actually in A. This means that partial meet contractions cannot be iterated, as the approach gives no guidance on how to produce a new selection function, presumably one is not given to the believing agent when he receives a new sentence to contract from, and after a single contraction the selection function the believing agent possesses becomes useless. This problem is shared by the entrenchment contractions discussed next, but the incisions of kernel contractions and the hierarchies of safe contractions discussed afterwards may be reused.

See [11] for a further discussion.

2.3 Entrenchment contractions

Entrenchment contractions are a radically different method of contraction, introduced in 1988 by Gärdenfors and Makinson [8]. The central idea is to utilize a preference structure the agent possesses on his beliefs directly, rather than on subsets thereof as in relational partial meet contractions. These are *entrenchment relations*, relations on an agent's beliefs with an intended reading of "willingness to give up."

In all the following, let A be a set, α, β, and γ be sentences, and $\leq$ be a relation on A.

transitivity If $\alpha \leq \beta$ and $\beta \leq \gamma$, then $\alpha \leq \gamma$.

dominance If $\alpha \vdash \beta$, then $\alpha \leq \beta$.

conjunctiveness Either $\alpha \leq \alpha \wedge \beta$ or $\beta \leq \alpha \wedge \beta$.

minimality If $\bot \notin \mathrm{Cn}(A)$, then $\alpha \notin A$ if and only if $\alpha \leq \beta$ for all β.

maximality If $\beta \leq \alpha$ for all β, then $\vdash \alpha$.

The meaning of transitivity is obvious, dominance says that entrenchment follows consequence, conjunctiveness says that a conjunction is as entrenched as one of its conjuncts, minimality says that every sentence not currently believed is minimally entrenched (as long as the beliefs are consistent), and maximality says that only logical truths are maximally entrenched.

Definition 2.3.1 ([8])**.** $\leq$ is an *entrenchment relation on A* if and only if it satisfies transitivity, dominance, conjunctiveness, minimality, and maximality.

These are used in a very direct way to produce contractions.

Definition 2.3.2 ([8])**.** $\div$ is the entrenchment contraction on a theory A generated by an entrenchment $\leq$ if and only if $A \div \alpha = \{\beta \in A | \alpha < \alpha \vee \beta \text{ or } \alpha \in \text{Cn}(\emptyset)\}$.

$\div$ is an entrenchment contraction on a theory A if and only if there is some entrenchment $\leq$ on A such that $\div$ is generated by $\leq$.

The only part of this definition which is not completely straightforward is why the disjunction $\alpha \vee \beta$ is present instead of simply β. In short, this guarantees that the contraction satisfies recovery. For a discussion, see page 100 of [16] (where these contractions are called *Gärdenfors entrenchment contractions*).

With the strong constraints placed on the relations, it is not surprising we get the following strong characterization result.

Theorem 2.3.3 ([8])**.** *$\div$ is an entrenchment contraction on a theory A if and only if $\div$ satisfies the basic and the supplementary Gärdenfors postulates.*

2.3.1 Basic entrenchments

It is worth noting that Hans Rott developed a weaker theory of entrenchment-based operations much later, which he calls "basic entrenchment operations" [21]. This work is done in terms of revision operations directly, without using the Levi identity and not considering pure contractions, so reporting those results would require too much background to do here, especially since they are not used by any current results on safe contractions.

However, it is worth noting that in that paper, Rott has found a class of entrenchment-based operations equivalent in the case of theories to the basic Gärdenfors postulates, so study of entrenchment can finally be broken into two "tiers" as discussed in the section on partial meet contractions. Perhaps safe contractions can be related to this new class.

2.3.2 Philosophical notes

Entrenchment contractions are at least somewhat more plausible than partial meet contractions, in that the relations on which they are based are relations on an agent's beliefs, and not on subsets thereof. However, there are lots of restrictions on the behavior of the relation. Of course, all of the restrictions have a clear logical motivation, but their philosophical motivation and normative status is less clear (it is obvious they are not descriptive claims, no believer can even pretend to obey them all). The hierarchies discussed in the next section on safe contractions have much weaker constraints.

Also, like selection functions for A being selection functions only for A (and no other sets) because of the case in their definition for an empty remainder, entrenchment relations for A are entrenchment relations for A only (and no other sets) because of minimality. This causes the same difficulties with iterated belief change that were discussed in Section 2.2.3.

Entrenchment contractions are also the least extensible of the methods discussed in this thesis. There has been little to no work on definitions of entrenchment contractions for arbitrary sets; it is extremely difficult to do so. If the background set for an entrenchment contraction is small, the restrictions on the entrenchment relation behave strangely. And, as discussed, only comparatively recently was a way to weaken entrenchment contractions to satisfy only the basic Gärdenfors postulates found.

The chief interest of entrenchment contractions in this thesis is the result of Hans Rott's, reported as Theorem 2.4.9 in the next section, that a particular kind of relation on which a subclass of safe contractions is based can be turned into an entrenchment relation and vice versa.

2.4 Safe contractions

Safe contractions were introduced in 1985 by Alchourrón and Makinson, making them the second-oldest class of contractions in the AGM tradition—yet they are by far the least understood. The idea is, like in entrenchment contractions, to use a preference structure on an agent's beliefs which roughly corresponds to "willingness to give up," or "relative safety" of pairs of beliefs. However, the relations are used in a more complex manner and have much weaker requirements.

Definition 2.4.1 ([4])**.** A relation $\prec$ on a set of sentences A which is finitely acyclic (i.e. there are not $\alpha_1, \ldots, \alpha_n \in A$ such that $\alpha_1 \prec \cdots \prec \alpha_n \prec \alpha_1$ for any finite n) is called a *hierarchy on A*.

Finite acyclicity is the only requirement placed on the relations, and it is extremely weak. Note transitivity is consistent with finite acyclicity but not required by it. $\alpha \prec \beta$ should be read as "α is less safe than β," i.e. the agent is more willing to give up α than β.

Definition 2.4.2 ([4])**.** The set A/α of elements of A *safe with respect to α* (modulo a hierarchy on A called $\prec$) is the set of all elements of A which are not $\prec$-minimal in any inclusion-minimal subset B of A such that $\alpha \in \mathrm{Cn}(B)$.

The safe elements of A with respect to α are all elements of A which cannot be "blamed" for α; either they do not contribute to the implication of α at all or, when they do, there is always another element contributing to the implication of α which the agent disprefers. So, when contracting from α, one need not remove any of the safe elements.

This usage of the hierarchy $\prec$ explains the requirement of acylicity; because of it, at least one element in each inclusion-minimal subset above will be $\prec$-minimal, and therefore fail to be safe. Moreover, we only need finite acyclicity because these inclusion-minimal subsets will always be finite by compactness of the consequence operation. This will suffice to allow us to define a contraction.

Definition 2.4.3 ([4])**.** $\div$ is the *safe contraction on A generated by the hierarchy* $\prec$ if and only if for all α, $A \div \alpha = A \cap \mathrm{Cn}(A/\alpha)$.

$\div$ is a safe contraction on A if and only if there is some hierarchy $\prec$ on A such that $\div$ is generated by $\prec$.

Safe contractions are intuitive in this way. When contracting from α, we simply keep all the safe elements of A with respect by α, and anything they imply. Alchourrón and Makinson found immediately that safe contractions on theories satisfied the basic postulates.

Theorem 2.4.4 ([4])**.** *If $\div$ is a safe contraction for a theory A, $\div$ satisfies the basic Gärdenfors postulates.*

2.4.1 Regularity and virtual connectivity

Like transitively relational partial meet contractions, additional structure must be specified for safe contractions to satisfy the supplementary postulates. This comes in placing restrictions on the generating hierarchy. Two related properties are given below.

continuing-up If $\alpha \prec \beta$ and $\beta \vdash \gamma$, then $\alpha \prec \gamma$.

continuing-down If $\alpha \vdash \beta$ and $\beta \prec \gamma$, then $\alpha \prec \gamma$.

Both continuing-down and continuing-up say that the hierarchy coheres with consequence in natural but distinct ways. It turns out that these are all that is needed to achieve conjunctive overlap.

Definition 2.4.5. If $\div$ is a safe contraction for A generated by a hierarchy which satisfies a property, we also say the safe contraction has that property. For example, if $\div$ is generated by a continuing-up hierarchy, $\div$ is a continuing-up safe contraction.

Theorem 2.4.6 ([4])**.** *If $\div$ is a continuing-up or continuing-down safe contraction for a theory A, then $\div$ satisfies conjunctive overlap.*

See Section 3.5 for my own results on safe contractions and conjunctive overlap more generally.

Conjunctive inclusion, the other supplementary postulate, is a little trickier. To achieve it, we need virtual connectivity, a strong postulate that says the agent has as many preferences as possible. (Full connectivity, having every unordered pair of beliefs in a preference relationship, is too strong, as it implies reflexivity, which violates finite acyclicity and is thus inconsistent with being a hierarchy.)

virtual connectivity If $\alpha \prec \beta$ then either $\alpha \prec \gamma$ or $\gamma \prec \beta$.

regularity Both continuing-up and continuing-down are satisfied.

The continuing-up and continuing-down distinction collapses in the presence of virtual connectivity and closure.

Lemma 2.4.7 ([4]). *If $\prec$ is a virtually connected hierarchy over a theory A, $\prec$ satisfies continuing-up if and only if it satisfies continuing-down.*

Alchourrón and Makinson were able to produce the following satisfaction result.

Theorem 2.4.8 ([4]). *If $\div$ is a regular and virtually-connected safe contraction for a theory A, then $\div$ satisfies the basic and the supplementary Gärdenfors postulates.*

In a paper the following year [5], Alchourrón and Makinson got part of the characterization result (for the case of a finite language), but the full theorem was only given by Hans Rott later in 1992.

Theorem 2.4.9 ([20]). *$\div$ is a regular and virtually connected safe contraction for a theory A if and only if it satisfies the basic and the supplementary Gärdenfors postulates.*

The technique Rott used was to demonstrate a conversion technique between regular and virtually connected hierarchies on theories and the entrenchments discussed in the last section. Continuing-up and -down and virtual connectivity can be exploited to achieve transitivity, dominance, conjunctiveness, minimality, and maximality.

This is an ingenious technique, but unfortunately it has proven to be inextensible. Without those strong properties, there is no easy connection between hierarchies and entrenchments; perhaps some work could be done linking hierarchies to the weaker basic entrenchments discussed in Section 2.3.1.

As of 2009, this result is still the most advanced technical result and the only characterization result related to safe contractions in existence.

2.4.2 Revision equivalence

There is a concept called *revision equivalence* which is of interest to philosophers and logicians working with contractions, and it is worth noting how it relates to safe contractions in particular. In Section 2.1.1 we discussed the three types of belief change, and mentioned that assuming the (philosophically controversial) Levi identity (Definition 2.1.6), we could focus on contractions exclusively as the arbitrary belief changes of revisions could be modelled in a two-step procedure. However, the revisions obtained by the use of the Levi identity are still of great importance, since examples of pure contractions where no beliefs are added are somewhat difficult to produce, and their existence is also considered philosophically debatable.

The Levi identity is a conversion procedure transforming contraction operators into revision operators. It turns out to be non-injective; that is, some contraction operators are identified after passing through the Levi identity. These are called revision-equivalent.

Definition 2.4.10 ([19])**.** Two contraction operators $\div$ and $\div'$ for a theory A are *revision-equivalent* if and only if for all α we have that $(A \div \neg\alpha) + \alpha = (A \div' \neg\alpha) + \alpha$, i.e., if using the Levi identity to transform them into revision operators transforms them into the same revision operator.

In the paper where he introduced revision equivalence, Makinson defined the following very general class of operations.

Definition 2.4.11 ([19])**.** An operator $\div$ on a theory A is a *withdrawal* if and only if it satisfies success, inclusion, closure, extensionality, and vacuity.

In Section 2.1.2, a completely implausible contraction operator was given which satisfied success, inclusion, closure, and extensionality, but which failed to satisfy vacuity. Satisfying vacuity as well is not much more difficult: define $\div$ a withdrawal on a theory A such that if $\alpha \notin A$ then $A \div \alpha = A$, and if $\alpha \in A$ then $A \div \alpha = \mathrm{Cn}(\emptyset)$ ([16], p. 73). That is, remove all non-tautological beliefs any time you must remove any belief. This shows just how permissive the withdrawal concept is (and thus how much some postulate like recovery is needed).

Withdrawals are disconcerting because of the following result.

Theorem 2.4.12 ([19])**.** *Let A be a theory and $\div$ be a withdrawal for A. Then there is a partial meet contraction $\div'$ for A which is revision-equivalent to $\div$.*

Corollary. *Let A be a theory and $\div$ a contraction for A that satisfies the basic Gärdenfors postulates except recovery. Then there is a contraction $\div'$ for A that satisfies the basic Gärdenfors postulates including recovery and is the same as $\div$ from the standpoint of revisions.*

This result, on first reading, seems to say that much of the technical differences we have discussed are irrelevant from the standpoint of revisions. However, it only says that the classes of withdrawals over theories and partial meet contractions over theories are revision-equivalent; for restricted subclasses of partial meet contractions this result may not apply. As safe contractions over theories do form a subclass of partial meet contractions over theories, a further result showing the applicability or lack of applicability of the revision-equivalence worry is needed.

2.4.3 Philosophical notes

Safe contractions are much simpler than the other types of contraction discussed in the AGM tradition. Hierarchies are relations on beliefs which do not require an agent to adhere to nearly any coherence requirements. The motivation and definition of "safe elements" is intuitive and reasonable. They are not without philosophical problems (for example, when taken over theories they still satisfy recovery), but any they do have are shared by all other approaches to belief revision discussed in this thesis. (There are more modern approaches, such as "mild" and "Levi-" contractions, which do not fit so neatly into this picture.)

Also, note that safe contractions are easily iterable, as a relation can be a hierarchy for many sets. Thus, an agent can contract and contract again. Unfortunately, using the same hierarchy all the time means that an agent's future beliefs are determined by his current belief state, which is not ideal. However, for non-philosophical applications of belief revision techniques, such as database updating in computer science, this may not be a problem. In any case, safe contractions are better off than most. In the next section we will see that (with the falsification of relevance) safe contractions over arbitrary sets are also immune to still other philosophical issues.

2.5 Kernel contractions

Kernel contractions were introduced by Hansson in 1994 in [14], though some ideas go back to [6]. He intended them to be a non-relational superclass of the safe contractions introduced in the previous section; to stand in the same relationship to safe contractions as partial meet contractions stand with regard to relational partial meet contractions in standard AGM theory.

During the presentation of partial meet contractions, we reflected on the following fact: the central idea of the partial meet contraction is the remainder set, the set containing all inclusion-maximal subsets which fail to imply a sentence. Informally, these are all the "ways" we could *fail* to imply something. From the remainder set we choose the result of the contraction. But we also mentioned the dual idea, focusing on all the "ways" something *is* implied.

Definition 2.5.1 ([14]). Let A be a set and α a sentence. The *kernel set* (or *entailment set*) of A with α, written $A \perp\!\!\!\perp \alpha$, is the set of all inclusion-minimal subsets of A entailing α. Its elements are the *α-kernels of A*.

The set of all kernel sets of A will be written KerSets(A), and the set of all kernels of A (for *all* α!) will be written Kers(A). (Notice KerSets(A) $\subseteq \mathcal{P}(\mathcal{P}(A))$, just as RemSets($A$) $\subseteq \mathcal{P}(\mathcal{P}(A))$. Notice also Kers($A$) $\subseteq \mathcal{P}(A)$.)

The second component of partial meet contractions is the selection function which chooses which elements of the remainder set to use. The counterpart in the kernel contraction framework is the incision function, which chooses which elements of the kernels in the kernel set to stop using.

Definition 2.5.2 ([14])**.** An *incision function* σ for a set A is a function σ : KerSets(A) $\to \mathcal{P}(A)$ such that for all α

1. $\sigma(A \perp\!\!\!\perp \alpha) \subseteq \bigcup(A \perp\!\!\!\perp \alpha)$,
2. for all $X \in A \perp\!\!\!\perp \alpha$, if $X \neq \emptyset$, then $X \cap \sigma(A \perp\!\!\!\perp \alpha) \neq \emptyset$.

Notice that there is no clause in the definition of an incision function which explicitly mentions the background set, and thus an incision can be reused like a hierarchy can. This is in direct contrast to entrenchment relations and selection functions, which mention the background explicitly, but agrees with hierarchies, as expected. (For more information, see [16] section 2.11.)

An incision function "cuts into" each kernel, selecting at least one sentence for removal. We use this idea to define the notion of a kernel contraction.

Definition 2.5.3 ([14])**.** Let A be a set and σ an incision function for A. The *kernel contraction $\approx_\sigma$ for A generated by σ* is as follows

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

An operation $\div$ is a *kernel contraction for A* just in case there is some incision function σ for A such that $\div$ is the kernel contraction for A generated by σ.

Notice immediately that a kernel contraction is associated with *exactly one* incision function, and vice versa. Thus, when studying kernel contractions, we can restrict our attention to incision functions (in much the same way that when studying partial meet contractions we may restrict our attention to selection functions).

In [14], Hansson found a characterization result for the class of kernel contractions over arbitrary sets. It turns out we need an even weaker postulate than relevance for the result. This postulate is core retainment.

core retainment If $\beta \in A$ and $\beta \notin A \div \alpha$ then there is $B \subseteq A$ such that $\alpha \notin \mathrm{Cn}(B)$ and $\alpha \in \mathrm{Cn}(B + \beta)$.

Theorem 2.5.4 ([14]). *$\div$ is a kernel contraction for a set A if and only if $\div$ satisfies success, inclusion, uniformity, and core retainment.*

As mentioned, core retainment is a weakening of relevance, so we have the following result as well.

Lemma 2.5.5 ([16]). *If $\div$ is an operation on A which satisfies relevance, $\div$ satisfies core retainment.*

Corollary. *If $\div$ is a partial meet contraction on A, $\div$ is a kernel contraction on A.*

Kernel contractions in general do not satisfy closure, even over theories. The portion of Lemma 2.2.9 relating to closure does not apply here because relevance may not be satisfied. We will need an additional constraint on incision functions to guarantee closure, which is the subject of Section 2.5.1.

2.5.1 Smoothness and closure

The problem is that we may remove sentences without removing other sentences which entail them. We obtained closure "for free" in the partial meet framework, because inclusion-maximal subsets of closed sets will clearly be closed, and the intersection of several closed sets is itself closed. Here, we must introduce the smoothness condition:

Definition 2.5.6 ([14]). Let σ be an incision function for a set A. σ is *smooth* if and only if for all α and β and $B \subseteq A$ that if $\beta \in \mathrm{Cn}(B)$ and $\beta \in \sigma(A \perp\!\!\!\perp \alpha)$ then $B \cup \sigma(A \perp\!\!\!\perp \alpha) \neq \emptyset$.

$\div$ is a *smooth kernel contraction* if and only if it is generated by a smooth incision function.

This is enough to give us another new postulate which Hansson defined in [14], a version of closure which is meaningful for arbitrary sets.

relative closure $A \cap \mathrm{Cn}(A \div \alpha) \subseteq A \div \alpha$.

The following characterization result tells us that relative closure is exactly what the restriction to smooth incisions has obtained.

Theorem 2.5.7 ([14]). *$\div$ is a smooth kernel contraction for a set A if and only if $\div$ satisfies success, inclusion, uniformity, core retainment, and relative closure.*

Moreover, this class is still a superclass of partial meet contractions, even for arbitrary sets.

Lemma 2.5.8 ([16]). *If $\div$ is an operation on A which satisfies relevance, then $\div$ satisfies relative closure.*

Corollary. *If $\div$ is a partial meet contraction on A, then $\div$ is a smooth kernel contraction on A.*

Proof. We start with success, inclusion, uniformity, and relevance. We must obtain core retainment and relative closure. By Lemma 2.5.5, relevance gives us core retainment. By the above lemma, we have relative closure. ◻

Over arbitrary sets, the converse claim fails. (Hansson gives an example on p. 91 of [16].) However, in the case of theories, the converse claim holds. To see why, here is a lemma.

Lemma 2.5.9 ([16])**.** *Let $\div$ be an operation on an arbitrary set A.*

1. *If $\div$ satisfies inclusion and relative closure, $\div$ satisfies closure.*
2. *If $\div$ satisfies inclusion and core retainment, $\div$ satisfies vacuity.*
3. *If A is a theory and $\div$ satisfies core retainment, $\div$ satisfies recovery.*

Corollary. *If $\div$ is a smooth kernel contraction on A, $\div$ satisfies the basic Gärdenfors postulates except recovery. If in addition A is a theory, $\div$ satisfies recovery and is a partial meet contraction.*

Proof. We start with success, inclusion, uniformity, core retainment, and relative closure. We need to obtain closure, extensionality and vacuity. By the above lemma, we have closure from relative closure and inclusion, and we have vacuity from inclusion and core retainment. By Lemma 2.2.10, we have extensionality from uniformity.

If A is a theory, by the above we also have recovery from core retainment and therefore $\div$ is a partial meet contraction. ◻

The above and the previously reported Lemma 2.5.8 that all partial meet contractions are smooth kernel contractions yield the characterization and identification result.

Theorem 2.5.10 ([14])**.** *$\div$ is a smooth kernel contraction for a theory A if and only if $\div$ satisfies the basic Gärdenfors postulates.*

2.5.2 Saturation and smoothness

We can achieve closure (and thus the basic postulates) in the kernel contraction framework in another way. Smoothness is a constraint on the incisions which generate a kernel contraction, which restricts them to the "correct" class to obtain relative closure. We may also modify a kernel contraction directly via a process Hansson called saturation.

Definition 2.5.11 ([14]). Let $\approx_\sigma$ be a kernel contraction for A. The *saturation of* $\approx_\sigma$, written $\hat{\approx}_\sigma$, is defined in the following way. For all α

$$A \hat{\approx}_\sigma \alpha = A \cap \mathrm{Cn}(A \approx_\sigma \alpha)$$

$\div$ is a *saturated kernel contraction for A* if and only if it is the saturation of some kernel contraction for A.

Saturating a contraction satisfies relative closure in the most obvious way—we simply place all the entailed sentences in it directly. These two approaches happily coincide.

Theorem 2.5.12 ([14]). *$\div$ is a smooth kernel contraction for a set A if and only if it is a saturated kernel contraction for A.*

Corollary. *$\div$ is a saturated kernel contraction for a set A if and only if it satisfies success, inclusion, uniformity, core retainment, and relative closure. $\div$ is a saturated kernel contraction for a theory A if and only if it satisfies the basic Gärdenfors postulates.*

2.5.3 Kernel selections and cumulativity

Notice that we do not require that incision functions treat "similar" entailment sets in "similar" ways, which seems like a reasonable idea. For example, if p and q are independent sentences in A, $\{p, p \to q\}$ is an element of $A \perp\!\!\!\perp (p \wedge q)$ and also of $A \perp\!\!\!\perp (p \leftrightarrow q)$, but the definition of an incision function (Definition 2.5.2) does not require that $\sigma(A \perp\!\!\!\perp (p \leftrightarrow q))$ have any relationship whatsoever to $\sigma(A \perp\!\!\!\perp (p \wedge q))$.

This prompts the following idea, mentioned briefly in [14] but developed more in Hansson's book [16].

Definition 2.5.13 ([14]). s is a *kernel selection function* (or *ksf*) for a set A if and only if it is a function $s : \mathrm{Kers}(A) \to \mathcal{P}(A)$ such that for all α and for all $X \in A \perp\!\!\!\perp \alpha$,

1. $s(X) \subseteq X$,
2. if $X \neq \emptyset$, then $s(X) \neq \emptyset$.

The most important thing to notice immediately about kernel selection functions is that they are defined *simultaneously* over all elements of all kernel sets. So if a subset X of A is both an α-kernel and a β-kernel for $\alpha \neq \beta$, X must be treated the same way regardless of whether α or β is being contracted.

We may define an incision function in terms of a ksf in the following way:

Definition 2.5.14 ([16]). Let s be a ksf for a set A. Then an incision function σ is the *cumulation of s* if and only if for all α

$$\sigma(A \perp\!\!\!\perp \alpha) = \bigcup \{s(X) | X \in A \perp\!\!\!\perp \alpha\}$$

An incision function σ is *cumulative* if and only if there is some ksf such that σ is the cumulation of that ksf. An operation $\approx_\sigma$ is a cumulative kernel contraction if and only if it is generated by a cumulative incision function σ.

In this way, we have defined a new class of contractions—the cumulative kernel contractions. Being able to do this is why it is useful to talk about Kers(A) instead of just KerSets(A). We could also have defined a Rems(A), the set of all remainders of A, regardless of whether they're α- or β-remainders, but because of how remainder sets are used to define partial meet contractions, talking about Rems(A) would have been pointless, and we can restrict our attention to RemSets(A). This is because a selection function selects several remainders to use at once, but does not "look inside" the remainders at their elements. By contrast, an incision "looks inside" each kernel of the kernel set it is defined on.

Hansson notes in Section 2.9 of [16] that there are non-cumulative kernel contractions, but his example uses a background set which fails to be closed. I will further investigate the connections between cumulativity and other classes of contractions in Sections 3.3 and 3.4, and extend this example to the case of theories.

2.5.4 Safe kernel contractions

Now we are finally at a point where we can define safe contractions in Hansson's more general kernel contraction framework. This idea was clearly present in [14], but was not mentioned as more than an aside. A comprehensive treatment had to wait until [16].

Safe contractions are essentially relational restrictions of cumulative kernel contractions, so we make the following definition.

Definition 2.5.15 ([16])**.** A ksf s for a set A is *based on a relation $\prec$ on A* if and only if for all X, s selects the $\prec$-minimal elements of X. That is, for all α, for all $X \in A \perp\!\!\!\perp \alpha$, $\beta \in s(X)$ if and only if $\beta \in X$ and there is no $\delta \in X$ such that $\delta \in \beta$.

A ksf is *relational* if and only if it is based on some relation.

Notice that if a relation generates a ksf, it generates a unique ksf. (See my own results in Section 3.4 for when exactly a relation generates a ksf.) Moreover, a hierarchy always suffices to generate a ksf:

Theorem 2.5.16 ([16])**.** *Let $\prec$ be a hierarchy on a set A. Then there is a ksf s for A based on $\prec$. This is the ksf* generated by *$\prec$.*

This is because, as remarked in the section on Alchourrón and Makinson's safe contractions, at least one element will be $\prec$-minimal in every inclusion-minimal subset. In the new ksf terminology, $s(X) \neq \emptyset$ for any $X \neq \emptyset$ by finite acyclicity of $\prec$ and the fact that each kernel X will be finite since Cn is assumed to be compact.

A safe kernel contraction, then, is a cumulative kernel contraction based on a hierarchy which is then saturated.

Definition 2.5.17 ([16])**.** Let $\prec$ be a hierarchy on a set A. Let s be the ksf generated by $\prec$, and let σ be the cumulation of s. Then $\approx_\sigma$, the saturation of the kernel contraction generated by σ, is the *safe contraction based on* $\prec$.

An operator $\div$ is a *safe kernel contraction* for a set A if and only if there is some hierarchy $\prec$ on A such that $\div$ is the safe kernel contraction based on $\prec$.

This definition clearly shares the same motivation as the definition of Alchourrón and Makinson discussed in Section 2.4, but no proof that these definitions were the same appears in the literature. In fact, in the only existing presentation of safe kernel contractions, which is in Section 2.9 of Hansson's textbook [16], there is a further restriction on hierarchies not present in the original defining paper [4], even though Hansson cites that paper in the definition. These issues are the focus of Section 3.1 of the next chapter.

2.5.5 Revision equivalence

Revision equivalence is not an especially pressing matter for kernel contractions, but precisely why that is so should be noted. Of course, kernel contractions exhibit significantly more interesting and varied behavior than partial meet contractions, and so it would be more obnoxious to have that behavior collapse under the Levi identity, but this behavior occurs only when they are taken over arbitrary sets. In that case, the class of all kernel contractions, of smooth kernel contractions, and of partial meet contractions forms a strict descending chain of subsets, as we will see in the next chapter.

However, Makinson's revision equivalence results of [19] only apply to the case of theories. The class of kernel contractions over theories is largely unstudied because the postulate of closure may fail, which actually means it is a *larger* class than even the class of withdrawals, and thus the revision equivalence worry does not apply. Moreover, the distinction between partial meet and smooth kernel contractions collapses, so the worry there is exactly the same as it is for partial meet.

Cumulative kernel contractions over theories form neither a subclass nor a superclass of partial meet contractions over theories, as we shall see in Sections 3.4 and 3.3. It is unclear how they relate to withdrawals, and thus how revision equivalence relates to them. More research is needed in this area.

2.5.6 Philosophical notes

Kernel contractions are somewhat less intuitive than safe contractions are, which is unsurprising since they are so much more general. Much as it is a little strange for a believer to possess a selection function (regardless of whether or not it is relational), it is a little strange for a believer to possess an incision, much less a smooth one. The extra worry about smoothness is alleviated somewhat by the fact that it can be dispensed with via saturation. In contrast to other areas of the AGM tradition, both the normative status and the descriptive status of kernel contractions are relatively open, since they restrict a believer's behavior only very slightly.

As for iterability, kernel contractions share the advantage of safe contractions in that incision functions do not make explicit mention of the background set on which they are defined, and thus may be reused exactly as hierarchies may. They likewise share the drawback that future changes are determined by that incision and that it is not and cannot be modified, and that the history of how an agent arrives at a belief state is not taken into account when constructing a new belief state for him.

Now that all the definitions needed have been introduced and the current state of the art has been presented, let us turn to new results.

Chapter 3

Results

3.1 Safe kernel contractions and safe contractions

As is obvious from the presentation in the last section, Hansson's kernel contractions are supposed to be nonrelational generalizations of safe contractions. It is worth showing, though, that he achieved this goal, and the kernel-based version of safe contraction is the same as the original. No such result appears in the literature, so I derived the below.

Theorem 3.1.1. *Let $\prec$ be a hierarchy on a set A. Then the safe kernel contraction and the safe contraction generated by $\prec$ are identical.*

Proof. Let $\prec$ be a hierarchy on A. $\prec$ generates a ksf s which generates an incision σ which generates a safe kernel contraction, call it $\div_H$ (for Hansson). $\prec$ also generates a safe contraction, call it $\div_{AM}$ (for Alchourrón and Makinson). To show $\div_H$ and $\div_{AM}$ are the same, it is enough to show $A \div_H \alpha = A \div_{AM} \alpha$ for all α. Let α be given. $A \div_{AM} \alpha = A \cap \mathrm{Cn}(A/\alpha)$ and $A \div_H \alpha = A \cap \mathrm{Cn}(A \approx_\sigma \alpha)$. It is enough to show $A \approx_\sigma \alpha = A/\alpha$.

$A \approx_\sigma \alpha = A \setminus \sigma(A \perp\!\!\!\perp \alpha)$, and moreover by the cumulation in the generation of the safe kernel contraction $A \approx_\sigma \alpha = A \setminus \bigcup\{s(X) | X \in A \perp\!\!\!\perp \alpha\}$. An element β is a member of $A \setminus \bigcup\{s(X) | X \in A \perp\!\!\!\perp \alpha\}$ when $\beta \in A$ and $\beta \notin s(X)$ for any $X \in A \perp\!\!\!\perp \alpha$. But recall by the generation of the ksf s in the safe kernel contraction that $\beta \notin s(X)$ means β is not $\prec$-minimal in X. Further, recall that $X \in A \perp\!\!\!\perp \alpha$ means X is a subset of A minimal under inclusion which proves α. So $\beta \in A \approx_\sigma \alpha$ exactly when $\beta \in A$ and β is not $\prec$-minimal in any subset of A minimal under inclusion which proves α. This is exactly the definition of β being safe in A with respect to α, so $\beta \in A/\alpha$ exactly when $\beta \in A \approx_\sigma \alpha$. Thus, $A/\alpha = A \approx_\sigma \alpha$ and $A \div_{AM} \alpha = A \div_H \alpha$. □

Corollary. *The classes of safe contractions and safe kernel contractions are the same. Moreover, any subclass of safe contractions formed by restricting the behavior of the hierarchies is identical to any subclass of safe kernel contractions formed by the same restriction, and vice versa.*

This result, although extremely straightforward once presented, really ought to appear in the literature. It justifies the initial project of this thesis to study safe contractions via kernel contractions, and means Hansson has succeeded in extending safe contractions to a non-relational concept, since the class of safe contractions is obviously a subclass of saturated kernel contractions, and thus of kernel contractions.

As can be seen in the proof above, the safe elements of A with respect to α in Alchourrón and Makinson's terminology are exactly the value on α of the contraction $\approx_\sigma$ which is saturated to become the safe contraction on A. This motivates introducing a definition for this kind of contraction.

Definition 3.1.2. $\approx_\sigma$ is the *presafe contraction* for a set A generated by a hierarchy $\prec$ for A if and only if σ is the cumulation of the ksf s generated by $\prec$.

$\div$ is a presafe contraction for A if and only if there is a hierarchy $\prec$ for A such that $\div$ is the presafe contraction for A generated by $\prec$.

The class of presafe contractions is clearly a subclass of the cumulative kernel contractions, but its precise properties will be explored in the remainder of this thesis.

What makes the above Theorem 3.1.1 less trivial is, if not an error, at least a significant oversight. Hansson remarks already in [14] that one can define safe contractions in the kernel contraction framework, as reported in this thesis and verified above. In the literature, though, this definition is not ever done until Hansson's textbook [16], and there is a difference between the presentation given there and the presentation reported here. On page 94 of [16], Hansson defines a hierarchy, citing Alchourrón and Makinson's [4] as we have. However, he also requires a hierarchy to satisfy the following condition.

intersubstitutivity If $\alpha \leftrightarrow \alpha', \beta \leftrightarrow \beta' \in \mathrm{Cn}(\emptyset)$, then $\alpha \prec \beta$ if and only if $\alpha' \prec \beta'$.

The definition of a hierarchy requiring intersubstitutivity is presented as a citation of the definition in [4] even though it differs from that presented there and reported here, and the condition of intersubstitutivity does not appear elsewhere in the literature (e.g. [5] and [20] also only require finite acyclicity, as I do in this thesis).

Clearly the classes of all plain (merely finitely acyclic) hierarchies and all intersubstitutive hierarchies are not equal, and it is not obvious that this difference should collapse when the hierarchies are used to generate safe contractions. The strongest result I have been able to obtain is the following.

Lemma 3.1.3. *Let $\prec$ be a hierarchy on A. Then there is an intersubstitutive hierarchy $\prec'$ such that if $\div$ is the safe contraction generated by $\prec$ and $\div'$ that generated by $\prec'$, we have $A \div' \alpha \subseteq A \div \alpha$ for all α.*

Proof. Let $\prec$ be a hierarchy on A and $\div$ be the safe contraction generated by $\prec$. We will define an associated intersubstitutive hierarchy on A by removing problematic relationships from $\prec$. Define $\prec'$ such that $\alpha \prec' \beta$ if and only if $\alpha \prec \beta$ and for all $\alpha', \beta' \in A$ such that $\alpha \leftrightarrow \alpha', \beta \leftrightarrow \beta' \in \mathrm{Cn}(\emptyset)$ it is the case that $\alpha' \prec \beta'$. Clearly $\prec'$ is an intersubstitutive hierarchy.

Let a sentence α be given. Let $\beta \in A/'\alpha$, the set of elements of A safe with respect to α modulo $\prec'$. Therefore, in any inclusion-minimal subset B of A such that $\alpha \in \mathrm{Cn}(B)$, if $\beta \in B$ then there is $\gamma \in B$ such that $\gamma \prec' \beta$. Notice every $\prec'$ relationship between members of A is a $\prec$ relationship, by definition. So in any such B, if $\beta \in B$ then there is $\gamma \in B$ such that $\gamma \prec \beta$, and thus $\beta \in A/\alpha$, the set of elements of A safe with respect to α modulo $\prec$. $A/'\alpha \subseteq A/\alpha$ and therefore $A \div' \alpha \subseteq A \div \alpha$, by monotonicity of Cn and intersection. □

Two aspects of this lemma and proof should be noticed. Firstly, a lemma of this character showing $A \div' \alpha = A \div \alpha$ for all α is exactly what is needed to show that the intersubstitutivity condition is simply a mistake in citation and not a mistake in mathematics. It is not obvious how any other technique could be used to show the classes of intersubstitutive safe contractions and plain safe contractions are equal.

Secondly, the technique in the first part of the proof is the only way one can hope to construct an intersubstitutive hierarchy from a plain hierarchy for this purpose. It is clear the construction of the intersubstitutive hierarchy must start with the plain hierarchy and that hierarchy can be modified in only two ways: relationships can be added to force intersubstitutivity ($\alpha \prec' \beta$ if and only if $\alpha' \prec \beta'$ for *any* equivalent α' and β'), or relationships can be removed to force intersubstitutivity as above ($\alpha \prec' \beta$ if and only if $\alpha' \prec \beta'$ for *all* equivalent α' and β'). But the first approach cannot work because, using it, one may in general introduce cycles, violating the definition of a hierarchy. Thus, we are left with the second approach, as used in the above proof.

So it is clear an extension of the above argument that shows $A \div \alpha \subseteq A \div' \alpha$ is needed, but this is not so easy. The direction above is easy to show because in defining the new hierarchy we remove relationships. This introduces more minimal elements, which means more things could be removed. However the direction to be shown requires us to show that everything which could be removed is added back via the saturation construction. It is difficult to see how this could be the case, but I have not been able to produce a counterexample. The saturation step requires one to work with $\beta \in \mathrm{Cn}(A/\alpha)$ instead of simply $\beta \in A/\alpha$, and this is much more difficult.

Another reason for doubt is that intersubstitutivity is not implied by acyclicity at the level of ksfs.

Lemma 3.1.4. *There is a ksf on a theory such that any relation on which that ksf is based is not an intersubstitutive hierarchy.*

Proof. We will construct the ksf. Let p, q be independent beliefs and $A = \mathrm{Cn}(\{p, q\})$. $\{p, q\}$ and $\{p \wedge p, q\}$ are both elements of $A \perp\!\!\!\perp (p \wedge q)$. Define s such that $s(\{p, q\}) = \{p\}$ but $s(\{p \wedge p, q\}) = \{q\})$. This is clearly permitted by Definition 2.5.13. Let $\prec$ be a relation such that s is based on $\prec$. Clearly $p \prec q$ and $q \prec p \wedge p$. Then, either $p \wedge p \prec q$ or not $p \wedge p \prec q$. In the first case, $\prec$ is not acyclic, in the second case, $\prec$ is not intersubstitutive as obviously $p \leftrightarrow (p \wedge p) \in \mathrm{Cn}(\emptyset)$. Thus, in either case, $\prec$ is not a intersubstitutive hierarchy. □

This is an open question which should be resolved. My intuition is that a counterexample will be found, perhaps only for sets which fail to be logically closed. The class of intersubstitutive safe contractions, should it fail to be equal to the class of all safe contractions, is also an interesting class of contractions worthy of study.

3.2 Safe contractions on arbitrary sets

The fact that that safe contractions can be defined in the kernel contraction framework gives us reason to study their action on arbitrary sets which may fail to be theories. As reported in Section 2.4, Alchourrón and Makinson only studied safe contractions over theories, but the kernel contraction framework allows us to expand our inquiries. (Of course, the structure of the proof in the previous section shows that one could have studied the original safe contractions over non-closed sets just as well.)

Over theories, Theorem 2.4.4 shows us that safe contractions are all partial meet contractions. Over arbitrary sets, their definition implies that safe contractions are all smooth kernel contractions (because they are saturated kernel contractions and by Theorem 2.5.12), and Lemma 2.5.8 shows us that all partial meet contractions are also smooth kernel contractions.

It is natural to ask, then, whether safe contractions over arbitrary sets are also partial meet contractions. By Theorem 2.2.8, this is the case if and only if they satisfy success, inclusion, relevance, and uniformity. By definition and Theorem 2.5.10, safe contractions satisfy success, inclusion, core retainment, relative closure, and uniformity. Therefore, the question is simple.

Remark 3.2.1. Safe contractions over arbitrary sets form a subclass of partial meet contractions over arbitrary sets if and only if they satisfy relevance.

Lemmas 2.5.5 and 2.5.8 show that relevance implies core retainment and relative closure, but in the general case the converse is not true (not every saturated kernel contraction is a partial meet contraction, as seen on p. 91 of [16]). The only difference between a safe contraction and a saturated kernel contraction is the hierarchy on which the safe contraction is based, and it is difficult to see how the hierarchy could help with the crucial step that changes a proof of satisfaction of core retainment to a proof of satisfaction of relevance. Recall the definitions of the two properties:

core retainment If $\beta \in A$ and $\beta \notin A \div \alpha$, then there is A' such that $A' \subseteq A$ and that $\alpha \notin \mathrm{Cn}(A')$ but $\alpha \in \mathrm{Cn}(A' \cup \{\beta\})$.

relevance If $\beta \in A$ and $\beta \notin A \div \alpha$, then there is A' such that $A \div \alpha \subseteq A' \subseteq A$ and that $\alpha \notin \mathrm{Cn}(A')$ but $\alpha \in \mathrm{Cn}(A' \cup \{\beta\})$.

Both core retainment and relevance require that for any element β of a set A which fails to be in $A \div \alpha$, there be a subset A' of A which fails to imply α alone, but for which $A' \cup \{\beta\}$ implies α. In the case of safe contractions, like kernel contractions in general, β failing to be in $A \div \alpha$ means that it was a member of some subset B of A such that B implies α but $B \setminus \{\beta\}$ does not. Using $B \setminus \{\beta\}$ for A' is a straightforward way to satisfy core retainment.

However, relevance places the additional requirement that the subset A' in question is a superset of $A \div \alpha$, which $B \setminus \{\beta\}$ may not be. Further, the obvious idea of using $(A \div \alpha) \cup (B \setminus \{\beta\})$ to get around this problem fails. Although this A' will certainly satisfy that $A' \cup \{\beta\}$ implies α, it cannot be shown that it does not imply α before β is added. This is why smooth kernel contractions are characterized only by core retainment and not relevance.

It turns out that the generation by a hierarchy does not sufficiently restrict the behavior of the contraction to satisfy relevance. The main idea of the following counterexample was supplied by Horacio Arló-Costa in correspondence; I further simplified it.

Theorem 3.2.2. *There is a safe contraction which fails to satisfy relevance.*

Proof. Take $A = \{p, p \to q, q, q \to r\}$. Define a hierarchy on A such that the only relationship is $q \to r \prec q$. Let $\div$ be the safe contraction on A generated by $\prec$.

First, notice $A \perp\!\!\!\perp r = \{\{p, p \to q, q \to r\}, \{q, q \to r\}\}$. All elements of the first r-kernel are $\prec$-minimal, q is the only element of the second r-kernel which fails to be $\prec$-minimal, and all elements of A appear in some r-kernel, so $A/r = \{q\}$. Therefore, $A \div r = A \cap \mathrm{Cn}(\{q\}) = \{p \to q, q\}$.

Notice then that $p \in A$ but $p \notin A \div r$. So relevance requires that there be an A' with $A \div r \subseteq A' \subseteq A$ such that $r \notin \mathrm{Cn}(A')$ but $r \in \mathrm{Cn}(A' \cup \{p\})$. There are four A' which satisfy the subset requirements, and all fail the consequence requirements:

$$r \notin \mathrm{Cn}(\{p \to q, q\} \cup \{p\})$$
$$r \notin \mathrm{Cn}(\{p, p \to q, q\} \cup \{p\})$$
$$r \in \mathrm{Cn}(\{p \to q, q, q \to r\})$$
$$r \in \mathrm{Cn}(\{p, p \to q, q, q \to r\})$$

Therefore, there is no possible A' and relevance is falsified. (To see why core retainment is still satisfied in this case, notice that $\{p \to q, q \to r\} \subseteq A$ is such that $r \notin \mathrm{Cn}(\{p \to q, q \to r\})$ but $r \in \mathrm{Cn}(\{p \to q, q \to r\} \cup \{p\})$. It is just that $A \div r \not\subseteq \{p \to q, q \to r\}$.) □

Corollary. *Not all safe contractions are partial meet contractions.*

This confirms that the already difficult behavior of safe contractions on theories is even more complex on arbitrary sets; in the case of theories, safe contractions form a subclass of the partial meet contractions (moreover, we shall see that this relationship is strict in the next section), but in the general case neither forms a subclass of the other.

3.3 Cumulative kernel contractions and saturation

The result presented in Section 3.1 that safe contractions can be studied as a subclass of saturated kernel contractions is quite useful. With it, we have several other kinds of contractions which are related to safe contractions, and we can explore their properties.

The first area I investigated is the relationship between saturated kernel contractions and cumulative kernel contractions. Hansson notes in Section 2.9 of [16] that not all kernel contractions are cumulative, but his example is stated for a set that is not closed, and makes no reference to saturation or smoothness.

The counterexample is as follows. Let $A = \{p, p \to q, q \to p\}$ for logically independent p and q. Then $A \perp\!\!\!\perp q = \{\{p, p \to q\}\}$ and $A \perp\!\!\!\perp (p \leftrightarrow q) = \{\{p, p \to q\}, \{p \to q, q \to p\}\}$. Define σ an incision for A such that $\sigma(A \perp\!\!\!\perp q) = \{p \to q\}$ and $\sigma(A \perp\!\!\!\perp (p \leftrightarrow q)) = \{p, q \to p\}$. But if σ cumulated any ksf s, then $s(\{p, p \to q\}) = \{p \to q\}$ and therefore $p \to q \in \sigma(A \perp\!\!\!\perp (p \leftrightarrow q))$ which contradicts the definition of σ. So σ does not cumulate any ksf ([16], p. 93).

I extended this counterexample, adding saturation and basing it on a theory in the below result.

Theorem 3.3.1. *There is a saturated kernel contraction $\div$ for a theory A which is not a cumulative kernel contraction.*

Proof. We will exhibit such a kernel contraction. Let p and q be logically independent beliefs, and suppose $A = \mathrm{Cn}(\{p, q\})$.

Notice that every $\{\alpha, \beta\}$ such that α is equivalent to p and β is equivalent to $p \to q$ is in $A \perp\!\!\!\perp q$, and any such β must be in some element of some entailment set of that form. So, define σ such that for any sentence β such that $\beta \leftrightarrow (p \to q) \in \mathrm{Cn}(\emptyset)$, $\beta \in \sigma(A \perp\!\!\!\perp q)$. Now notice that for any $X \in A \perp\!\!\!\perp (p \leftrightarrow q)$, if $\beta \in X$ such that $\beta \leftrightarrow (p \to q) \in \mathrm{Cn}(\emptyset)$, then $X = \{\beta, \delta\}$ for some δ such that either $\delta \leftrightarrow p \in \mathrm{Cn}(\emptyset)$

or $\delta \leftrightarrow (q \to p) \in \mathrm{Cn}(\emptyset)$. Therefore, we may define σ such that for every such β, $\beta \notin \sigma(A \perp\!\!\!\perp (p \leftrightarrow q))$.

We may further require that σ be smooth, since for any subset $B \subseteq A$ comprised of sentences equivalent to $p \to q$, that subset entails nothing except other sentences equivalent to $p \to q$, so smoothness is vacuously satisfied by σ acting on those elements, and we may assume it is satisfied by σ's action on all other elements without contradicting our definition above.

But this σ is not cumulative. Assume for contradiction that s is a ksf such that σ cumulates s. Then since every β equivalent to $p \to q$ is in $\sigma(A \perp\!\!\!\perp q)$, for every $\{\alpha, \beta\}$ such that α is equivalent to p, $s(\{\alpha, \beta\}) = \{\beta\}$ by the definition of a ksf. But $\{\alpha, \beta\} \in A \perp\!\!\!\perp (p \leftrightarrow q)$, as already observed. So, by the definition of cumulation, $\beta \in \sigma(A \perp\!\!\!\perp (p \leftrightarrow q))$, contradicting our definition of σ. So σ cannot be cumulative.

Thus, the kernel contraction generated by σ is not cumulative, but by Theorem 2.5.12, that kernel contraction is saturated. □

Corollary. *The class of saturated kernel contractions is not a subclass of the class of cumulative kernel contractions. The class of saturated kernel contractions over theories is not a subclass of the class of cumulative kernel contractions over theories.*

In some sense, this is unsurprising. Kernel contractions fail to be cumulative because their action is "too different" on inputs that are "too similar." One could hope that requiring them to respect closure via smoothness removed all of these differences, but it is not so. Unfortunately, we also have the following.

Theorem 3.3.2. *There is a cumulative kernel contraction on a theory which is not saturated.*

Proof. We will exhibit such a kernel contraction. Let p and q be logically independent beliefs, and suppose $A = \mathrm{Cn}(\{p, q\})$. Notice that we may define s such that $p \notin \sigma(A \perp\!\!\!\perp (p \wedge q))$ for σ the cumulation of s. (This is because p is not equivalent to $p \wedge q$, thus $\{p\} \notin A \perp\!\!\!\perp (p \wedge q)$, so we may always select other elements and ensure p is not selected.) Since $\{p \wedge p, q\} \in A \perp\!\!\!\perp (p \wedge q)$ as well, we may define s such that $p \wedge p \in \sigma(A \perp\!\!\!\perp (p \wedge q))$. But then σ is not smooth, as $\{p\} \vdash p \wedge p$, but $\{p\} \wedge \sigma(A \perp\!\!\!\perp (p \wedge q)) = \emptyset$. Thus, we have defined a σ which is cumulative and not smooth (and thus not saturated, by Theorem 2.5.12). □

Corollary. *The class of cumulative kernel contractions is not a subclass of the class of saturated kernel contractions. The class of cumulative kernel contractions over theories is not a subclass of the class of saturated kernel contractions over theories.*

This is not that surprising either, though it is similarly unfortunate.

These two results establish that:

1. The classes of saturated and cumulative kernel contractions are both strict subclasses of the class of kernel contractions.

2. Neither class is a subclass of the other.

3. All of these relationships are preserved when we restrict our attention exclusively to kernel contractions over theories.

Another question immediately comes to mind: are these classes (essentially) disjoint? Of course there will be trivial contractions, say on $\mathrm{Cn}(\emptyset)$, which will be definable both as saturated and cumulative kernel contractions, but the existence of nontrivial saturated cumulative kernel contractions must be shown. We will see an example of such a contraction in the next section, as a consequence of Theorem 3.4.2.

Where do safe contractions lie, given these results? One could hope that safe contractions were saturated cumulative contractions, since they are by definition saturated kernel contractions, and they are the saturations of cumulative kernel contractions. So, in essence, we may ask: does saturation preserve cumulativity? I have been unable to get a full result, but the following lemma reduces the question to whether or not we can construct an incision of a particular kind:

Lemma 3.3.3. *Let $\approx_\sigma$ be a kernel contraction for a theory A, and let τ be an incision function for A. Then τ generates $\hat{\approx}_\sigma$ if and only if for all α, $\tau(A \perp\!\!\!\perp \alpha) = \sigma(A \perp\!\!\!\perp \alpha) \setminus \mathrm{Cn}(A \approx_\sigma \alpha)$.*

Proof. Let A, $\approx_\sigma$, and τ be as indicated. Notice that by Definition 2.5.3, τ generates $\hat{\approx}_\sigma$ if and only if for all α we have that $A \hat{\approx}_\sigma \alpha = A \setminus \tau(A \perp\!\!\!\perp \alpha)$. So our claim is that the above is equivalent to for all α that $\tau(A \perp\!\!\!\perp \alpha) = \sigma(A \perp\!\!\!\perp \alpha) \setminus \mathrm{Cn}(A \approx_\sigma \alpha)$. It's enough to show this holds for every individual α, so fix α. The following are all equivalent:

$$\begin{array}{rcl}
\tau(A \perp\!\!\!\perp \alpha) & = & \sigma(A \perp\!\!\!\perp \alpha) \setminus \mathrm{Cn}(A \approx_\sigma \alpha) \\
A \setminus \tau(A \perp\!\!\!\perp \alpha) & = & A \setminus (\sigma(A \perp\!\!\!\perp \alpha) \setminus \mathrm{Cn}(A \approx_\sigma \alpha)) \\
A \setminus \tau(A \perp\!\!\!\perp \alpha) & = & (\mathrm{Cn}(A \approx_\sigma \alpha) \cap A) \cup (A \setminus \sigma(A \perp\!\!\!\perp \alpha)) \\
A \setminus \tau(A \perp\!\!\!\perp \alpha) & = & (A \cap \mathrm{Cn}(A \approx_\sigma \alpha)) \cup (A \approx_\sigma \alpha) \\
A \setminus \tau(A \perp\!\!\!\perp \alpha) & = & A \cap \mathrm{Cn}(A \approx_\sigma \alpha) \\
A \setminus \tau(A \perp\!\!\!\perp \alpha) & = & A \hat{\approx}_\sigma \alpha
\end{array}$$

Line 1 is equivalent to line 2 as both $\tau(A \perp\!\!\!\perp \alpha)$ and $\sigma(A \perp\!\!\!\perp \alpha)$ must be subsets of A since they are incisions for A. Line 2 is equivalent to line 3 by a set theoretic identity. Line 3 is equivalent to line 4 since $A \approx_\sigma \alpha$ must be a subset of both A and $\mathrm{Cn}(A \approx_\sigma \alpha)$ since kernel contraction satisfies success and Cn satisfies monotonicity. The rest of the equivalences are by definition. $\square$

Using this lemma to show saturation preserves cumulativity would be an extremely fruitful result.

Remark 3.3.4. If for an incision σ which cumulates a ksf s we can define a ksf t such that t is cumulated by the incision which takes value $\sigma(A \perp\!\!\!\perp \alpha) \setminus \mathrm{Cn}(A \approx_\sigma \alpha)$ for each α, then:

1. Saturation preserves cumulativity.
2. Safe contractions form a subclass of cumulative kernel contractions (they form a subclass of saturated kernel contractions by definition).
3. Moreover, safe contractions form a strict subclass of cumulative kernel contractions (since we know every safe contraction is saturated and not all cumulative kernel contractions are saturated).
4. In fact, safe contractions form a subclass of saturated cumulative kernel contractions.
5. All of the above relationships hold when we restrict our attention to kernel contractions over theories.

This is perhaps the open question with the most immediate and interesting consequences that my research has uncovered.

The above remark mentions that safe contractions would form a (strict, in fact) subclass of cumulative kernel contractions if saturation preserves cumulativity. I also investigated the relationship between safety and cumulativity separately, which is the subject of the next section.

3.4 Relational cumulative kernel contractions

Safe contractions are saturations of cumulative kernel contractions which are based on hierarchies. At first glance, it might seem that the cumulation, rather than the relation, is doing all the work since the requirement of finite acyclicity is so weak. As previously noted, finite acyclicity is required so that the relation actually generates a ksf; to generate a ksf from a relation $<$ at least one element must be $<$-minimal in each inclusion-minimal set which implies a sentence. By compactness, every-inclusion minimal set which implies a sentence is finite, so if there are no finite cycles, every set must have a $<$-minimal element.

Of course, only some sets will be inclusion-minimal; the elements of those sets are all independent. Therefore, we could have weakened the definition of a hierarchy to have no cycles among independent elements, as seen below.

Lemma 3.4.1. *Let s be a ksf for a set A. Then any relation on which s is based has no cycles among independent elements of A.*

Proof. Let A and s be as indicated. Let $\prec$ be a relation such that s is based on $\prec$. Suppose for contradiction there are $\alpha_1, \ldots, \alpha_n$ all independent beliefs such that $\alpha_1 \prec \cdots \prec \alpha_n \prec \alpha_1$. Notice that as they are all independent, $\{\alpha_1, \ldots, \alpha_n\} \in A \perp\!\!\!\perp (\alpha_1 \wedge \cdots \wedge \alpha_n)$. Thus, since s is a ksf and that set is nonempty, $s(\{\alpha_1, \ldots, \alpha_n\}) \neq \emptyset$. But since s is based on $\prec$, $s(\{\alpha_1, \ldots, \alpha_n\}) = \emptyset$ by Definition 2.5.15. This is a contradiction, so there can be no such $\alpha_1, \ldots, \alpha_n$. □

Corollary. *s is a relational ksf for a set A if and only if s is based on a relation $\prec$ on A which has no cycles among independent elements of A.*

Notice that the structure of the proof critically depends on the fact that the elements in question are independent of one another. One cannot derive the same contradiction if the elements are dependent. Yet even the weak requirements placed on hierarchies do restrict the behavior of contractions.

Theorem 3.4.2. *There is a smooth and cumulative kernel contraction for a theory which is not based on a hierarchy.*

Proof. We will construct the incision on which this contraction is based. Let p, q, r be logically independent beliefs and let $A = \mathrm{Cn}(\{p, q, r\})$. Let σ be an incision function for A such that:

$$\begin{array}{rcl} p & \notin & \sigma(A \perp\!\!\!\perp (p \wedge q)) \\ q & \notin & \sigma(A \perp\!\!\!\perp (q \wedge r)) \\ r & \notin & \sigma(A \perp\!\!\!\perp (p \wedge r)) \end{array}$$

We can do this while preserving smoothness as if the consequences of p (resp. q and r) are not in $\sigma(A \perp\!\!\!\perp (p \wedge q))$ (resp. the corresponding incisions for $q \wedge r$ and $p \wedge r$) then there will be no elements present which are entailed by the singleton $\{p\}$ (resp. $\{q\}$ and $\{r\}$), and thus we may maintain smoothness without the presence of p (resp. q and r).

We can do this while preserving cumulativity because no set will be a member of the kernel set of any two of $p \wedge q$, $q \wedge r$, and $p \wedge r$ by the assumption of independence. So there will clearly be ksfs which this cumulates.

However, any ksf which is cumulated by this incision cannot be based on a hierarchy, as we have that for any such ksf s:

$$\begin{array}{rcl} s(\{p, q\}) & = & \{q\} \\ s(\{q, r\}) & = & \{r\} \\ s(\{p, r\}) & = & \{p\} \end{array}$$

This is so because clearly each of these doubletons is present in the entailment set for the incisions above, so the value of the ksf applied to them must be nonempty and cannot include the other element.

But then any $\prec$ on which this ksf is based must be such that $p \prec r \prec q \prec p$, contradicting acyclicity, and thus must not be a hierarchy. Therefore, σ cannot be based on a hierarchy. □

Corollary. *The class of presafe contractions is a strict subclass of the class of cumulative kernel contractions. The class of saturated presafe contractions is a strict subclass of the class of saturated cumulative kernel contractions. Both of these relationships are preserved if we restrict our attention to kernel contractions over theories.*

This also implies that the class of save contractions over theories is a strict subclass of the class of smooth kernel contractions, and thus of the partial meet contractions. Therefore, safe contractions over theories cannot be characterized by the basic Gärdenfors postulates.

Even more distressingly:

Theorem 3.4.3. *There is a smooth and cumulative incision for a theory which is not based on any relation.*

Proof. In the proof of the above Theorem 3.4.2, we constructed a ksf such that any relation which that ksf is based on must have a cycle, and thus cannot be a hierarchy. However, by Theorem 3.4.1, if a ksf is based on a relation, then that relation has no cycles among independent elements. The elements used in the above proof are independent, thus, that ksf is not based on any relation. □

Corollary. *The class of relational cumulative kernel contractions is a strict subclass of the class of cumulative kernel contractions. The class of saturated relational cumulative kernel contractions is a strict subclass of the class of the saturated cumulative kernel contractions. Both of these relationships are preserved if we restrict our attention to kernel contractions over theories.*

Last section, we saw that the classes of saturated kernel and cumulative kernel contractions were strict subsets of the class of kernel contractions, and neither were subsets of each other (and these relationships were preserved over theories). Now it is clear also that they have a nonempty intersection, and that the classes of relational cumulative kernel and presafe contractions are successively more restrictive subclasses of cumulative kernel contractions, and saturated relational cumulative kernel and safe contractions are successively more restrictive subclasses of saturated kernel contractions.

What about the relationship between saturation and hierarchies? One could hope that at least the incisions generated by hierarchies were well-behaved, and that the saturation step in the definition was superfluous. However, it turns out to be wholly necessary:

Theorem 3.4.4. *There is a presafe contraction for a theory which is not smooth.*

Proof. We will construct the incision on which the presafe contraction is based. Let p, q, r be logically independent beliefs and let $A = \mathrm{Cn}(\{p, q, r\})$. Let $\prec$ be a hierarchy for A such that $r \prec r \to q$, $r \to q \prec p \to r$, and $r \to q \prec p$. (None of these are equivalent by assumption, so we may assume that the rest of $\prec$ is defined such that it is acyclic and intersubstitutive.) Let s be a ksf based on $\prec$, and let σ be the cumulation of s.

Clearly $\{p, p \to r\} \subseteq A$, and $\{p, p \to r\} \vdash r$. Also, clearly $r \in \sigma(A \perp\!\!\!\perp (p \to q))$ as $\{r, r \to q\} \in A \perp\!\!\!\perp (p \to q)$ and $r \prec r \to q$. But $\{p, p \to r\} \cap \sigma(A \perp\!\!\!\perp (p \to q)) = \emptyset$, as any element of $A \perp\!\!\!\perp (p \to q)$ containing p and $p \to r$ must contain $r \to q$, and thus neither p nor $p \to r$ will be selected for removal as $r \to q \prec p$ and $r \to q \prec p \to r$.

This contradicts the definition of smoothness, so σ is not smooth. □

Corollary. *If saturation preserves being based on a hierarchy, the class of safe contractions is a strict subclass of the class of presafe contractions. This relationship would be preserved when restricting our attention to kernel contractions over theories.*

The situation is then extremely complex. Every piece in the definition of safe contractions within the kernel contraction framework does some work: the fact that the ksf is based on a hierarchy restricts its behavior, the fact that the incision is cumulative restricts its behavior further, and the fact that the contraction is then saturated even further restricts its behavior.

Let us turn our attention now to the relationship between safe contraction and the supplementary postulates.

3.5 Safety, theories, and overlap

There are interesting connections between safe contractions and the supplementary postulates, as first explored in Theorem 2.4.9 due to Hans Rott. However, via a lemma noticed (in a different form) even by Alchourrón and Makinson, incisions based on hierarchies have a property very close to conjunctive overlap. I developed this very slightly, giving the result below:

Theorem 3.5.1. *If $\div$ is a presafe contraction, then it satisfies conjunctive overlap.*

Proof. Let A be set, and let $<$ be a hierarchy on A. Let σ be the incision function based on $<$. The first below line is given by a lemma in [4], and the following lines are equivalent to it by set theory and definition:

$$\begin{aligned}
\sigma(A \perp\!\!\!\perp (\alpha \wedge \beta)) &\subseteq \sigma(A \perp\!\!\!\perp \alpha) \cup \sigma(A \perp\!\!\!\perp \beta) \\
A \setminus (\sigma(A \perp\!\!\!\perp \alpha) \cup \sigma(A \perp\!\!\!\perp \beta)) &\subseteq A \setminus \sigma(A \perp\!\!\!\perp (\alpha \wedge \beta)) \\
(A \setminus \sigma(A \perp\!\!\!\perp \alpha)) \cap (A \setminus \sigma(A \perp\!\!\!\perp \beta)) &\subseteq A \approx_\sigma (\alpha \wedge \beta) \\
(A \approx_\sigma \alpha) \cap (A \approx_\sigma \beta) &\subseteq A \approx_\sigma (\alpha \wedge \beta)
\end{aligned}$$

Thus, $\approx_\sigma$ satisfies conjunctive overlap, as required. □

The presafe contraction which is saturated to become a safe contraction must satisfy conjunctive overlap. However, it is not clear whether or not its saturation must: it may be the case that saturating $A \approx_\sigma \alpha$ and saturating $A \approx_\sigma \beta$ yields beliefs not present in the saturation of their intersection. In any case, I have not been able to find a completely straightforward extension of the above result to saturations. Instead, there is the following.

Lemma 3.5.2. *If $\approx_\sigma$ is a kernel contraction for a theory A that satisfies closure and conjunctive overlap, then its saturation $\mathrel{\hat{\approx}}_\sigma$ satisfies conjunctive overlap.*

Proof. Let A be a theory, and $\approx_\sigma$ be a kernel contraction for A satisfying closure and conjunctive overlap. Let beliefs α and β be given. Consider $A \mathrel{\hat{\approx}}_\sigma \alpha$ and $A \mathrel{\hat{\approx}}_\sigma \beta$. By set theory and definitions:

$$\begin{aligned}(A \mathrel{\hat{\approx}}_\sigma \alpha) \cap (A \mathrel{\hat{\approx}}_\sigma \beta) &= (A \cap \mathrm{Cn}(A \approx_\sigma \alpha)) \cap (A \cap \mathrm{Cn}(A \approx_\sigma \beta)) \\ &= A \cap (\mathrm{Cn}(A \approx_\sigma \alpha) \cap \mathrm{Cn}(A \approx_\sigma \beta)) \\ &= A \cap \mathrm{Cn}(A \approx_\sigma \alpha \cap A \approx_\sigma \beta)\end{aligned}$$

The last step is the step where the requirement of closure (and thus also that A is a theory) is used. Further, notice that, since conjunctive overlap is satisfied, we can do the following:

$$\begin{aligned}A \approx_\sigma \alpha \cap A \approx_\sigma \beta &\subseteq A \approx_\sigma (\alpha \wedge \beta) \\ \mathrm{Cn}(A \approx_\sigma \alpha \cap A \approx_\sigma \beta) &\subseteq \mathrm{Cn}(A \approx_\sigma (\alpha \wedge \beta)) \\ A \cap \mathrm{Cn}(A \approx_\sigma \alpha \cap A \approx_\sigma \beta) &\subseteq A \cap \mathrm{Cn}(A \approx_\sigma (\alpha \wedge \beta)) \\ A \cap \mathrm{Cn}(A \approx_\sigma \alpha \cap A \approx_\sigma \beta) &\subseteq A \mathrel{\hat{\approx}}_\sigma (\alpha \wedge \beta)\end{aligned}$$

Putting both arguments together, we get that $\mathrel{\hat{\approx}}_\sigma$ satisfies conjunctive overlap, as required. □

Corollary. *If presafe contractions over theories satisfy closure, then all safe contractions over theories satisfy conjunctive overlap.*

This is an interesting almost-result. Unfortunately, its extension to the real desideratum turns out to be impossible.

Lemma 3.5.3. *There is a presafe contraction for a theory which fails to satisfy closure.*

Proof. By Theorem 3.4.4, there is a presafe contraction on a theory which is not a smooth kernel contraction. Therefore, since smooth kernel contractions on theories are characterized exactly by the basic Gärdenfors postulates, which include closure, there is a presafe contraction on a theory which fails to satisfy closure. □

There may be a way to modify this technique to no longer require closure and achieve the desired result, but it seems doubtful. Conjunctive overlap is tantalizingly close for safe contractions, but it seems that they need continuing-up or continuing-down to avoid adding undesirable beliefs to contractions of conjuncts in the saturation step, as seen in Theorem 2.4.6.

Chapter 4

Conclusions

To recap, we have seen that every step in the construction of safe contractions in the kernel contraction framework is crucial: there is a saturated kernel contraction which is not cumulative, there is a saturated cumulative kernel contraction which is not based on a hierarchy, and there is a cumulative kernel contraction based on a hierarchy which is not saturated. We have also seen that although safe contractions form a subclass of partial meet contractions in the case of theories, this relationship is not present in the general case. Also, even in the case of theories, this relationship is strict and we cannot hope to reverse it to provide a characterization.

Along the way, I have isolated the classes of cumulative, relational cumulative, and presafe contractions (and their saturations), which have not been studied in the literature thus far (excepting of course the saturations of presafe contractions, which are safe contractions by definition). These have complex behavior, and are worthy of further study.

At this point, it is worthwhile to summarize all the characterization results known thus far. I shall list all the relationships which are known, from the largest classes to the smallest.

First, for the simpler case of contractions over theories. We will focus only on saturated classes, since only then are we guaranteed closure.

1. Smooth kernel contractions, saturated kernel contractions, partial meet contractions, and basic entrenchment contractions are equivalent and characterized by the basic Gärdenfors postulates.
2. Saturated cumulative kernel contractions are a strict subclass of the saturated kernel contractions, no characterization result is known.
3. Saturated relational cumulative kernel contractions are a strict subclass of the saturated cumulative kernel contractions, no characterization result is known.

4. Safe contractions are a strict subclass of the saturated relational cumulative kernel contractions, no characterization result is known.
5. Regular and virtually connected safe contractions, transitively relational partial meet contractions, and entrenchment contractions are all equivalent, characterized by the basic and supplementary Gärdenfors postulates, and a strict subclass of the safe contractions.

The situation only worsens in the case of contractions over arbitrary sets. Very few of these classes have characterization results, so I will only note when they do.

1. Kernel contractions are the largest class, and are characterized by success, inclusion, uniformity, and core-retainment.
2. Smooth kernel contractions and saturated kernel contractions are equivalent and characterized by success, inclusion, uniformity, core-retainment, and relative closure, and are a strict subclass of the kernel contractions.
3. Cumulative kernel contractions are a strict subclass of the kernel contractions.
4. Relational cumulative kernel contractions are a strict subclass of the cumulative kernel contractions.
5. Presafe contractions are a strict subclass of the relational cumulative kernel contractions.
6. Saturating the above three classes preserves their relationships: saturated cumulative contractions are a strict superclass of the saturated relational cumulative contractions which are a strict superclass of the safe contractions.
7. Regular and virtually connected safe contractions are a strict subclass of the safe contractions.
8. Partial meet contractions are characterized by success, inclusion, uniformity, and relevance, and are a strict subclass of the smooth kernel contractions.
9. Transitively maximizingly relational partial meet contractions are a strict subclass of the partial meet contractions.

We have also seen an early indication that conjunctive overlap does not hold for all safe contractions, and that not all safe contractions are intersubstitutive. There are now many more relationships which still have yet to be shown.

Several important questions have been raised here. In a rough order of importance, they are:

1. Does saturation preserve cumulativity? (It seems likely, using Lemma 3.3.3.)
2. Does saturation preserve being based on a hierarchy (i.e. are the safe contractions exactly the saturated presafe contractions)?
3. Does saturation of presafe contractions preserve conjunctive overlap, even when they fail to satisfy closure? (It seems unlikely, as described in Section 3.5, but a counterexample should be found.)
4. Are all safe contractions intersubstitutive safe contractions? (Presumably not, at least for arbitrary sets, as discussed in Section 3.1, but a counterexample is needed.)
5. Does Makinson's worry about revision equivalence apply to safe contractions? (See Section 2.4.2.)

Returning to the original project, though, it is clear that we will have to take a different approach in order to characterize safe contraction. A postulate is neeeded which is reasonably strong and which is "relational" like the supplementary postulates are, but which does not entail them. Notice that we have also demonstrated that the basic postulates are not sufficient to characterize cumulative kernel contractions; it would be good to find a characterization. Combined with saturation preserving cumulativity, this would give us additional postulates that safe contractions all satisfy.

In any case, the kernel contraction framework and its subclasses has been shown to be the most complex situation in AGM theory by far. It is clear there is still much work to be done in this area.

Bibliography

[1] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[2] C. E. Alchourrón and D. Makinson. Hierarchies of regulation and their logic. In Hilpinen, editor, *New Studies in Deontic Logic*, pages 125–148. D. Reidel Publishing Company, Dordrecht, the Netherlands, 1981.

[3] C. E. Alchourrón and D. Makinson. On the logic of theory change: Conctraction functions and their associated revision functions. *Theoria*, 48:14–37, 1982.

[4] C. E. Alchourrón and D. Makinson. On the logic of theory change: Safe contraction. *Studia Logica*, 44:405–422, 1985.

[5] C. E. Alchourrón and D. Makinson. Maps between some different kinds of contraction function: The finite case. *Studia Logica*, 45:197–198, 1986.

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

[7] P. Gärdenfors. Rules for rational changes of belief. In T. Pauli, editor, *Philosophical Essays Dedicated to Lennart Åqvist on His Fiftieth Birthday*, pages 88–101. Uppsala University, Uppsala, Sweden, 1982.

[8] P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In M. Vardi, editor, *Theoretical Aspects of Reasoning About Knowledge*, pages 83–95. Morgan Kaufmann, Los Altos, California, 1988.

[9] S. O. Hansson. Belief contraction without recovery. *Studia Logica*, 50:251–260, 1991.

[10] S. O. Hansson. A dyadic representation of belief. In P. Gärdenfors, editor, *Belief Revision*, pages 89–121. Cambridge University Press, Cambridge, UK, 1992.

[11] S. O. Hansson. In defense of base contraction. *Synthese*, 91:239–245, 1992.

[12] S. O. Hansson. Changes on disjunctively closed bases. *Journal of Logic, Language and Information*, 2:255–284, 1993.

[13] S. O. Hansson. Reversing the Levi identity. *Journal of Philosophical Logic*, 22:637–669, 1993.

[14] S. O. Hansson. Kernel contraction. *Journal of Symbolic Logic*, 59(3):845–859, 1994.

[15] S. O. Hansson. Taking belief bases seriously. In D. Prawitz and D. Westerståhl, editors, *Logic and Philosophy of Science in Uppsala*, pages 13–28. Kluwer Academic Publishers, Dordrecht, the Netherlands, 1994.

[16] S. O. Hansson. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer Academic Publishers, Dordrecht, the Netherlands, 1999.

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

[18] I. Levi. *The Fixation of Belief and Its Undoing*. Cambridge University Press, Cambridge, UK, 1991.

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

[20] H. Rott. On the logic of theory change: More maps between different kinds of contraction function. In P. Gärdenfors, editor, *Belief Revision*. Cambridge University Press, Cambridge, UK, 1992.

[21] H. Rott. Basic entrenchment. *Studia Logica*, 73(2):257–280, 2003.