Funciones binarias para el cambio de teorías

Becher, Verónica Andrea

1999

Tesis presentada para obtener el grado de Doctor en Ciencias de la Computación de la Universidad de Buenos Aires

Departamento de Computación
Facultad de Ciencias Exactas y Naturales
Universidad de Buenos Aires

Binary Functions for Theory Change

Verónica A. Becher

A dissertation submitted in conformity with the requirements for the Degree of Doctor of Philosophy in the University of Buenos Aires.
Argentina
June 1999

To Carlos E. Alchourrón

Funciones Binarias para el Cambio de Teorías

Verónica A. Becher

Resumen

El problema del cambio sobre cuerpos de información es realmente interesante. La legislación se encuentra en constante modificación, nuevos descubrimientos modifican a las teorías científicas y los robots deben actualizar su representación del mundo cada vez que un sensor adquiere nuevos datos. La teoría del cambio de teorías ofrece un modelo para estos procesos bajo ciertas idealizaciones.

Se asume un lenguaje formal y una noción de consecuencia lógica. La nueva información es expresada como enunciados en el lenguaje lógico. Conjuntos de enunciados clausurados por la operación de consecuencia lógica, es decir, teorías, son modificadas por medio de funciones. Estas toman un conjunto y un enunciado y retornan el conjunto actualizado. Las funciones de cambio responden a un principio fundamental: consistencia lógica. El resultado de un cambio debe ser siempre un conjunto de enunciados mutuamente consistentes.

En 1985 Alchourrón, Gärdenfors y Makinson (en adelante AGM) fueron autores de lo que sería la referencia clásica sobre el tema [Alchourrón *et al.*, 1985]. Concibieron funciones que, bajo la máxima de consistencia, retornan teorías que preservan lo más posible de las teorías originales y contemplan la nueva información. Las teorías no deben ser modificadas más allá de lo necesario. La relación de inclusión entre conjuntos no es suficiente como criterio de mínima pérdida informacional porque, en general, hay inifinitas teorías mutuamente incomparables con respecto a la relación de inclusión. Por lo tanto puede resultar imposible seleccionar una única como la más preservativa. En consecuencia las funciones AGM deben realizar una selección no determinística o codificar algún otro criterio de selección.

Al menos en dos respectos la teoría AGM está indefinida. En primer lugar las teorías de cambio de teorías deben enfrentarse con *el problema de la iteración*

del cambio. Las funciones AGM modelan cambios singulares, toman una teoría y retornan una teoría actualizada, realizan un solo paso. Pero tarde o temprano, habrá otro cambio que inducirá una nueva teoría. Es decir, se deberá actualizar la teoría ya actualizada. Aunque el formalismo AGM no prohíbe la iteración de sus funciones omite toda especificación de cómo debe realizarse o cuáles son las propiedades del cambio sucesivo.

El otro frente indefinido en el formalismo AGM es *el problema del cambio en múltiples teorías*. Si dos teorías no son independientes, es de esperar que las operaciones de cambio respectivas tampoco lo sean. Por ejemplo, si un teoría está incluída en la otra podría esperarse que el cambio de la primera esté incluído en el cambio de la segunda. Algunas propiedades de coherencia deberían vincular la operación de cambio sobre distintas teorías. Este es el tema central de esta tesis.

Aunque las funciones AGM proveen una noción coherente de cambio para teorías tomadas separadamente, estas funciones no necesariamente son conjuntamente coherentes. Esta es una limitación seria del formalismo AGM y el presente trabajo está dedicado a superar esta limitación. Este problema no ha sido considerado en la teoría AGM y tampoco ha sido objeto de investigación en el área. Tras argumentar que las funciones AGM son en realidad funciones unarias (de aridad uno) relativas a la teoría a ser modificada, en este trabajo se proponen auténticas funciones binarias para el cambio de teorías. El término binario se refiere a funciones de dos argumentos, es decir de aridad dos.

Las funciones binarias resuelven el problema del cambio en múltiples teorías, y siendo definicionalmente simples, también resuelven en cierta medida el problema de la iteración del cambio. Dado que las funciones binarias están definidas para toda teoría, el resultado de aplicar una función es a su vez otra teoría que puede ser puesta como argumento de la *misma* función. En consecuencia, las funciones binarias inducen un esquema de cambio iterado que es determinístico respecto de los argumentos de la función. Este comportamiento, que ha sido interpretado como carente de memoria histórica, no siempre resulta deseable. Es una preocupación actual entre los investigadores del área la búsqueda de un modelo general de iteración, un único conjunto de postulados que gobiernen el cambio reiterado, en el mismo espíritu que los de AGM gobiernan cambios singulares. Luego de catorce años de la formulación de la teoría AGM se han planteado varias formalizaciones alternativas que difieren en sus virtudes y defectos, pero se desconoce si tales postulados únicos han de existir; tal vez no haya una única regularidad que deba ser expuesta.

Dentro de la teoría AGM existen dos funciones binarias que gozan de múltiples propiedades, pero se corresponden con casos límites de funciones de cambio aceptables; estas son la función de expansión y las funciones AGM *full meet* [Alchourrón and Makinson, 1982]. Fuera de la tradición AGM, Katsuno y Mendelzon han formalizado su operación de *update* [Katsuno and Mendelzon, 1992] como una función binaria para el cambio de teorías. Las operaciones de update y revisión AGM denotan dos tipos de cambio que han sido considerados fundamentalmente diferentes. La función de revisión se ha considerado propicia para modelizar el proceso de refinamiento o corrección de una representación de objetos que permanecen estáticos. En contraste, la operación de update modeliza la noción de cambio sobre la representación de objetos que están en evolución. En esta tesis se estudian en detalle las vinculaciones formales de ambas operaciones.

Apartándose de la tradición AGM la operación de update está definida como un conectivo binario sobre un lenguaje basado en un conjunto finito de variables proposicionales. Se demuestra que nada crucial depende de ésto, ya que es posible reformular la operación de update como una funcion binaria que toma una teoría y un enunciado y retorna una teoría. Sin embargo se exhibe un resultado inesperado: los postulados de Katsuno y Mendezon son incompletos para caracterizar la función de update para lenguajes proposicionales infinitos. Se provee un conjunto apropiado de postulados, reforzando los originales, y se demuestra el correspondiente teorema de representación para lenguajes posiblemente infinitos. De esta manera se extiende el trabajo original de Katsuno y Mendelzon que estaba definido solo para el caso finito. Los resultados encontrados completan y clarifican los de [Peppas and Williams, 1995], quienes ya habían notado que los postulados originales de update eran incompletos para lenguajes de primer orden. Adicionalmente se consigue que las operaciones de revisión AGM y la de update queden en una misma base definicional, permitiendo su comparación y mejor comprensión, cuando la naturaleza de la diferencia es aún una pregunta abierta en la literatura de lógica filosófica.

En este trabajo se proponen dos familias de funciones binarias que extienden el formalismo AGM: las funciones *AGM iterables* y las *AGM analíticas*. Ambas se definen sobre lenguajes posiblemente infinitos, mediante postulados que extienden a los de AGM y para ambas se demuestran teoremas de representación sobre distintas estructuras formales.

Las funciones AGM iterables tienen la peculiaridad de ser funciones casi constantes sobre el primer argumento cuando el segundo está fijo. A pesar de su simpleza proveen una fuerte noción de coherencia con respecto al cambio en distintas teorías. De acuerdo con las funciones AGM iterables el cambio en una teoría depende del cambio de la teoría más grande de todas, que es el conjunto de todos los enunciados del lenguaje. Se demuestra que las funciones AGM iterables satisfacen muchas propiedades, tanto para el cambio de múltiples teorás como para el cambio reiterado.

Las funciones AGM analíticas son funciones binarias de mayor complejidad definicional que las iterables. Son casi monótonas sobre su primer argumento cuando el otro está fijo, sin ser funciones constantes ni casi constantes. La operación de cambio analítica puede calcularse por medio de un análisis por casos, con la propiedad de que si una teoría es extensión de otra, los casos considerados para la primera son también casos para la segunda. Una subclase de funciones analíticas es la de las maxi-analíticas, cuya característica es que mapean teorías completas en teorías completas. Las funciones analíticas son candidatas interesantes para el cambio en distintas teorías y también satisfacen relevantes propiedades del cambio iterado.

Pero las funciones analíticas poseen además otro interés. Proveen una conexión formal entre la operación de update de Katsuno y Mendelzon y la revisón de AGM. La revisión AGM analítica se basa en el aparato semántico de update, y de este modo establece un puente entre dos formalizaciones aparentemente incomparables.

Por último la tesis provee un resultado de unificación de dos cálculos lógicos para la teoría AGM: las lógicas CO [Boutilier, 1992] y DFT [Alchourrón, 1995]. A partir de la noción de consecuencia lógica ambas lógicas pueden usarse para calcular cambios en diferentes teorías. Y aunque las dos son lógicas condicionales, difieren. La semántica de CO es relacional mientras que la de DFT no lo es. También difieren en la definición del conectivo condicional. Se demuestra que, bajo condiciones restrictivas apropiadas, las dos lógicas son equivalentes. En su rol de lógicas para el cambio de teorías, las ocurrencias anidadas del condicional sugieren una función de cambio que admite iteración. Pero resulta claro rapidamente que dicha función es trivial.

La tesis plantea direcciones de trabajo futuro, principalmente sugiriendo la definición de nuevas funciones binarias para el cambio de teorías y la posibilidad de proveer un calculo lógico para las funciones AGM iterables y analíticas.

Preface

This thesis is about the logic for representing changes, a topic in theoretical Artificial Intelligence. This work was performed in two stages, quite distant in time of one another. The first was from 1994 to 1996, after my studies in Canada, as a doctoral student of Carlos Alchourrón under a scholarship from the University of Buenos Aires. I met a privileged intellect, a true maestro and an affectionate man. All the problems and solutions I discuss in this thesis were initiated in those days, in long conversations, while having some tea. I deeply regret, and so does every one who met him, his passing away in January 1996. At that time David Makinson guided and encouraged my work. This prestigious and intimidating figure was a key appearance for my studies. I am extremely grateful to him for his wise recommendations and detailed teaching. Untiredly, he worked through my papers in various stages, corrected my proofs, added new results and sent me references.

The second stage was delayed until the end of 1997. Although I remained close to the theory of theory change, in part because of the process that involved my publications, in part because of being in a research project, I went into a crisis with the subject of my thesis. Mostly with by my close friend Carlos Areces we started to warry about decidability and expressibility in formal languages. We studied non-standard modal logics and renamable Horn sets, and wrote about "Characterization Results for d-Horn Formulas, or On formulas that are true on Dual Reduced Products", that will appear in Lecture Notes in Computer Science, CSLI series, in 1999.

Somehow, in January 1998, after understanding some old results of theory change I was ready to write my thesis on binary functions for theory change.

Undoubtedly Carlos Areces is the person I am most indebted to. I could not have written this thesis without his acute constructive criticism and his invaluable technical help. I specially acknowledge his total dedication to our work during his visits to Buenos Aires. And of course I thank Gladys Palau for her wise guiding and support, and my peers, Eduardo Fermé and Ricardo Rodriguez, for giving me the encouragement to write this thesis. I'm also grateful to Hans Rott for his insightful comments during his visit to our Department in April 1999, and to José Alvarez for his comments on a draft of this work. Lastly, I deeply thank the indispensable support I received from my husband and I apologize to my loving children for the innumerable hours without them.

Much of the contents of this thesis has been published or has been submitted for publication. Chapters 3 and 4 treat the problems discussed in "Iterable AGM functions", written in collaboration with Carlos Areces, which appear in Rott H. Williams M.(eds), Frontiers in Belief Revision, Kluwer Applied Logic Series, to be published in 1999.

The initial ideas of Chapter 6 date back to "Unified Semantics for Revision and Update, or the Theory of Lazy Update" which appeares in the Proceedings of the 24 Jornadas Argentinas de Informatica e Investigacion Operativa (JAIIO), pp.641-650, Buenos Aires, Argentina, 1995.

Chapter 7 expands ideas of the following two articles. "Two Conditional Logics for Defeasible Inference: A Comparison", which appeared in Notes in Artificial Intelligence, pp.49-58. Wainer, J. Carvalho A. Advances in Artificial Intelligence. Springer Verlag, 1995; and "Some Observations on Carlos Alchourrn's Theory of Defeasible Conditionals", written in collaboration with E.Fermé, S.Lazzer, R. Rodriguez, C. Oller and G.Palau; in Mc Namara P. Prakken H. (eds), Norms Logics and Information Systems, New Studies on Deontic Logic and Computer Science, IOS Press, 1998.

Chapters 6 and 7, almost in their present form, have been submitted for publication.

Verónica A. Becher

June 1999

Contents

Chapter 1

Introduction

The problem of change in corpora of information is indeed interesting. Legislation is under constant modification, new discoveries reshape scientific theories and robots have to update their representation of the world each time a sensor gains new data. The theory of theory change offers a model for these processes under certain idealizations.

A formal language and a notion of logical consequence relation are assumed. New information is expressed as sentences in the logical language. Sets of sentences closed under logical consequence, i.e. theories, are changed by *functions*. These take a theory and a sentence and return an updated theory. There is a leading principle for change functions: consistency. The result of a change by a consistent sentence should always be a consistent theory.

In 1985 Alchourrón, Gärdenfors and Makinson (henceforth "AGM") published the article that became the classical reference in the literature on theory change [Alchourrón *et al.*, 1985]. They conceived change functions that, under the maxim of consistency, preserve *as much as possible* of the original theory while accounting for the new information; theories should not be changed beyond necessity. Subset inclusion among theories alone is not enough as a criterion of minimal information loss because, in general, infinitely many theories are incomparable with each other with respect to set inclusion. Hence, it may be impossible to select a single one as the most preservative. As a result AGM functions must commit to a nondeterministic choice or else encode some other criteria for selection. The work of Alchourrón, Gärdenfors and Makinson created a whole new area of research, also referred to as *belief revision* (see for example, [Gärdenfors, 1988],[Gärdenfors, 1992]).

At least in two respects the AGM theory is underdefined. One is *the problem of iterated change.* AGM functions model single changes, they take one given theory to an updated theory, they perform one single step. But there will be yet another change after the one just considered that will induce yet another theory. That is, we will have to update the already updated theory. Although the AGM formalism does not forbid the iteration of change functions, it omits any specification of how it should be performed or what the properties of successive change are.

The other is *the problem of change in multiple theories.* If two theories are not independent of one another we may expect the respective change operations not to be independent either. For example, if one theory is included in another, we may expect that the change of the first be included in the change of the other. Some coherence properties should linking the change operation over different theories. This is the central topic of this thesis. This problem has hardly been addressed by the AGM theory and it has not been the object of much investigation in the theory change community either.

In this thesis we will propose change functions that are defined for every theory and every formula. As we will stay within the AGM framework, we will refer to them as *binary AGM functions.* We use the term binary to mean that they are functions of arity two; they take a theory and a formula and return a theory. In particular we will provide two specific formulations, two significant subclasses of binary AGM functions. We will establish a relation between the problem of iterated change and the problem of change in multiple theories and we will propose binary AGM functions as a definitionally simple scheme of iterated change. Clearly, binary functions can account for successive change because a theory returned by one application of a binary function is yet a possible argument of the same function. If our simple solution for iterated change possesses enough virtues (for a class of problems at hand) then the maxim of parsimony in science will have been achieved. Otherwise, if it oversimplifies the problem, it will be justified to commit to a more complex solution.

1.1 The Theory of Theory Change in Computer Science

Computer programs are finite sequences of symbols that are expected to perform some task. Thus, programs can be taken to be a symbolic representation of their output. We may face two different reasons for changing a program. One is when the program's output differs from what we expected it to do. We usually say that the program is incorrect with respect to its specifications, or that it "has bugs". Correcting the program, also referred as "debugging", leads to new versions of the program until one (hopefully) reaches a final version that performs the desired task.

A different reason to modify a program is when we are given a new specification of what our program should do. Even though our program was sound with respect to some original specification, it should now be changed to match an updated specification. It is not that our program was incorrect, but there is something different that should be accounted for.

The two examples above illustrate two different forms of changing representations. The theory of theory change offers a model of the dynamics of representations. AGM change functions, specially AGM revisions, have been considered suitable for correcting representations, but not for modeling changes produced by evolving specifications. A suitable operation for these kinds of changes has been proposed by Katsuno and Mendelzon [1992], and the two operations have been taken as representative of fundamentally different forms of theory change.

The theory of theory change was rapidly included in Artificial Intelligence (AI). According to the the declarative or logical school, as portrayed by [Hayes, 1985] or [Moore, 1982], when solving problems in AI we start from a representation of a problem. But such a representation may only be applicable if we can understand and model how to update it in light of new information. The state of a program is expected to be in constant change, reflecting the diverse inputs from the world. Theories of theory change are relevant to AI addressing this issue.

But why a representation must be a set of sentences in some logic? As explained by Boutilier [1992a], of course any formal system will do when it comes to characterizing in a principled manner the reasoning performed by a program, and logic should be accorded no special status in this regard. If a set of differential equations will accurately model the behaviour of a program, why bother with logical accounts? While prediction of behaviour might be accurate within any formal system, it is the model-theoretic semantics of logics that give logical representations their advantage in understanding behaviour. Clearly formal semantics provides no real meaning to sentences (see [Putnam, 1970]), it is merely the mapping of one mathematical structure (the logical language) into another (an interpretation of the language). These so-called models may be any structure whatsoever. What different but equivalent representations are useful for is to grasp a problem from different perspectives. Logics are equipped with formal semantics that justify the notion of consequence in the logic. However, we require no actual commitment by a system to give an explicit representation in terms of logical sentences and to reason with a general purpose theorem prover, only that such a system be able to be understood in such terms.

1.2 Thesis Overview

Throughout we will assume some familiarity with classical propositional logic and with the AGM theory. In Chapter 2 we will introduce notational conventions and review the background concepts that will be needed. We will briefly present the definitions and results of theory change that we will be concerned with.

In Chapter 3 we will formally present the two main problems discussed in this thesis. On the one hand, the problem of change in multiple theories, which was originally considered by Alchourrón and Makinson in their article on Safe Contractions [Alchourrón and Makinson, 1985]. Since then, this problem has not been the object of much attention in the literature and the existing examples of functions that provide coherent change in multiple sets were motivated by unrelated concerns. The problem arises because AGM defined functions that are relative to a specific given theory and may be inapplicable to another. On the other hand, the problem of iterated change. Also in the same article on Safe Contractions appears the very first reference to successive application of change operators. In contrast to the problem of change in multiple theories, this has become a very relevant theme and there is already an important amount of literature about it. Instead of surveying the different approaches to iterated change we will isolate a set of significant properties arising from different proposals. Our contribution in this chapter will be to show that although the AGM model has been criticized for not addressing the problem of iterated change, it is in fact compatible with iteration. In particular we will show how the subclass of AGM functions that solves the problem of change in multiple sets, provides a simple scheme of iterated change. We will dub the functions in this class *binary AGM functions.*

A main concern among researchers studying iterated change is whether there is a unique general model, a single set of properties in the same spirit AGM postulated functions for single changes. Fourteen years after the inception of the AGM theory we find several alternative formalizations differing in their virtues and defects, but remains unknown whether there could exist such a uniform set of properties of iterated change; perhaps there is no unique regularity to be exposed.

We will devote Chapter 5 to present a result about a distinctive binary function authored by Katsuno and Mendelzon [1992], the update operation. In contrast to the AGM tradition, Katsuno and Mendelzon have formalized their operation as a connective in a finite language — namely, a propositional language over a finite set of propositional variables—. In this chapter we will reformulate the update operation as a binary function, taking a theory and a formula to an updated theory. Then we will exhibit an unexpected result: Katsuno and Mendelzon's postulates are incomplete to characterize the update function for infinite propositional languages. We will then provide the appropriate set of postulates, strengthening the original ones, and prove the corresponding representation theorem for possibly infinite propositional languages. This result extends and clarifies previous results in the area.

We will define two families of binary AGM functions. Chapter 4 considers AGM functions that are almost constant (on their first argument, the second argument held fixed); we will name them *iterable AGM functions.* We will show that despite their definitional simplicity they satisfy a number of significant properties of iterated change.

Then, in Chapter 6 we will present binary functions that are almost monotonous (on their first argument, the second argument held fixed); we will name them *analytic AGM functions*. They too satisfy many of properties of iterated change but they are definitionally more complex than iterable AGM functions. For both, iterable and analytic functions, we will give formulations for possibly infinite languages, provide alternative representations and prove representation theorems.

Analytic AGM functions have two main interests. As AGM functions for changing multiple theories they possess a significant property. The analytic change operation is decomposable in the sense that it can can be calculated by means of simpler operations. The other main interest of analytic AGM functions is that they provide a formal link between the AGM revision operation and Katsuno and Mendelzon's update, when the two have been traditionally taken as incomparable frameworks.

We will devote Chapter 7 to provide a unification result about two logical calculus for the AGM theory: DFT [Alchourrón, 1995] and CO [Boutilier, 1992a]. By appealing to the notion of consequence they both allow to calculate changes in different theories. Although the two are modal conditional logics, they differ. Boutilier's semantics is relational while Alchourrón's is not. They also differ in the definition of their conditional connective. We will prove that, under restricting conditions, the two logics are indeed equivalent. In both logics the nested occurrences of the conditional connective suggest a function of iterated change. Unfortunately, DFT and CO are of no help to the problem of iterated change since such a function is truly trivial.

Finally, in Chapter 8 we will summarize the contributions of this thesis and examine avenues for further research.

Chapter 2

The AGM Theory of Theory Change

Throughout this thesis we will assume knowledge of the AGM theory. In this chapter we will briefly present the definitions and results that will be needed in subsequent chapters, making emphasis on the alternative presentations of the AGM theory. We shall start introducing notational conventions and basic definitions.

2.1 Preliminaries

If X and Y are sets, a relation R between X and Y is a set of ordered pairs, $R = \{(x, y) | x \in X \text{ and } y \in Y\}$, a subset of the Cartesian product of $X \times Y$. If $(x, y) \in R$ we shall write xRy.

A function from X to Y is a relation f such that the domain of f is X and for each $x \in X$ there is a unique element y in Y with $(x, y) \in f$. For each $x \in X$ the unique $y \in Y$ is denoted by $f(x)$. From now on we shall write $f(x) = y$ instead of $(x, y) \in f$. The element y is called the value that the function assumes at the argument x. The words map or mapping and operator are sometimes used as synonymous for function. The range of f consists of those elements y of Y for which there exists an x in X such that $f(x) = y$. If the range of f is equal to Y, then f is surjective. If f maps different elements of the domain to different elements in the range, then f is injective. If f is surjective and injective, then it is bijective, establishing a one to one correspondence between X and Y. The symbol $f : X \to Y$ is used as an abbreviation for "f is a function from X to Y." Given a function $f : X \times Y \to Z$, we will refer to the function $f_{x_0} : Y \to Z$ defined $f_{x_0}(y) = f(x_0, y)$ as the projection of f for a fixed value x_0 of the first argument. Similarly for the second argument, $f_{y_0} : X \to Z$ defined $f_{y_0}(x) = f(x, y_0)$, as the projection of f for a fixed $y_0 \in Y$.

A function is *unary* when it is a function on a single argument. A function is called *binary* (n-ary), or of two (n) arguments, if it is defined on a set of ordered pairs (n-tuples); for example, the sum on the natural numbers is binary.

We will refer to the following properties of binary relations. Let X be a set, and R be a binary relation over elements of X.

R is *irreflexive* in X if and only if for all $x \in X$, not xRx.

R is *reflexive* in X if and only if for all $x \in X, xRx$.

R is *symmetric* in X if and only if for all $x, y \in X$ if xRy, then yRx.

R is *antisymmetric* in X if and only if for all $x, y \in X$, xRy and yRx only if $x = y$.

R is *transitive* in X if and only if for all $x, y, z \in X$, if xRy and yRz then xRz.

R is *connected* in X if and only if for all $x, y \in X$ if $x \neq y$, then xRy or yRx.

R is *totally connected* in X if and only if for all $x, y \in X : xRy$ or yRx.

Notice that total connectedness implies reflexivity.

A relation R over X is *virtually connected* over $Y \subseteq X$ if and only if for every $x, y, z \in Y$ if xRy then either xRz or zRy. Equivalently, $R \subseteq X \times X$ is virtually connected over $Y \subseteq X$ iff its complement $\bar{R} = X \times X - R$ is transitive over Y.

R is a *preorder* on X if and only if R is reflexive and transitive.

R is a *partial order* on X if and only if R is reflexive, transitive and antisymmetric.

R is a *total order* on X if an only if R is antisymmetric, transitive and totally connected.

A relation R is *well founded* on X if every non empty subset of X has a non empty subset of R-minimal elements; equivalently, if R is free of infinite descending chains.

A relation R on X is *acyclic* if for any set of elements $x_1, \ldots, x_n \in X$, it is not the case that $x_1 R x_2 R \ldots x_n R x_1$. Let's notice that for $n = 1$, acyclicity implies irreflexivity. For $n = 2$ acyclicity implies asymmetricity.

To denote arbitrary relations that are orders we will use the symbols $\prec$, $\preceq$, $<$ and $\leq$, sometimes with subscripts. We will write $\mathbb{N}$ for the set of natural numbers, $\mathcal{O}$ for the set of ordinals and $\mathbb{R}$ for the reals.

We assume familiarity with basic notions of propositional logic. We consider a classical propositional language L and denote with P the set of all its propositional letters. If P is finite we will call L a finite propositional language. The symbols $\wedge, \vee, \neg, \supset, \equiv$ will denote the usual truth functional connectives. Indistinguishably, we will use the terms formula and sentence to refer to an element of L. As we only deal with propositional languages the two terms are indeed equivalent. Capital letters A, B, C will be used to denote arbitrary formulae of L. We consider Cn a Tarskian consequence operation, a function that takes each subset of L to another subset of L such that:

(inclusion) $X \subseteq \mathrm{Cn}(X)$.

(monotony) If $X \subseteq Y$ then $\mathrm{Cn}(X) \subseteq \mathrm{Cn}(Y)$.

(idempotency) $\mathrm{Cn}(X) \subseteq \mathrm{Cn}(\mathrm{Cn}(X))$.

In addition, following [Alchourrón *et al.*, 1985] we assume Cn on L satisfies:

(supra classicality) If A can be derived from X by classical truth functional logic, then $A \in \mathrm{Cn}(X)$.

(compactness) If $A \in \mathrm{Cn}(X)$, then $A \in \mathrm{Cn}(Y)$ for some finite subset $Y \subseteq X$.

(introduction) If $C \in \mathrm{Cn}(X \cup \{A\})$ and $C \in \mathrm{Cn}(X \cup \{B\})$ then $C \in \mathrm{Cn}(X \cup \{A \vee B\})$. (introduction of disjunction into the premisses).

Under these assumptions the consequence operation Cn also satisfies the deduction theorem, that $B \in \mathrm{Cn}(X \cup \{A\})$ if and only if $(A \supset B) \in \mathrm{Cn}(X)$, A theory is a subset of L closed under Cn. Capital letters K, K', H are used for theories of L, and we denote by $\mathbb{K}$ the set of all theories of L. While L is the largest theory, $\mathrm{Cn}(\emptyset)$ is the smallest. A subset X of L is consistent (modulo Cn) iff for no formula A do we have $(A \wedge \neg A) \in \mathrm{Cn}(X)$. A theory is complete if it sanctions a truth value for each propositional letter.

We take W as the set of all maximal consistent subsets of L, that is, the set of all complete consistent extensions of L. The valuation function $[\ ] : L \rightarrow \mathcal{P}(W)$ is defined as usual, for any propositional letter p, $w \in [p]$ iff $p \in w$. Given $A \in L$ we denote by $[A]$ the proposition for A, or the set of A-worlds, the set of elements of W satisfying A. For the purposes of this work we consider the terms maximal consistent subset of L, valuation on L and possible world, interchangeable. This, of course, amounts to working with models that are injective with respect to the interpretation function (no two distinct worlds satisfy exactly the same formulae) and full (every consistent set of formulae is satisfied by some world). If K is a theory, $[K]$ denotes the set of possible worlds including K. Given U a set of possible worlds, $\mathrm{Th}(U)$ returns the associated theory.

We will say that a subset X of W is *L-nameable* whenever there exists a formula A in L such that $X = [A]$. When working with relations on W, we will often refer to a property that Lewis [1973] called the *limit assumption.* A preorder relation R on W satisfies the limit assumption if and only if for any satisfiable formula A in L there exists a set of R-minimal A-words. This requirement is in general weaker than the well foundedness condition. The limit assumption just requires that L-nameable non empty subsets of W have set of minimal elements, as opposed to requiring so for every subset of W.

2.2 AGM Functions

A comprehensive introduction to the AGM theory can be obtained in [Gärdenfors, 1988; Gärdenfors, 1992; Hansson, 1998].

Three are the operations advocated by the AGM model [Alchourrón *et al.*, 1985]: expansions, revisions and contractions. The first two deal with "accommodating" a new formula into the current theory, while the third is a "removing" operation. Expansion is the simplest form of theory change. It is a simple addition function, where a new formula A, hopefully consistent with a given theory K, is set theoretically added to K and this expanded set is closed under logical consequence. The function $+ : \mathbb{K} \times L \rightarrow \mathbb{K}$ is defined as $K + A = \mathrm{Cn}(K \cup \{A\})$. The expansion function can also be characterized by the following postulates [Gärdenfors, 1982].

(K+1) $K + A$ is a theory. (closure)

(K+2) $A \in K + A$. (success)

(K+3) $K \in K + A$. (inclusion)

(K+4) If $A \in K$ then $K + A = K$. (vacuity)

(K+5) If $K \subseteq H$, $K + A \subseteq H + A$. (monotony)

(K+6) $K + A$ is the $\subseteq$-smallest theory that satisfies closure, success, and inclusion. (minimality)

The AGM contraction and revision operations have a more subtle definition. The contraction function $-$ takes a theory K and a formula A and returns the contracted theory, notated as $K - A$. Contractions are changes in a theory that involve giving up some formulae without incorporating new ones. When retracting a formula A from K, there may be other formulae in K that entail A (or other formulae that jointly entail A without separately doing so). In order to keep $K - A$ closed under logical consequence, it is necessary to give up A and other formulae as well. The problem is to determine which formulae should be given up and which should be retained. In contrast to expansion, the explicit construction of AGM contraction functions is not so direct. AGM developed postulates that fully characterize the contraction functions. The first six postulates, (K-1)-(K-6), are called the basic postulates for contraction and they characterize *partial meet contraction functions.* Postulates (K-7) and (K-8) are called supplementary, and they impose additional conditions, which give rise to *transitively relational contraction functions.* These functions will be our focus of attention. The names of partial meet functions originated in the method for constructing the functions, that we shall review in the next section.

(K-1) $K - A$ is a theory. (closure)

(K-2) $K - A \subseteq K$. (inclusion)

(K-3) If $A \notin K$, then $K - A = K$. (vacuity)

(K-4) If not $\mathrm{Cn}(A) = \mathrm{Cn}(\emptyset)$ then $A \notin K - A$. (success)

(K-5) If $A \in K$, then $K \subseteq (K - A) + A$. (recovery)

(K-6) If $\mathrm{Cn}(A) = \mathrm{Cn}(B)$ then $K - A = K - B$. (preservation)

(K-7) $(K - A) \cap (K - B) \subseteq K - (A \wedge B)$. (conjunctive overlap)

(K-8) If $A \notin K - (A \wedge B)$, then $K - (A \wedge B) \subseteq K - A$. (conjunctive inclusion)

The conjunction of postulates (K*7) and (K*8) is equivalent to the Ventilation property reported in [Alchourrón *et al.*, 1985], which provides a factoring on the contraction by a conjunction from a theory.

(Ventilation) For all A and B, $K-(A\wedge B) = K-A$, or $K-(A\wedge B) = K-B$ or $K - (A \wedge B) = K - A \cap K - B$.

The AGM revision function $*$ takes a theory K and a formula A to a revised theory $K * A$. The problem here is that the formula A should be added under the requirement that the resulting theory be consistent (whenever A is); hence, A can not just be set theoretically added to K. Revisions are constrained by the following eight postulates [Alchourrón *et al.*, 1985].

(K*1) $K * A$ is a theory. (closure)

(K*2) $A \in K * A$. (success)

(K*3) $K * A \subseteq K + A$. (inclusion)

(K*4) If $\neg A \notin K$ then $K + A \subseteq K * A$. (vacuity)

(K*5) $K * A = \mathrm{Cn}(\perp)$ only if $\mathrm{Cn}(\neg A) = \mathrm{Cn}(\emptyset)$. (consistency)

(K*6) If $\mathrm{Cn}(A) = \mathrm{Cn}(B)$ then $K * A = K * B$. (preservation)

(K*7) $K * (A \wedge B) \subseteq (K * A) + B$. (superexpansion)

(K*8) If $\neg B \notin K * A$ then $(K * A) + B \subseteq K * (A \wedge B)$. (subexpansion)

As for contractions, the first six are called the basic postulates for revision, and they characterize *partial meet revision functions.* Postulates (K*7) and (K*8) are supplementary and they give rise to *transitively relational partial meet revision functions.*

A nice feature about revisions and contractions is that they are inter definable. By the *Levi identity* revisions can be defined in terms of contractions and expansions. This identity defines revisions as first pruning away all potential inconsistencies, and then adding the new formula.

(Levi-id) $K * A = (K - \neg A) + A$.

The counterpart of the Levi identity is the *Harper identity,* which provides a definition of contractions in terms of revisions. The formulae in $K - A$ is captured as what K and $K * \neg A$ have in common.

(Harper-id) $K - A = K \cap (K * \neg A)$.

It is not hard to verify that the two identities commute. Given the interdefinability of revisions and contractions throughout this thesis we will present change functions in either the contraction or revision version, indistinctly.

A crucial remark about the AGM postulates for contraction and revision is that they indicate nothing about the behaviour of the functions when applied to different theories $K \in I\!K$. Although change functions are supposedly defined as binary functions taking two arguments, a theory K and a formula A, they are in fact a *family of independent unary functions*:

$$\{*^K : L \rightarrow I\!K : K \in I\!K \text{ and } *^K \text{ satisfies K*1-K*8}\}.$$

The AGM postulates never refer to revision of different theories, hence the functions $*^K$ can be totally independent. There are no properties shaping the joint behaviour of the different unary functions. To discover what are these properties and how they interact is a fundamental issue, and we will return to it in subsequent chapters. For the moment it should be clear that postulates (K*1)-(K*8) aim to characterize only the single unary functions $*^K$ at a time, or equivalently, consider $*$ for a fixed, theory K (similarly so for contractions). That is, the postulates constrain the behaviour of the change function with respect to all kinds of propositional input but do not deal with varying theories (see [Rott, 1999] and [Areces and Becher, 1999]).

The situation is quite different for expansions. The function $+$ is indeed a binary function $+ : \mathbb{K} \times L \to \mathbb{K}$, which provides one and the same definition for every theory $K \in \mathbb{K}$. A strong coherence property links the expansion of different theories because set theoretical addition is monotone.

However, in general $*$ and $-$ are not monotone. If one theory is included in another, the revision of the first is not necessarily included in the revision of the second:

(Monotony *): If $H \subseteq K$ then $H * A \subseteq K * A$.

Observation 2.1 (follows from [Alchourrón *et al.*, 1985]) If $*$ is a revision operation satisfying postulates (K*1),(K*4) and (K*5), in a language admitting at least two mutually independent formulae A, B (neither $A \in \mathrm{Cn}(B)$ nor $B \in \mathrm{Cn}(A)$), then monotony fails for $*$.

PROOF. Let $K = \mathrm{Cn}(A, B)$, $H_1 = \mathrm{Cn}(A), H_2 = \mathrm{Cn}(B)$. Assume monotony.

As $H_i \subseteq K$ for $i \in \{1, 2\}$, by monotony, $H_1 * \neg(A \wedge B) \subseteq K * \neg(A \wedge B)$ and $H_2 * \neg(A \wedge B) \subseteq K * \neg(A \wedge B)$.

By independence, $H_1 = \mathrm{Cn}(A)$ is consistent with $\neg(A \wedge B)$, so $H_1 * \neg(A \wedge B) = \mathrm{Cn}(H_1 \cup \{\neg(A \wedge B)\}) = \mathrm{Cn}(A \wedge \neg B)$.

Likewise, $H_2 * \neg(A \wedge B) = \mathrm{Cn}(H_2 \cup \{\neg(A \wedge B)\}) = \mathrm{Cn}(B \wedge \neg A)$. Hence, both $(A \wedge \neg B)$ and $(B \wedge \neg A)$ are included in $K * \neg(A \wedge B)$.

Therefore, $K * \neg(A \wedge B)$ is inconsistent. By postulate (K*5), $\neg(A \wedge B)$ is then inconsistent, contradicting the independence of A and B. QED

2.3 Constructions of AGM functions

In the words of Alchourrón and Makinson [1982], the postulates characterize the change operations by formulating conditions of a more or less inclusional or equational nature. They allow for clear intuitions about the processes under study and the web of interrelations between them. But another approach to defining the functions is to seek for explicit constructions. These provide some kind of foundation for justifying the intuitions. Originally the work on contraction functions and their associated revision functions in terms of explicit constructions was given by Alchourrón and Makinson in [1982]. The representation theorem linking the explicit functions with the postulates was given in the celebrated joint paper by the three authors, Alchourrón, Gärdenfors and Makinson [1985].

2.3.1 Partial Meet Functions

Let K be a theory and A a language formula. The process of eliminating A from the theory is not uniquely defined unless additional specifications are given. In general there are many subsets of a set that do not imply a given formula, and indeed many maximal such subsets.

Alchourrón and Makinson [1982] base the construction of a contraction function for theory K and sentence A on the set of maximal subsets of K that fail to imply A. They define $K \perp A$ as the set of all these maximal subsets.

Definition 2.2

$$K \perp A = \{K' \subseteq K | A \notin \mathrm{Cn}(K') \text{ and } K' \text{ is } \subseteq\text{-maximal with this property}\}.$$

By the compactness of Cn it follows that $K \perp A$ is not empty unless $\mathrm{Cn}(A) = \mathrm{Cn}(\emptyset)$; in addition, the elements of $K \perp A$ are theories.

Alchourrón and Makinson give two natural ways to to define contraction functions : by intersection and by choice. The *full meet contraction* is defined by putting $K - A = \bigcap(K \perp A)$, when $K \perp A$ is non-empty, and to be K itself otherwise.

Definition 2.3 (Full meet contraction)

$$K - A = \begin{cases} \bigcap(K \perp A), & \text{if } K \perp A \neq \emptyset. \\ K, & \text{otherwise .} \end{cases}$$

In contrast, they define *maxichoice contraction function* by putting $K - A$ equal to a single element in $K \perp A$, whenever $K \perp A$ is non empty, and $K - A = K$, otherwise. To come up with the single element of $K \perp A$ they require a choice function that makes the selection (actually, in [Alchourrón and Makinson, 1982] this function is referred as a *choice contraction* that they rename as maxichoice in the AGM joint paper [Alchourrón *et al.*, 1985]). As they observed, maxichoice functions have some rather disconcerting properties. In particular the maxichoice revision has the property that for every theory K, whether complete or not, the maxichoice revision of K by A will be complete whenever A is a proposition inconsistent with K. So in general the result of a maxichoice revision is a set that is too large. In contrast, the result of a full meet function is in general too small. In particular when K is a theory with $\neg A \in K$ the full meet revision of K by A yields $\mathrm{Cn}(A)$, just the consequences of A.

Observation 2.4 If * is a full meet revision then for every $A \in L$,

$$K * A = \begin{cases} \mathrm{Cn}(A), & \text{if } \neg A \in K. \\ \mathrm{Cn}(K \cup \{A\}), & \text{otherwise.} \end{cases}$$

AGM explain that the full meet operation is very useful as a point of reference, as it serves as a natural lower bound of any reasonable change function. As a result in [Alchourrón *et al.*, 1985] they propose *partial meet functions*, which yield the intersection of some nonempty family of maximal subsets of the theory that fail to imply the formula being eliminated.

A partial meet function is based on a selection function s^K which returns a nonempty subset of a given nonempty set $K \perp A$. Let K be a theory, we note as $s^K : L \rightarrow \mathcal{P}(\mathcal{P}(K)) \setminus \{\emptyset\}$, a selection function for $K \perp A$, for $A \in L$. We furthermore require that $s^K(A) = \{K\}$ whenever $K \perp A = \emptyset$. The AGM partial meet contraction function $-$ is then defined, for a theory K, as follows.

Definition 2.5 (Partial meet contraction)

$$K - A = \bigcap s^K(A), \text{ where } s^K \text{ is a selection function for } K.$$

Under this definition the contraction function $-$ is formally characterized by the basic AGM postulates (K-1) to (K-6) ([Alchourrón *et al.*, 1985], Observation 2.5.) For $-$ to be characterized by the extended set of postulates, (K-1) to (K-8), it suffices that s^K be *transitively relational*, i.e. for each $A \in L$ the selection function returns the smallest elements according to some transitive relation defined over $K \perp A$ ([Alchourrón *et al.*, 1985], Corollary 4.5.) Explicit constructions of (transitively relational) partial meet revisions are definable via the Levi identity, so the representation results apply for revisions as well.

The AGM theory enjoys three other presentations over quite different formal structures. We will briefly present them here to be revisited in the chapters to follow. We will first concentrate on Alchourrón and Makinson's [1985] *contraction* function, where they start from a hierarchical ordering over the formulae in the theory under change. The connection between safe contractions and transitively partial meet contractions was studied by [Alchourrón and Makinson, 1986] in the finite case, and extended by [Rott, 1992a] for the general case. In addition, safe contraction functions were later generalized by [Hansson, 1994] under the name of incision functions.

Then we will focus on epistemic entrenchment orderings, originally defined by Gärdenfors [1984]. The representation theorem linking the change functions based on epistemic entrenchment relations and partial meet functions was proved in [Makinson and Gärdenfors, 1988].

Next we will concentrate on the systems of spheres proposed by [Grove, 1988]. They provide a kind of possible worlds semantics for AGM functions. This representation result allowed later for the connection established by [Boutilier, 1992a] between AGM functions and modal conditional logics. Grove's formalization has been of great insight for us too, most of our definitions were firstly considered in systems of spheres.

2.3.2 Safe Contraction Functions

Alchourrón and Makinson [1985] construct a contraction function based on a hierarchical ordering in the language. They based the idea on their previous work on hierarchies of regulations and their logic [Alchourrón and Makinson, 1981].

Let K be a theory, $<_{s,f}$ a non-circular relation over K and A a formula in L we wish to eliminate from K. An element is *safe* with respect to A (modulo $<_{s,f}$ and given some background Cn) iff it is not a minimal element under $<_{s,f}$ of any minimal subset (under set inclusion) $H \subseteq K$ such that $A \in Cn(H)$. They define the *safe contraction* of K by A as the set of safe elements of K with respect to A. Let's study some details. .

A binary relation $<_{s,f}$ over a set K is a *hierarchy* if it is acyclic: for any set of elements $A_1, \ldots, A_n \in K, n \geq 1$, it is not the case that $A_1 <_{s,f} A_2 <_{s,f} \ldots A_n <_{s,f} A_1$.

A relation $<_{s,f}$ over K *continues up* Cn if for every $A_1, A_2, A_3 \in K$, if $A_1 <_{s,f}$

A_2 and $A_3 \in \mathrm{Cn}(A_2)$ then $A_1 <_{s,f} A_3$.

A relation $<_{s,f}$ over K *continues down* Cn if for every $A_1, A_2, A_3 \in K$, if $A_2 \in \mathrm{Cn}(A_1)$ and $A_2 <_{s,f} A_3$ then $A_1 <_{s,f} A_3$.

A relation $<_{s,f}$ over K is *virtually connected* if for every $A_1, A_2, A_3 \in K$ if $A_1 <_{s,f} A_2$ then either $A_1 <_{s,f} A_3$ or $A_3 <_{s,f} A_2$.

Let $<_{s,f}$ be a virtually connected hierarchy over a theory K that continues up and down Cn, and let A be a sentence in L. The safe contraction function $-_{s,f}$ is defined as:

Definition 2.6 (Safe contraction)

$$K -_{s,f} A = \mathrm{Cn}(\{B| \quad \forall K' \subseteq K, \text{ s.t. } A \in \mathrm{Cn}(K') \text{ and } K' \text{ is } \subseteq\text{-minimal with this property, } B \notin K' \text{ or there is } C \in K' \text{ s.t. } C <_{s,f} B\}).$$

The elements of $K -_{s,f} A$ are called the safe elements of K with respect to A since they can not be "blamed" for implying A. An element is safe for A if it does not belong to any of the $\subseteq$-minimal subsets of K that imply A, or else it is not $<_{s,f}$-minimal in the hierarchy in such subsets.

Alchourrón and Makinson [1985] show that every safe contraction over a theory K is a partial meet contraction function over K. They also prove the converse result for finite theories (in the sense that the consequence operation Cn partitions the elements of K into a finite number of equivalence classes). The general case (finite and infinite theories) was proved by [Rott, 1992a]. The following representation theorem links safe contractions contractions and partial meet functions.

Observation 2.7 ([Alchourrón and Makinson, 1986; Rott, 1992a])
Every contraction function $-$ over K satisfying the (K-1)-(K-8) can be represented as a safe contraction function $-_{s,f}$ generated by a hierarchy $<_{s,f}$ that is virtually connected and continues up and down Cn.

2.3.3 Epistemic Entrenchments

An *epistemic entrenchment* for a theory K is a total relation among the formulae in the language reflecting their degree of relevance in K and their usefulness when performing inference. The following five conditions are required for an epistemic entrenchment relation $\leq_{ee}$ for a theory K [Gärdenfors, 1984; Makinson and Gärdenfors, 1988]:

(EE1) If $A \leq_{ee} B$ and $B \leq_{ee} D$ then $A \leq_{ee} D$.

(EE2) If $B \in \text{Cn}(A)$ then $A \leq_{ee} B$.

(EE3) $A \leq_{ee} (A \wedge B)$ or $B \leq_{ee} (A \wedge B)$.

(EE4) If theory K is consistent then $A \notin K$ iff $A \leq_{ee} B$ for every B.

(EE5) If $B \leq_{ee} A$ for every B then $A \in \text{Cn}(\emptyset)$.

The associated interpretation is that the epistemic entrenchment of a sentence is tied to its overall informational value within the theory. For example, lawlike sentences generally have greater epistemic entrenchment than accidental generalizations. When forming contractions, the formulae that are retracted are those with the lowest epistemic entrenchment. Tautologies are the most entrenched, hence they are never given up.

For any given relation $\leq_{ee}$ for a consistent theory K, the formulae in K are ranked in $\leq_{ee}$, while all the formulas outside K have the $\leq_{ee}$-minimal epistemic value. That is, by (EE4) for a consistent theory K, all the formulas outside K are zeroed in $\leq_{ee}$. However, (EE4) is vacuous for the contradictory theory L. (EE1)-(EE3) imply connectivity, namely, either $A \leq_{ee} B$ or $B \leq_{ee} A$ (the epistemic entrenchment ordering will cover *all* the sentences).

The AGM contraction function $-_{ee}$ based on an epistemic entrenchment relation $\leq_{ee}$ for K, is defined as follows.

Definition 2.8 (Epistemic entrenchment contraction) For every formula A in L,

$$K -_{ee} A = \{B \in K | A \in \text{Cn}(\emptyset) \text{ or } A <_{ee} (A \vee B)\},$$

where $<_{ee}$ is the strict relation obtained from $\leq_{ee}$.

The representation result shows that a revision function can be constructed by means of an epistemic entrenchment ordering on the language.

Observation 2.9 ([Makinson and Gärdenfors, 1988]) A contraction function $-$ for K satisfies (K-1)-(K-8) iff there exists an epistemic entrenchment relation for K satisfying (EE1)-(EE5) such that for all $A \in L$, $K - A = K -_{e} A$.

2.3.4 Systems of Spheres

Among the alternative presentations of the AGM theory, Grove's [1988] provides a possible worlds semantics systems of spheres. A system of spheres S^K centered on a theory K is a subset of $\mathcal{P}(W)$ containing W, totally ordered under set inclusion, such that $[K]$ is the $\subseteq$-minimal element of S^K. A system S^K should validate the limit assumption, in the sense that for every satisfiable formula A in the language there exists a $\subseteq$-minimal sphere in S^K (written as $c^K(A)$) with non-empty intersection with $[A]$.

Definition 2.10 (System of spheres) A system of spheres S^K centered on theory K is a set of sets of possible worlds that verifies the properties:

(S1) If $U, V \in S$ then $U \subseteq V$ or $V \subseteq U$. (totally ordered.)

(S2) For every $U \in S$, $[K] \subseteq U$. (minimum.)

(S3) $W \in S$. (maximum.)

(S4) For every sentence A such that there is a sphere U in S^K with $[A] \cap U \neq \emptyset$, there is a $\subseteq$-minimal sphere V in S such that $[A] \cap V \neq \emptyset$. (limit assumption.)

For any sentence A, if $[A]$ has a non-empty intersection with some sphere in S^K then by (S4) there exists a minimal such sphere in S^K, say $c^K(A)$. But, if $[A]$ has an empty intersection with all spheres, then it must be the empty set (since (S3) assures W is in S^K), in this case c^K is put to be just W. Given a system of spheres S^K and a formula A, c^K is defined as:

$$c^K(A) = \begin{cases} W & \text{if } [A] = \emptyset \\ \text{the } \subseteq\text{-minimal sphere } S' \text{ in } S^K \text{ s.t. } S' \cap [A] \neq \emptyset & \text{otherwise.} \end{cases}$$

A system S^K determines a contraction function $-_{s}$ for K in the sense that for every formula $A \in L$ and every $w \in W$, $w \in [K -_{s} A]$ iff $w \in (c^K(A) \cap [A]) \cup [K]$.

Definition 2.11 (Sphere contraction) Let S^K be a system of spheres centered on K. For every formula A in L,

$$K -_s A = \mathrm{Th}((c^K(A) \cap [A]) \cup [K]).$$

Grove proves the following representation result.

Observation 2.12 ([Grove, 1988], Theorems 1,2) $-$ is a revision function for K satisfying (K-1)-(K-8) iff there exists a system of spheres S^K centered on K such that for all formulas $A \in L$, $K - A = K -_s A$.

The same approach can be used to model revision functions. If we define $K * A$ as the theory of $(c^K(A) \cap [A])$, by means of the Levi identity we obtain the representation theorem for contraction.

Let's turn now to a subclass of AGM functions, the subclass generated by well founded systems of spheres. A system of spheres S^K is well founded if $\subset$ is a well founded relation on S^K, that is for *every subset* of $X \subseteq W$ there exists $\subset$-minimal sphere in S^K intersecting X. In contrast general systems of spheres establish the requirement only for *nameable* subsets of W - actually we require nameability by a single formula rather than a set of formulae -. Following [Peppas, 1993] we refer to revision functions definable over a well founded system of spheres as *well behaved revision functions.* All revision functions for theories over a finite propositional language are well behaved. But it is well known that well founded systems of spheres do not capture all AGM revision functions. This is perspicuously proved by Peppas in [1993] who exhibits a first order theory K and a revision function $*$ for K such that no well founded system of spheres represents $*$. Peppas characterizes well behaved revision functions the following postulate.

(K*WB) For every nonempty set X of consistent formulae of L there exists a formula $A \in X$ such that $\neg A \notin K * (A \vee B)$, for every $B \in X$.

Peppas proves the following.

Observation 2.13 ([Peppas, 1993], Theorem 5.4.3) Let $*$ be a revision function satisfying (K-1)-(K-8). Then $*$ is well behaved iff it satisfies (K*WB) for every theory K of L.

It is up to now unclear how strong is the restriction to well behaved revision functions.

It is possible to recast a system of spheres centered in $[K]$ as a total preorder $\preceq$ over W, having the elements of $[K]$ as minimal elements, and satisfying the limit assumption (every L-nameable subset of W must have some $\preceq$-minimal element). Without loss of generality then a system of spheres centered in $[K]$ can be seen as a function from W to any totally ordered set with smallest element. This set can be taken to be $\mathbb{R}^+$, be the set of positive real numbers including 0, but not necessarily so. We define $d_K : W \rightarrow \mathbb{R}^+$ that decorates with real numbers the nested spheres of a Grove system.

Observation 2.14 For every system of spheres S^K there is a function d_K on $\mathbb{R}^+$ such that

$d_K(v) < d_K(w)$ iff $(\exists S_1, S_2 \in S^K)(v \in S_1, w \in S_2$ and $S_1 \subset S_2)$, and

$d_K(v) = d_K(w)$ iff $(\forall S_i \in S^K)(w \in S_i \Leftrightarrow v \in S_i)$.

These functions provide a notion of distance from theories to worlds: If $d_K(w) < d_K(v)$ then w is closer than v or "more consistent" with the current theory K. And this measure can be naturally extended to functions over sets of worlds (propositions), by requiring the value assigned to a set X to be the smallest value assigned to the worlds in X. Special consideration is required if X is empty. Let now S^K be any system of spheres and d_K any real function corresponding to it as in Observation 2.14 above. We first extend d_K to any subset of W as follows. Define $d_K : \mathcal{P}(W) \rightarrow \mathbb{R}^+$ as:

$$d_K(X) = \begin{cases} \min\{d_K(w) : w \in X\} & ,\text{if } X \neq \emptyset. \\ 0 & ,\text{if } X = \emptyset. \end{cases}$$

In order to represent a system of spheres by a function d_K we should impose the limit assumption on d_K. For every nameable subset X , $d_K(X)$ must be defined. But if X is not nameable by a single formula then the set $\{d_K(w) : w \in X\}$ can be infinite, with infinite descending values where the min may be undefined.

The function d_K induces a revision function $*$ such that $K * A$ is the theory entailed by the set of A-worlds that are closest to K according to the function d_K. Then, if we take

$$K * A = \mathrm{Th}(\{w \in [A] : d_K(w) = d_K([A])\})$$

the revision operation so obtained coincides with the original $*$ operation whose semantic model was S^K.

Well founded systems of spheres are free of infinite descending chains of spheres. Consequently for these systems the function d_K can be defined over the ordinals as opposed to be defined over the reals. For instance Spohn's ordinal functions [Spohn, 1987] $k_K : W \rightarrow \mathcal{O}$ straightforwardly represent well founded systems of spheres that are centered on a consistent theory.

Chapter 3

Binary Functions for Theory Change

The AGM model has as points of departure a theory to be modified, a formula to be considered as new information, and change functions. The framework seems to be that of binary functions that when applied to a theory K and a formula A return a new theory K'. However, when we study the AGM model, we immediately realize that change functions are *relative to some given theory*. Then, a function is applicable to theory K, but in general, it is not applicable to another theory K'. The situation can be compared, for example, to the square root function on the set of Natural numbers. If we define $f : \mathbb{N} \rightarrow \mathbb{N}$ by $f(n) = \sqrt{n}$, the actual domain of f is just the set of perfect squares, because the value of f is only defined for natural numbers that are perfect squares, and it is undefined otherwise. So f is only partially defined over $\mathbb{N}$. Hence, when AGM functions are regarded as binary functions they are just partially defined on the set $\mathbb{K} \times L$. Once the first argument has been fixed to be a given theory K, the function is well defined for every language formula. But it may be undefined when the first argument is any other theory. Henceforth the behaviour of AGM functions is asymmetrical with the two arguments. For this reason we argue that AGM functions are not truly binary and they should be taken to be *unary functions* $-^K : L \rightarrow \mathbb{K}$. It is possible to consider a binary $- : \mathbb{K} \times L \rightarrow \mathbb{K}$ as the family of independent unary functions $-^K$, one per theory K, such that

$$- = \{-^K : L \to \mathbb{K} : K \in \mathbb{K} \text{ and } -^K \text{ satisfies } (K-1)-(K-8)\}$$

But this family can be arbitrary. The AGM postulates only aim to constrain the behaviour of the indexed unary functions separately without trying to correlate them with each other.

3.1 Change in Multiple Theories

The problem of *change in multiple theories* is simply the problem of an appropriate definition of unary change functions that are jointly coherent. Suppose we possess some complete set of unary change functions, one for each possible theory $K \in \mathbb{K}$. Since the AGM postulates provide no correlation between the different unary functions we should not expect that the change of one theory be significantly related to the change of another.

But not all AGM functions are alike in this respect. Expansions are substantially different to general revisions and contractions. The definition of $+ : \mathbb{K} \times L \to \mathbb{K}$ states that for every $K \in \mathbb{K}$ and for every $A \in L$, $K + A = \mathrm{Cn}(K \cup \{A\})$. Expansion is really a function of two arguments, identically defined for every theory and every formula just in terms of the consequence relation. Being based on Cn, + inherits monotony, which definitely counts as a coherence property over the change in different theories.

As it stands in the original partial meet construction presented in [Alchourrón *et al.*, 1985], AGM contraction is a unary function relative to a theory K, $-^K : L \to \mathbb{K}$ based on a selection function s^K depending on K.

$$\text{Partial Meet: } -^K(A) = \begin{cases} \bigcap s^K(K \perp A) & \text{if } K \perp A \neq \emptyset \\ K & \text{otherwise} \end{cases}$$

where the set $K \perp A$ contains the maximal subsets of K that do not imply A and the function $s^K : L \to \mathcal{P}(\mathcal{P}(K)) \setminus \{\emptyset\}$ selects a nonempty subset of $K \perp A$.

The limiting case in which the function s^K returns the whole set $K \perp A$ gives rise to the full meet contraction function. The selection function s^K relative to K disappears, yielding a contraction that depends solely on the explicit arguments K and A, i.e. if $-$ is a full meet, again we have a binary function $- : \mathbb{K} \times L \to \mathbb{K}$.

Moreover, the representation result states that the full meet contraction function is defined as $K - A = K \cap \mathrm{Cn}(\neg A)$. Like expansions, they depend on no underlying structure, relative order or selection function, and are applicable to every theory.

Alchourrón, Gärdenfors and Makinson have argued that full meet functions suffer from too much loss of information and have taken them as a demarcation of the limiting case. The question now is whether it is possible to provide binary AGM functions which are more interesting than full meet functions. As we will remark in Chapter 5, outside the AGM framework there is a vivid example of a binary function with a strong coherence property, Katsuno and Mendelzon's update function. In Chapters 4 and 6 we will present two different formulations of binary AGM functions, one is based on Alchourrón and Makinson's safe contraction [1985], the other inspired in the update function.

3.2 Iterated Change

The motivation to consider successive change is indisputable. A change operation takes a theory to a modified theory. But eventually there will be yet another change after the one just considered that will induce yet another theory. Hence we will have to update the already updated theory. This problem has been dubbed the problem of *iterated theory change.*

A pertinent criticism of the AGM formalism is its lack of definition with respect to iterated change (see [Halpern and Friedman, 1996] and Rott [1999; 1998]). The iteration of revisions, contractions and expansions separately is significant, and even more so the consideration of sequences of different kinds of change. Although the AGM formalism does not forbid the iteration of change functions, it omits any specification of how it should be performed or what the properties of successive change are.

Consider any two formulae A, B, a particular theory K and any AGM change function $\circ_1$ for K (for example $\circ_1$ may stand as a transitively relational partial meet revision for K). In order to calculate the successive changes of K, first by A and then by B, we need $\circ_1$ for K but also the change function $\circ_2$ relative to $(K \circ_1 A)$. The result of the successive change is the theory $(K \circ_1 A) \circ_2 B$.

The application of a change function over a theory that is the result of another change operation is referred as an *iterated change.*

Once we have understood that AGM change functions are really indexical (relative to the theory to be changed), an obvious first attempt to deal with iteration presents itself. If we possess *beforehand* the complete set of unary change functions, one for each possible theory, we can freely perform successive changes. But beware, if there are no coherence properties linking the different change functions the result obtained can be unexpected and the corresponding behavior erratic. The whole point is then to investigate ways to coordinate these different change functions.

Clearly, binary AGM functions can be trivially iterated, that is, $(K \circ A) \circ B$ is well defined. In particular AGM expansions inherit their capacity of iteration directly from the consequence operation. For instance, for any theory K and formulae A, B, we have $(K + A) + B = \mathrm{Cn}(\mathrm{Cn}(K \cup \{A\}) \cup \{B\}) = \mathrm{Cn}(K \cup \{A\} \cup \{B\}) = K + (A \wedge B)$.

For similar reasons, full meet functions (revisions as well as contractions) also validate that $(K \circ A) \circ B = K \circ (A \wedge B)$. The fact that full meet functions can be iterated can be taken as an evidence for the compatibility of the AGM theory with iterated change. However, these are too specific binary AGM functions and the properties they satisfy we do not want them to hold as properties of binary AGM functions in general.

3.2.1 The Property of Historic Memory

Some advantages of binary functions as a scheme of iterated change are evident: they are mathematically elegant, definitionally simple and remain close to the AGM model (each time the theory argument is fixed a standard AGM indexical function is obtained). But, while formally attractive, binary functions make a strong simplifying assumption. Each theory is modified in a predetermined way independently of how we have obtained such a theory. A binary change function $\circ : \mathbb{K} \times L \rightarrow \mathbb{K}$ is deterministic with respect to the theory to be modified, i.e. it satisfies:

(Functionality) If $K = ((H \circ A_1) \ldots \circ A_n)$. then $K \circ A = ((H \circ A_1) \ldots \circ A_n) \circ A$, But if K is really considered an *argument* of the function $\circ$, this is to be expected. If f is a function, it is required that $f(a) = f(b)$ whenever $a = b$. This functional behavior has been interpreted as a *lack of historic memory*. Lehmann in [1995] refers to this property as a "Non Postulate" for he considers that interesting systems should not make this simplifying assumption.

In spite of the modesty of binary functions as operators for iterated change, they vary in the subtlety of their associated behavior. According to the representation results, an AGM change function can be characterized by some ordered structure. It can be an ordering of formulae in the object language as in the epistemic entrenchment approach, an ordering of possible worlds as in systems of spheres or an ordering over maximal consistent sets for partial meet functions. From this representation perspective binary AGM functions vary according to the sophistication of their associated structure.

In the simplest case we have binary functions that depend on no order at all, as expansion and the full meet functions [Alchourrón and Makinson, 1982]. A quite elaborate binary function, outside the AGM framework, is Katsuno and Mendelzon's update [Katsuno and Mendelzon, 1992]. Based on a fixed set of orders of possible worlds (one order relative to each possible world), the update function is obtained as a fixed combination of such multiple orders.

Proposals for iterated change that possess historic memory ought to expand the AGM model in such a way that change functions return not only the modified theory but also a modified version of the change function, or equivalently, return enough information to construct a new change function. Usually a method or algorithm to construct the new change function based on the original theory, the input formula and the *previous change function* is specified.

This can be done in a qualitative way as in [Boutilier, 1996; Nayak, 1994; Segerberg, 1997], or by enriching the model with numbers [Spohn, 1987; Williams, 1994; Darwiche and Pearl, 1997]. Rott in [1998] englobes them under the name of *iterative functions* and gives a thorough comparison. These are not really going back to a binary function and returning the theory K to its original role of argument. The "construction" method is more flexible than considering a binary function. These approaches are very rich — they can avoid the functional behaviour of the change. But they are usually complex. In these frameworks, given a theory K and a formula A, the change function associated with $K \circ A$ is not uniquely determined and depends really on a third argument: the change function for K. But, as insightfully discussed by [Rott, 1999], this is a circular description.

There are two alternative formalizations of iterative functions that circumvent the circularity. One is to consider iterative functions as

$$\circ : \langle \mathbb{K}, \circ_{\mathbb{K}} \rangle \times L \rightarrow \langle \mathbb{K}, \circ_{\mathbb{K}} \rangle$$

that operate not just over a theory and a formula but on a more complex structure: a theory together with the AGM function relative to such theory. Depending on the chosen representation, each change function relative to a given theory boils down to some ordering relation, over subsets, over formulae, or over possible worlds. Let us observe that when working with theories these orderings are always infinite relations, which may or may not be finitely specified. The AGM idealization has been altered so that these iterative functions are binary functions whose first argument is quite complex. They return also a complex structure encoding the resulting theory and enough information as to define a standard AGM function for it.

The other alternative formalization for iterative functions is presented by [Rott, 1999]. He defines iterative functions as unary functions that take a *sequence* of logical formulae and return a plain theory. An iterative change function

$$\circ : L^{\omega} \rightarrow \mathbb{K}$$

assigns for each sequence of input formulae the theory resulting after all the successive changes indicated in the sequence. Rott explains that these unary functions are relative to a state, a complex structure consisting of a theory together with its "changing criteria". Like in the previous formalization a state can be regarded as a theory of together with the standard AGM change function relative to such a theory.

Although the signatures of the two formalizations are quite different it is possible to visualize any of the above mentioned methods in both of them. We regard Rott's formalization as more elegant and closer in spirit to AGM's.

3.3 Some Properties

Clearly, we expect that not any binary function will qualify as coherent. Only those constrained in a certain way ought to be admissible theory change functions. In the words of [Rott, 1999], "...The most general idea to express conditions of coherence seems to be that the change function should be a structure preserving function, a morphism, in the sense that the values of the function stand in some special relation whenever the arguments of the function stand in a special relation." As we already argued, AGM postulates can be regarded as coherence constraints over the unary functions separately. In this section we will review many properties that have been presented in the literature, which will be of interest in the next chapters.

Properties of Binary Functions

We will first examine a property that was originally studied in [Alchourrón and Makinson, 1982] as an intuitive property for change functions. By means of the Levi identity we know that AGM revisions can be defined from contractions.

Levi. $K * A = (K - \neg A) + A$.

Alchourrón and Makinson [1982] wondered under which conditions, an AGM revision function could validate the following intuitive condition.

(Permutability) $(K - \neg A) + A = (K + A) - \neg A$.

In particular, full meet revisions functions are permutable, but the question under which conditions an AGM function is permutable was left open in that paper.

Hansson [1998] proposes reversing the Levi identity as an alternative and plausible way to define revision when change functions are applied to sets of formula that are not closed under logical consequence (bases).

R-Levi. $K * A = (K + A) - \neg A$.

Thus, permutable revisions are equivalently defined by the Levi and the R-Levi identity.

Consider now $\circ$ to be a generic binary revision function $\circ : \mathbb{K} \times L \to K$. Let's first refer to the Monotony property, which is indeed a postulate of the

AGM expansion function. a postulate of the AGM expansion function.

(Monotony) If $K \subseteq H$ then $K \circ A \subseteq H \circ A$.

For an arbitrary change function, it is strong coherence property that it may not always be desirable. For instance, as we have already shown in Observation 2.1 general partial meet functions do not validate it, because monotony collapses with the preservation property. However, full meet revisions satisfy:

(Weak Monotony) If $\neg A \in K$ and $K \subseteq H$ then $K \circ A \subseteq H \circ A$.

In the context of their safe contraction functions [Alchourrón and Makinson, 1985] have considered properties of the intersection and union of theories.

(Weak Intersection) If $\neg A \in (K_1 \cap K_2)$ then $(K_1 \cap K_2) \circ A = (K_1 \circ A) \cap (K_2 \circ A)$.

(Weak Union) If $\neg A \in K_1 \cap K_2$, then $(K_1 \cup K_2) \circ A = (K_1 \circ A) \cup (K_2 \circ A)$.

These properties truly relate the change of arbitrary theories. Quite trivially it can be proved that AGM expansions satisfy and full meet functions validate the two. In addition, they also validate the following D-Ventilation condition that is dual to the Ventilation condition of [Alchourrón *et al.*, 1985]. AGM provided the Ventilation as a factoring condition on the contraction by a conjunction from a theory; they proved it to be equivalent to postulates (K*7) and (K*8).

(D-Ventilation) $(K_1 \cap K_2) \circ A \in \{(K_1 \circ A) \cap (K_2 \circ A), K_1 \circ A, K_2 \circ A\}$.

However, full meet functions do not in general validate the following two properties:

(Intersection) $(K_1 \cap K_2) \circ A = (K_1 \circ A) \cap (K_2 \circ A)$.

(Union) $(K_1 \cup K_2) \circ A = (K_1 \circ A) \cup (K_2 \circ A)$.

Full meet functions validate:

(Commutativity) $(K \circ A) \circ B = (K \circ B) \circ A$.

(Elimination) $(K \circ A) \circ B = K \circ (A \wedge B)$.

The interest of the Elimination property is that it provides a way to reduce the iteration of functions to a plain single application of the change function.

Areces and Rott [1999] have recently devised iterative functions based on the principle of Elimination, where the change of a theory by a sequence of formulae is recast to a standard AGM operation by a formula obtained from the sequence itself. In the case of expansions and full meet functions, Commutativity follows from Elimination. However, in Rott and Areces' framework this is not the case. Their novel function is a binary function $\circ : L^{\omega} \rightarrow I\!\!K$, having as first argument the theory to be revised and the second argument is a sequence of formulae, the initial segment of the history so far. The revision of the original K by a sequence of formulae $[A_1, \ldots, A_n]$ yields a theory $K \circ [A_1, \ldots, A_n]$. This is obtained by applying a standard AGM revision function to the theory K, but the formula to revise it by is obtained as a boolean combination of the formulae in the input sequence. Rott and Areces provide various algorithms to calculate such a formula. It is interesting that Rott and Areces' function may possess historic memory even though it is based on a standard AGM revision function relative to the original theory.

Properties of Iterated Change

We will now inspect some properties of iterated change arising from different proposals. Let's assume now that $\circ$ is a generic iterative function. As we have argued above, trivially, binary functions over $I\!\!K \times L$ give rise to iterative functions; henceforth, some of the properties we will discuss are candidate properties for binary functions.

The following two conditions have been reported by Schlechta Lehmann and Magidor in [1996] as plausible properties for iterated change. For any pair of theories K_1, K_2 and sentences A, B, C, D,

(Or-Right) If $D \in (K \circ A) \circ C$ and $D \in (K \circ B) \circ C$ then $D \in (K \circ (A \vee B)) \circ C$.

(Or-Left) If $D \in (K \circ (A \vee B)) \circ C$ then $D \in (K \circ A) \circ C$ or $D \in (K \circ B) \circ C$.

For (Or-Right) suppose that after successive changes that differ only at step i (step i being A in one case and B in the other), one concludes that D holds. Then, one should also conclude D after identical successive changes when step

i is replaced by the disjunction $A \vee B$. We expect D to hold because knowing which of A or B is true at step i can not be crucial.

The case for (Or-Left) is similar. If one concludes D from the change by a disjunction, one should conclude it from at least one of the disjuncts.

Lehmann in [1995] argues that certain structures that he calls widening ranked orders are suitable for iterative change and proposes seven postulates that fully characterize revision functions based on these structures. In our notation they are:

(I1) $K \circ A$ is a consistent theory.

(I2) $A \in K \circ A$.

(I3) If $B \in K \circ A$, then $A \supset B \in K$.

(I4) If $A \in K$ then $K \circ B_1 \circ \ldots \circ B_n = K \circ A \circ B_1 \circ \ldots \circ B_n$ for $n \geq 1$.

(I5) If $A \in \mathrm{Cn}(B)$, then $K \circ A \circ B \circ B_1 \circ \ldots \circ B_n = K \circ B \circ B_1 \circ \ldots \circ B_n$.

(I6) If $\neg B \notin K \circ A$ then $K \circ A \circ B \circ B_1 \circ \ldots \circ B_n = K \circ A \circ (A \wedge B) \circ B_1 \circ \ldots \circ B_n$.

(I7) $K \circ \neg B \circ B \subseteq \mathrm{Cn}(K \cup B)$.

Postulates (I1-I4) are a direct transcription of AGM's. (I5) states that certain steps in a sequence of changes are negligible. The sequence containing a formula immediately followed by a logically stronger formula produces the same result as the counterpart sequence that lacks the logically weaker formula. Intuitively it says that if immediately after learning some information we obtain more specific information, the first learning is inconsequential. We consider that this condition is controversial, or without enough grounds to be a generally valid principle. Postulate (I6) also asserts that under certain circumstances two sequences give the same result; in particular, when new information is consistent with the theory obtained so far. In this case the formulae at steps i and $i-1$ can be replaced by the the single formula that is conjunction of the two. (I7) implies dependency between two revision steps and consequently enforces (at least to some extent) the property of historic memory, which in general binary functions lack.

Darwiche and Pearl [1997] have proposed a number of properties for iterated change. In our notation:

(C1) If $A \in \mathrm{Cn}(B)$ then $(K \circ A) \circ B = K \circ B$.

(C2) If $\neg A \in \mathrm{Cn}(B)$ then $(K \circ A) \circ B = K \circ B$.

(C3) If $A \in K \circ B$ then $A \in (K \circ A) \circ B$.

(C4) If $\neg A \notin K \circ B$ then $\neg A \notin (K \circ A) \circ B$.

(C5) If $\neg B \in K \circ A$ and $A \notin K \circ B$ then $A \notin (K \circ A) \circ B$.

(C6) If $\neg B \in K \circ A$ and $\neg A \in K \circ B$ then $\neg A \in (K \circ A) \circ B$.

While (C1)-(C4) have been proposed as desirable properties of iterated revisions, (C5) and (C6) have been considered too demanding. Condition (C1) amounts to Lehmann's (I5) and condition (C2) has been proved inconsistent with the AGM postulates (K*7) and (K*8) for binary change functions [Lehmann, 1995].

We shall now consider four postulates for iterated change. We call the first one a *trivial revision*, and it will be studied in Chapter 7. It reduces the revision by a sequence just to the revision by the last sentence.

(T) $K \circ A \circ B = K \circ B$.

It is quite obvious that the scheme it induces not only lacks historic memory, but is actually a fake scheme of iterated change. In addition it conflicts with the AGM postulate (K*4), which requires that it $\neg B \notin K \circ A$ then $(K \circ A) \circ B = (K \circ A) + B$. (T) and (K*4) would require that for all $A, B \in L$ such that $\neg B \notin K \circ A$, $(K \circ A) + B = K \circ B$, which is not generally valid.

For the next three postulates for iterated change we follow the presentation of [Rott, 1998]. The first one is a *conservative revision*. It has been firstly proposed as a possible worlds construction by Boutilier [1996] that he called a natural revision. Darwiche and Pearl have supplied the missing completeness theorem for Boutilier's operation ([Darwiche and Pearl, 1997], Theorem 11), providing the following postulate.

(C) If $\neg B \in K \circ A$, then $K \circ A \circ B = K \circ B$.

The conservative revision function validates all AGM postulates and (C), hence, when B is consistent with $K \circ A$, $K \circ A \circ B$ =Cn(($K \circ A) \cup \{B\}$). Rott has defined the same function as an epistemic entrenchment construction [1998], and proved the corresponding representation theorem. Boutilier in [1992c; 1993] has considered sequences of these revisions and has analyzed different reductions that can be involved in calculating the final result.

The criticism to conservative functions is that they privilege new information at the highest priority possible, but it is given up all to readily when more information comes in. In Rott's words [1998], " They provide no consistent attitude toward novelty. The most recent information is always embraced without reservation, but the last but one piece of information, however, is treated with utter disrespect."

The second iterative function is the *irrevocable revision* of [Segerberg, 1997]. Its characteristic postulate is just the Elimination property of the previous section.

(I) $K \circ A \circ B = K \circ (A \wedge B)$

Iterative functions are relative to a single AGM function relative to the original theory K. The postulate induces an irrevocable scheme because the sequence of revisions by contradictory formulae results in the inconsistent theory, and it is impossible to overcome the inconsistency by applying further revisions. So, in order to avoid an inconsistent result the conjunction of the formulae to revise by in successive revisions has to be logically consistent. Fermé in [1999] has given the characterization result of irrevocable revisions in terms of epistemic entrenchment in a form that is close to the constructions reprted in [Rott, 1991].

Finally, there are *moderate revisions*, as a compromise between the conservative and the irrevocable. They were proposed by Nayak [1994] as an epistemic entrenchment construction. Its characteristic postulate is:

(M)

$$K \circ A \circ B = \begin{cases} K \circ B & \text{, if } \neg B \in \text{Cn}(A) \\ K \circ (A \wedge B) & \text{, otherwise.} \end{cases}$$

Moderate functions always give priority to the new incoming information and, unless the new formula is logically inconsistent with the previous, the resulting theory should accommodate all the formulae in the sequence of revisions. Among the models that account for historic memory Nayak's seems to be the best model one can get.

Finally, the work of [Segerberg, 1995] and [Cantwell, 1997] on hypertheories promises a new perspective on iterative functions.

Chapter 4

Iterable AGM Functions

The aim of this chapter is to define *iterable AGM functions*, binary functions that satisfy all AGM postulates, but are close to being a constant function. We call them iterable because they provide a definitionally simple scheme for iterated change. We provide extended definitions for each of the five AGM presentations, meet functions, systems of spheres, postulates, epistemic entrenchments and safe hierarchies, and prove their equivalence.

The basic idea dates back to Alchourrón and Makinson's work on safe contractions [Alchourrón and Makinson, 1985]. Interestingly, in their paper they study some properties of the safe contraction function with respect to the intersection and union of theories, and also properties of "multiple contractions." They say [Alchourrón and Makinson, 1985], p. 419:

> "... we shall turn to questions that arise when K (the set of propositions) is allowed to vary. [...]But in the case of safe contraction the way of dealing with variations of K is quite straightforward. As we are working with a relation $<$ over K the natural relation to consider over a subset K' of K is simply the restriction $< \cap (A' \times A')$ of $<$ to A'."

They obtain a general result relating $K' - A$ to $K - A$, when $K' \subseteq K$. As a special case they apply it to $(K - B) - C$, since $K - B \subseteq K$ always holds. Although not explicit in their article, a particular case of Alchourrón and Makinson's proposal is to start with a hierarchical order over all the formulas of the language. The simple restriction of the hierarchy over L to the elements of any theory K provides for a hierarchy over such a theory, hence, an appropriate relation for the definition of a safe contraction function for K. This setting yields a binary contraction function based on a unique fixed order of all the formulae, the safe hierarchy.

Reusing the same fixed order makes sense, for example as pointed out by I. Levi (indirect personal communication), when involved in tentative reasoning: a fixed set of facts and laws which are known beforehand constitute the background knowledge from which a sequence of consistent, but tentative, inference steps are performed to reach a conclusion. We will come back to this idea in Section 4.1.5.

The following sections are devoted to the definition of iterable AGM functions in each of the classical presentations, following the ideas we just explained for safe contractions. Notice that since contraction and revision are interdefinable in the AGM framework via the Levi and Harper identities, the task of providing iterable change functions can be reduced to defining just one of them (see Section 4.3 for further details).

4.1 The Five Presentations

4.1.1 Extended Safe Contraction Functions

Let's recall the definitions. A relation $<_{s,f}$ over a set K is a *hierarchy* if it is acyclic: for any set of elements $A_1, \ldots, A_n \in K, n \geq 1$, it is not the case that $A_1 <_{s,f} A_2 <_{s,f} \ldots A_n <_{s,f} A_1$. A relation $<_{s,f}$ over K *continues up* Cn if for every $A_1, A_2, A_3 \in K$, if $A_1 <_{s,f} A_2$ and $A_3 \in \mathrm{Cn}(A_2)$ then $A_1 <_{s,f} A_3$. A relation $<_{s,f}$ over K is *virtually connected* if for every $A_1, A_2, A_3 \in K$ if $A_1 <_{s,f} A_2$ then either $A_1 <_{s,f} A_3$ or $A_3 <_{s,f} A_2$. Let $<_{s,f}$ be a virtually connected hierarchy over a theory K that continues up Cn, and let A be a sentence. The safe contraction function over K, $-^K : L \to I\!K$, is defined as:

$$-^K(A) = \mathrm{Cn}(\{B \mid \forall K' \subseteq K, \text{ s.t. } A \in \mathrm{Cn}(K') \text{ and } K' \text{ is } \subseteq\text{-minimal with this property, } B \notin K' \text{ or there is } C \in K' \text{ s.t. } C <_{s,f} B\}).$$

The formulae B of $-^K(A)$ are called the safe elements of K with respect to A since they can not be "blamed" for implying A. An element is safe for A if it does not belong to any of the $\subseteq$-minimal subsets of K that imply A, or else it is not $<_{sf}$-minimal in the hierarchy in such subsets.

Following Alchourrón and Makinson's idea of restricting the hierarchical order, we can define the iterable safe contraction function based on a hierarchy over all the sentences of L.

Definition 4.1 (Derived Order) Let $<_{sf}$ be a hierarchy over the language L. Then for any theory K the derived hierarchy $<^K_{sf}$ is defined as $<^K_{sf} = <_{sf}|K$ (where $R|X$ is the restriction of R to the elements in X).

Observation 4.2 Let $<_{sf}$ be a virtually connected hierarchy that continues up Cn in L, then for any theory K the relation $<^K_{sf}$ is a virtually connected hierarchy and continues up Cn in K.

PROOF. Trivial. The properties of being acyclic, virtually connected and continuing up Cn are preserved under taking restrictions to theories. QED

Once this result is obtained, to define an iterable safe contraction is straightforward. We define the binary function $-_{sf}: \mathbb{K} \times L \to \mathbb{K}$.

Definition 4.3 (Iterable Safe Contraction) Let $<_{sf}$ be a virtually connected hierarchy that continues up Cn in L. The iterable AGM contraction $-_{sf}: \mathbb{K} \times L \to \mathbb{K}$ is defined as

$$K -_{sf} A = \mathrm{Cn}(\{B \mid \forall K' \subseteq K, \text{ s.t. } A \in \mathrm{Cn}(K') \text{ and } K' \text{ is } \subseteq\text{-minimal with this property, } B \notin K' \text{ or there is } C \in K' \text{ s.t. } C <^K_{sf} B\})$$

where $<^K_{sf}$ is the derived safe hierarchy for K.

That $-_{sf}$ satisfies the AGM postulates (K-1) to (K-8) follows from Alchourrón and Makinson's original results stating that every safe contraction function generated by a virtually connected hierarchy that continues up Cn over a theory K is a transitively relational partial meet contraction function.

As a side remark, notice that definitions 4.1 and 4.3 can be merged in a unique definition and $-_{sf}$ defined then directly over $<_{sf}$ instead of over $<^K_{sf}$. This is just a matter of notation, as in both cases $-_{sf}$ is really a binary function as required. This remark applies as well to the definitions of iterable functions in the remaining presentations.

In the definitions above we started from a hierarchy $<_{sf}$ for L and defined its restriction $<_{sf}^{K}$. A relevant question is whether the converse can also be achieved. Given a hierarchy for K can a hierarchy for L be defined such that the iterable function agrees with $-^{K}$ when applied to K?

Observation 4.4 Let $-^{K}$ be an AGM safe contraction function for a given theory K. Then $-^{K}$ can be extended to an iterable AGM safe contraction $-_{sf}$, such that for every A, $K -_{sf} A = -^{K}(A)$.

PROOF. Given $<_{sf}^{K}$ the order associated to $-^{K}$, define $<_{sf}$ as follows: $A <_{sf} B$ iff either $(A \notin K)$ or $(A, B \in K$ and $A <_{sf} B)$. Intuitively, when extending the order to the whole language, elements in K are promoted in their safeness while elements outside K are minimally safe. From the definition $<_{sf}^{K} = <_{sf}|K$, and it is not hard to check that $<_{sf}$ is a virtually connected hierarchy that continues up Cn over L. QED

In [Hansson, 1994] the safe contraction approach is generalized to a "kernel contraction". Instead of implementing a relational way of defining "safe elements", selection functions (called incision functions) are introduced. Our results for safe contraction can easily be extended to kernel contraction.

4.1.2 Extended Partial Meet Contraction Functions

The principle of information economy requires that $K - A$ contains as much as possible from K without entailing A. For every theory K and sentence A, the set $K \perp A$ of maximal subsets of K that fail to imply A is the definitional basis for partial meet contraction functions.

$$K \perp A = \{K' \subseteq K \mid A \notin \mathrm{Cn}(K') \text{ and } K' \text{ is } \subseteq\text{-maximal with this property}\}.$$

A selection function is a function which returns a nonempty subset of a given nonempty set. Let K be a theory, we note as $s^{K} : L \rightarrow \mathcal{P}(\mathcal{P}(K)) \setminus \{\emptyset\}$, a selection function for $K \perp A$, for $A \in L$. We furthermore require that $s^{K}(A) = \{K\}$ whenever $K \perp A = \emptyset$. The original AGM partial meet contraction function $-^{K}$

is then defined, for a theory K, as

$$-^{K}(A) = \bigcap s^{K}(A), \text{ where } s^{K} \text{ is a selection function for } K.$$

Under this definition the contraction function $-^{K}$ satisfies the basic AGM postulates (K-1) to (K-6). To satisfy the extended set of postulates, (K-1) to (K-8), it suffices that s^{K} be *transitively relational*, i.e. for each $A \in L$ the selection function returns the smallest elements according to some transitive relation defined over $K \perp A$.

In order to define an iterable version of $-^{K}$ richer than the full meet contraction, we need to obtain somehow the selections functions s^{K}, one for each eventual K. Of course, we might assume to have all the selection functions beforehand. But following the ideas presented in the extension of safe contraction functions, we would rather synthesize the different s^{K} out of a unique structure.

The largest possible theory is L, the whole language. Then s^{L} provides for each formula A a selection function over all the maximal consistent sets of L that do not imply A. It is possible to extract from s^{L} the corresponding s^{K} for each theory K. This is a consequence of the following two observations: (a) If $A \notin K$, then, trivially, the maximal consistent subset of K that fails to imply A is K itself. (b) If $A \in K$, each maximal consistent subset of K that fails to imply A is included in a maximal consistent subset of L that fails to imply A (by a Lindenbaum-style argument, each element in $K \perp A$ can be extended to an element of $L \perp A$). Therefore, we can derive a selection function $s^{K}(A)$ by just restricting the result of $s^{L}(A)$ to its common part with K.

Definition 4.5 (Derived Selection Functions) Let s^{L} be a selection function for L. Then, for any theory K the selection function s^{K} is

$$s^{K}(A) = \begin{cases} \{K\} & \text{if } A \notin K \\ \{K' \in K \perp A \mid K' = K \cap H' \text{ with } H' \in s^{L}(A)\} & \text{otherwise.} \end{cases}$$

It is immediate that each derived s^{K} is indeed a selection function. What is more interesting is that each s^{K} is transitively relational whenever s^{L} is.

Observation 4.6 If s^{L} is a transitively relational selection function, then for any theory K, s^{K} is a transitively relational selection function.

PROOF. The intuition is as follows, as s^L is transitively relational there is a transitive relation R defined over $L \perp A$ whose smallest elements are selected by $s^L(A)$. This relation R can be projected over each $K \perp A$ to show that $s^K(A)$ selects the smallest elements of a transitive relation. QED

Given that s^L is a transitively relational selection function we are able to define an iterable AGM contraction function $-_{pm}$ based on the partial meet construction.

Definition 4.7 (Iterable Partial Meet Contraction) Let s^L be a transitively relational selection function over L. The iterable AGM contraction $-_{pm} : \mathbb{K} \times L \to \mathbb{K}$ is defined as $K -_{pm} A = \bigcap s^K(A)$, where s^K is the derived selection function for K.

By construction $-_{pm}$ is an AGM transitively relational partial meet contraction. It is iterable as it is applicable to any theory K. We now prove that every AGM partial meet contraction function can be extended to an iterable partial meet.

Observation 4.8 Let $-^K$ be an AGM transitively relational partial meet contraction function for a given theory K. Then $-^K$ can be extended to an iterable AGM partial meet contraction $-_{pm}$, such that for every A, $K -_{pm} A = -^K(A)$.

PROOF. Given a selection function s^K we have to come up with a selection function s^L. As we previously said, for each $H \in K \perp A$ there is $H' \in L \perp A$ such that $H \subseteq H'$. Hence, we can define $s^L(A) = \{H' \in L \perp A \mid \exists H \in s^K(K \perp A) \text{ and } H \subseteq H'\}$. Notice that there can be some $H' \in L \perp A$ such that there exists no subset H of K and $H \subseteq H'$, so that H' is not selected.

Since s^K is transitively relational there is a relation R over $K \perp A$ which can be lifted to $L \perp A$. If $R(H_1, H_2)$ then $R'(H_1', H_2')$ for $H_1', H_2' \in L \perp A$ such that $H_i \subseteq H_i'$. For every $H' \in L \perp A$ such that there exists no subset H of K and $H \subseteq H'$, we define $R'(H'', H')$ for every $H'' \in L \perp A$. Now $s^L(A)$ selects the smallest elements of R'. It follows from the definition that R' is transitive, hence s^L is transitively relational. QED

4.1.3 Extended Systems of Spheres

In this section we develop a definition of an iterable contraction function based on Systems of Spheres, which turns out to be equivalent to an early unpublished result of Makinson (personal communication). We first turn to Grove's original framework [Grove, 1988] for AGM functions.

A system of spheres S centered on a theory K is a set of sets of possible worlds that verifies the properties:

(S1) If $U, V \in S$ then $U \subseteq V$ or $V \subseteq U$. (Totally Ordered.)

(S2) For every $U \in S$, $[K] \subseteq U$. (Minimum.)

(S3) $W \in S$. (Maximum.)

(S4) For every sentence A such that there is a sphere U in S with $[A] \cap U \neq \emptyset$, there is a $\subseteq$-minimal sphere V in S such that $[A] \cap V \neq \emptyset$. (Limit Assumption.)

For any sentence A, if $[A]$ has a non-empty intersection with some sphere in S then by (S4) there exists a minimal such sphere in S, say $c_s(A)$. But, if $[A]$ has an empty intersection with all spheres, then it must be the empty set (since (S3) assures W is in S), in this case $c_s(A)$ is put to be just W. Given a system of spheres S and a formula A, $c_s(A)$ is defined as:

$$c_s(A) = \begin{cases} W & \text{if } [A] = \emptyset \\ \text{the } \subseteq \text{-minimal sphere } S' \text{ in } S \text{ s.t. } S' \cap [A] \neq \emptyset & \text{otherwise.} \end{cases}$$

Using the function c_s, the function $f_s : L \to \mathcal{P}(W)$ is defined as $f_s(A) = [A] \cap c_s(A)$. Given a sentence A, $f_s(A)$ returns the closest elements (with respect to theory K) where A holds. Grove shows that the function defined as $-^{K}(A) = \text{Th}([K] \cup f_s(\neg A))$ is an AGM contraction function. And conversely, for any AGM contraction function relative to a theory K there is a system of spheres S centered on K that gives rise to the same function.

We shall now extend Grove's construction to obtain an iterable function using the same strategy we used for partial meet. Again, the central idea is to consider the inconsistent theory. A system of spheres for L has the particular property that its innermost sphere is the empty set, since $[L] = \emptyset$. Given a system of spheres S centered in $\emptyset$ we define for any theory K a derived system S^K centered on K simply by "filling in" the innermost sphere of S with $[K]$.

Definition 4.9 (Derived System of Spheres) Let S be a system of spheres for L. Then for any theory K the derived system of spheres S^K is defined as $S^K = \{[K] \cup S_i \mid S_i \in S\}$.

Observation 4.10 Let S be a system of spheres for L. Then for any theory K, S^K is a system of spheres centered on K.

Having defined the method to derive a system of spheres S^K, the functions c_S^K and f_S^K are as above. We can now define the iterable contraction function $-_{ss} : \mathbb{K} \times L \to \mathbb{K}$, applicable to every theory K and every formula A.

Definition 4.11 (Iterable Sphere Contraction) Let S be a system of spheres for L. The iterable AGM contraction $-_{ss} : \mathbb{K} \times L \to \mathbb{K}$ is defined as $K -_{ss} A = \text{Th}([K] \cup f_S^K(\neg A))$, where f_S^K is the derived function for K.

It is clear that $-_{ss}$ is iterable. By Grove's characterization result it follows that $-_{ss}$ is an AGM contraction function. We prove that every AGM contraction function can be extended to an iterable sphere contraction function.

Observation 4.12 Let $-^K$ be an AGM contraction functions based on systems of spheres. Then $-^K$ can be extended to an iterable AGM contraction $-_{ss}$ based on systems of spheres, such that for every A, $K -_{ss} A = -^K(A)$.

PROOF. It is enough to prove that if S^K is a system of spheres for K, then it can be extended to a system of spheres for L. Define S centered in $\emptyset$ as $S = S^K \cup \{\emptyset\}$. Clearly, S validates (S1) to (S4) for L. QED

4.1.4 Extended Epistemic Entrenchments

An *epistemic entrenchment* for a theory K is a total relation among the formulae in the language reflecting their degree of relevance in K and their usefulness when performing inference. The following five conditions must hold for an epistemic entrenchment relation $\leq_{ee}$ for a theory K [Makinson and Gärdenfors, 1988]:

(EE1) If $A \leq_{ee} B$ and $B \leq_{ee} D$ then $A \leq_{ee} D$.

(EE2) If $B \in \text{Cn}(A)$ then $A \leq_{ee} B$.

(EE3) $A \leq_{ee} (A \wedge B)$ or $B \leq_{ee} (A \wedge B)$.

(EE4) If theory K is consistent then $A \notin K$ iff $A \leq_{ee} B$ for every B.

(EE5) If $B \leq_{ee} A$ for every B then $A \in \mathrm{Cn}(\emptyset)$.

The AGM contraction function $-^{K}$ based on an epistemic entrenchment relation $\leq_{ee}$ for K, is defined as follows. For every formula A in L,

$$-^{K}(A) = \{B \in K \mid A \in \mathrm{Cn}(\emptyset) \text{ or } A <_{ee} (A \vee B)\},$$

where $<_{ee}$ is the strict relation obtained from $\leq_{ee}$.

For any given relation $\leq_{ee}$ for a consistent theory K, the formulas in K are ranked in $\leq_{ee}$, while all the formulas outside K have the $\leq_{ee}$-minimal epistemic value. That is, by (EE4) for a consistent theory K, all the formulas outside K are zeroed in $\leq_{ee}$. However, (EE4) is vacuous for the contradictory theory L. If we consider as a point of departure an epistemic entrenchment over the contradictory theory L, there is an obvious way to derive an entrenchment order for any theory K: just depose the formulas not in K to a minimal rank.

Definition 4.13 (Derived Epistemic Entrenchment) Let $\leq_{ee}$ be an epistemic entrenchment relation for L. Then for any theory K the derived epistemic entrenchment relation $\leq_{ee}^{K}$ is defined as:

$$A \leq_{ee}^{K} B \text{ iff either } (A \notin K) \text{ or } (A, B \in K \text{ and } A \leq_{ee} B).$$

Again the first step is to establish that our definition is sound.

Observation 4.14 Let $\leq_{ee}$ be an epistemic entrenchment relation for L, then for any theory K, $\leq_{ee}^{K}$ is an epistemic entrenchment relation for K.

Definition 4.15 (Iterable Epistemic Entrenchment Contraction) Let $\leq_{ee}$ be an epistemic entrenchment relation for L. The iterable AGM contraction $-_{ee} : \mathbb{K} \times L \to \mathbb{K}$ is defined as $K -_{ee} A = \{B \in K \mid A \in \mathrm{Cn}(\emptyset) \text{ or } A <_{ee}^{K} (A \vee B)\}$, where $<_{ee}^{K}$ is the asymmetric part of $\leq_{ee}^{K}$, for $\leq_{ee}^{K}$ the derived epistemic entrenchment relation for K.

It remains to show that every contraction function based on epistemic entrenchments can be extended to an iterable contraction function.

Observation 4.16 Let $-^{K}$ be an AGM contraction function based on epistemic entrenchments for a given theory K. Then $-^{K}$ can be extended to an iterable

AGM contraction $-_{cc}$ based on epistemic entrenchments such that for every A, $K -_{cc} A = -^{K}(A)$.

PROOF. The key point is to prove that an epistemic entrenchment relation for $K \leq_{cc}^{K}$ can be extended to a relation for L.

If $K = L$ then we are done. Suppose $K \neq L$. We claim that $\leq_{cc}^{K}$ is also an epistemic entrenchment relation for L. Conditions (EE1), (EE2), (EE3) and (EE5) do not refer to the specific theory so they hold also trivially for L, while condition (EE4) does not apply as L is inconsistent. QED

4.1.5 Extended Postulates

One of the hallmarks of the AGM formalism is that a contraction operation always returns a consistent theory. The largest possible theory is the inconsistent theory L, the whole language. The contraction function over the inconsistent theory can be regarded as a generic removal procedure leading to consistency. As every theory is a subset of the inconsistent theory this generic removal procedure can be applied to any theory. We propose the following postulate:

$$(\mathbf{K-9}) \quad \text{If } A \in K, \text{ then } K - A = (L - A) \cap K.$$

Postulate $K-9$ is extremely simple and reveals the unsophisticated behavior of our iterable contraction function. Its dual iterable revision postulate is defined as:

$$(\mathbf{K*9}) \text{ If } \neg A \in K, \text{ then } K * A = (L * A)$$

In Section 4.3 we elaborate on the inter-definability of (K*9) and (K-9) via the Levi and Harper identities. It becomes obvious that a revision function $*$ satisfying (K*1)- (K*9) is in fact iterable: for any $A, B \in L$, $K * A * B$ is well defined: If $\neg B \in K*A$ then $K*A*B = (L*B)$; else $K*A*B = (K*A)+B$. An immediate observation is that (K*9) forces independence between two arbitrary revision steps. Namely, the result of revising a theory is independent of the preceding steps that lead to it, only the actual theory being revised matters. This is what we have described as lack of historic memory in Chapter 3, or as reported in [Friedman and Halpern, 1996], the qualitative analogue of the Markov Assumption.

The revision postulate (K*9) is sound with respect to the interpretation of revision as a kind of tentative reasoning. The revision function for L encodes a fixed and pre-established criteria, "the way things are" (facts) and "the way things work" (laws) in the actual world. A sequence of revisions is then performed in search of tentative explanations (of the facts) and conclusions (derived from them). When we detect an inconsistency between the hypothesis elaborated up to now and a new supposition we are trying to adjust to the reasoning, we lose confidence in the chain of hypothesis. We should then start it all over, and accommodate the latest piece of our tentative chain in accordance with our (fixed and pre-established) criteria, leaving behind our previous wrong conjectures.

We take (K-1) to (K-9) as defining iterable AGM contraction functions via postulates. We show in the next section that these functions coincide with the iterable AGM contraction functions defined above.

Lemma 7.4 in [Alchourrón and Makinson, 1985] can be considered as the first reference to the ideas put forward in postulate (K-9). But the connection with iteration is first elucidated by [Rott, 1992b]. He mentions explicitly (K-9) in connection with generalized entrenchment relations and considers it as a policy of iteration. He also proves that iterated theory change according to this method reduces to change of the inconsistent theory. Remarkably, [Freund and Lehmann, 1994] proposes precisely the same postulate (K*9) and shows the correspondence between an AGM revision operation satisfying it and a rational consistency-preserving consequence relation. Freund and Lehmann also show that such a revision function admits iteration. Although their postulate and ours turned out to be identical, the two works are indeed complementary. In the attempt to elucidate the meaning and effect of (K-9) we were driven to recast it in the four other standard presentations of AGM (safe hierarchies, partial meet functions, systems of spheres and epistemic entrenchments) and in the next section we will prove that they are indeed equivalent. Freund and Lehmann chose instead to consider the connection existing between theory change and non-monotonic reasoning [Makinson and Gärdenfors, 1991; Gärdenfors and Makinson, 1990] and study the effect of the new postulate on the (non-monotonic) inference relation. The main result in their paper is the proof that revisions satisfying (K*1) to (K*9) stand in one-to-one correspondence with rational, consistency-preserving non-monotonic inference relations.

4.2 Equivalences

In this section we will prove the equivalence of the five systems presented. We first prove that postulates (K-1) to (K-9) characterize the iterable AGM contractions based on systems of spheres.

Theorem 4.17 (Postulates/Systems of Spheres) Given an iterable AGM contraction $-$ satisfying (K-1) to (K-9), there exists a system of spheres S for L such that for every K and every A, $K - A = \text{Th}([K] \cup f_S^K(\neg A))$. Conversely, every $-_{ss}$ based on a system of spheres S for L satisfies postulates (K-1) to (K-9).

Proof. As $-_{ss}$ is a contraction based on systems of spheres it satisfies (K-1) to (K-8). It is trivial to check that it also satisfies (K-9).

By Grove's original result, for any AGM function for L that satisfies (K-1) to (K-8) there is a system of spheres S for L such that $L - A = \text{Th}(f_s(\neg A))$. By definition $S^K = \{[K] \cup S_i \mid S_i \in S\}$. There are two cases. For any $A \notin K$, clearly $f_S^K(\neg A) = [K] \cap [\neg A]$, then $\text{Th}([K] \cup f_S^K(\neg A)) = K$ and by postulate (K-3), $K = K - A$, so we are done. For $A \in K$, $f_S^K(\neg A) = f_S^L(\neg A)$, then $\text{Th}([K] \cup f_S^K(\neg A)) = \text{Th}([K] \cup f_S^L(\neg A)) = K \cap \text{Th}(f_S^K(\neg A))$ $= K \cap (L - A)$, and we are done. QED

We shall prove that $-_{ee}$ and the extended postulates are equivalent.

Theorem 4.18 (Postulates/Epistemic Entrenchments) Given an iterable AGM contraction $-$ that satisfies (K-1) to (K-9), there exists an epistemic entrenchment relation $\leq_{ee}$ for L such that for every K and every A, $K - A = \{B \in K \mid A \in \text{Cn}(\emptyset) \text{ or } A <_{ee}^K (A \vee B)\}$. Conversely, every $-_{ee}$ satisfies (K-1) to (K-9).

Proof. Again, by previous results, $-_{ee}$ satisfies (K-1) to (K-8) and it is easy to verify that it also satisfies (K-9).

Let $\leq_{ee}$ be the epistemic entrenchment guaranteed to exist for any contraction function satisfying (K-1) to (K-8). We already proved that it is an epistemic entrenchment for L.

If $A \notin K$ then by (K-3), $K - A = K$. As $\leq_{ee}$ satisfies (EE1) and (EE4), $A <^{K}_{ee} (A \vee B)$ for all $B \in K$. Hence $K - A = \{B \in K \mid A \in \mathrm{Cn}(\emptyset)$ or $A <^{K}_{ee} (A \vee B)\}$.

Suppose $A \in K$. As $\leq^{K}_{ee}$ is the restriction of $\leq_{ee}$, $K -_{ee} A = \{B \in K \mid A \in \mathrm{Cn}(\emptyset)$ or $A <^{K}_{ee} (A \vee B)\} = K \cap \{B \in L \mid A \in \mathrm{Cn}(\emptyset)$ or $A <_{ee} (A \vee B)\} = (L - A) \cap K = K - A$, if $-$ satisfies (K-9). QED

We have presented $-_{pm}$ and $-_{ss}$, and showed that they are both iterable AGM functions relative to some fixed order for the inconsistent theory L. We now prove that $-_{pm}$ and $-_{ss}$ are in fact equivalent.

Theorem 4.19 (Meet Functions/Systems of Spheres) For each iterable partial meet contraction $-_{pm}$ there exists a system of spheres S for L such that for every theory K and every A, $K -_{pm} A = \mathrm{Th}([K] \cup c^{K}_{S}(\neg A))$. Conversely, for each iterable contraction $-_{ss}$ defined by a system of spheres there exists a selection functions s^{L} such that for every theory K and every A, $K -_{ss} A = \bigcap s^{K}(A)$.

PROOF. The theorem is a direct consequence of the "Grove connection" [Makinson, 1993] relating consistent complete theories in the language of K including A and the elements in $\{K \cup \mathrm{Cn}(A) | K \in K \perp \neg A\}$ by a total injective mapping. In the particular case when we consider the inconsistent theory L, this mapping can be recast as a bijection between the set of all consistent complete theories (worlds) and $\bigcup_{A} \cup (L \perp A)$. Once this connection has been established, the order provided by a system of spheres centered in $\emptyset$ defines a transitively relational selection function s^{L} and vice versa. QED

Finally, by using results in [Rott, 1992a] we can establish the equivalence between iterated epistemic entrenchment contractions and iterated safe contractions functions, proving that the five approaches presented are indeed five faces of the same phenomenon.

Theorem 4.20 (Epistemic Entrenchments/Safe Hierarchies) For each iterable epistemic entrenchment contraction $-_{ee}$ there exists a virtually connected hierarchy $<_{sf}$ that continues up Cn in L, such that for every theory K and every A, $K -_{ee} A = K \cap \text{Cn}(\{B \mid \forall K' \subseteq K,$ s.t. $A \in \text{Cn}(K')$ and K' is a $\subseteq$ -minimal with this property, $B \notin K'$ or there is $C \in K'$ s.t. $C <^K_{sf} B\})$. Conversely, for each safe iterable contraction $-_{sf}$ there exists an epistemic relation $\leq_{ee}$ for L, such that for every theory K and every A $K -_{sf} A = \{B \in K \mid A \in \text{Cn}(\emptyset)$ or $A <^K_{ee} (A \vee B)\}$.

PROOF. The first part is immediate. As it is proved in [Rott, 1992a], an epistemic entrenchment is also a safe hierarchy. Furthermore the relativization to K used during iteration is preserved. For the second part, let $<_{sf}$ be the hierarchy for L associated to $-_{sf}$. Now using the main result in [Rott, 1992a] we can obtain an epistemic entrenchment relation $\leq_{ee}$ such that the associated contraction function behaves as $-_{sf}$ for L. Take $\leq_{ee}$ as the basis for our epistemic entrenchment iterable contraction function $-_{ee}$. If $A \in \text{Cn}(\emptyset)$ or $A \notin K$, then as both $-_{sf}$ and $-_{ee}$ are AGM functions, $K -_{sf} A$ $= K = \{B \in K \mid A \in \text{Cn}(\emptyset)$ or $A <^K_{ee} (A \vee B)\}$. If $A \notin \text{Cn}(\emptyset)$ and $A \in K$, as the functions satisfy (K-9), $K -_{sf} A = (L -_{sf} A) \cap K = (L -_{ee} A) \cap K$ $= \{B \in K \mid A \in \text{Cn}(\emptyset)$ or $A <^K_{ee} (A \vee B)\}$. QED

4.3 Properties

Postulate (K-9) immediately implies a weak form of monotony of the iterable function. This has been noticed by Makinson (personal communication).

(Weak Monotony) If $A \in K$ and $K \subseteq K'$ then $K - A \subseteq K' - A$.

This is a nice coherence property linking the contractions of different theories by the same formula. Also the iterable revision satisfies that,

(Almost Constant) If $\neg A \in K, K'$ then $K * A = K' * A$.

Namely, when the second argument is held fixed, the iterable revision behaves almost as a constant function on its first argument. By (K*9), if $\neg A$ is in K then $K * A$ is constant. But as iterable revisions satisfy AGM postulates (K*3) and (K*4), if $\neg A \notin K$, then $K * A = \text{Cn}(K \cup \{A\})$.

A key observation about iterable contractions is that they ought to be relative to the largest theory, L, because attempting to make them relative to a smaller theory collapses with the AGM recovery postulate.

Observation 4.21 (Makinson, personal communication) Let $-$ be an iterable AGM function. There is no value of theory H distinct from L for which

(K-9(H)): If $A \in K$ then $K - A = (H - A) \cap K$.

is consistent with (K-5) (the postulate of recovery).

PROOF. Suppose $H \neq L$. Choose any A in H (even $\top$ will do).

$L =$, by (K-5), $\mathrm{Cn}(L - A) \cup \{A\}) =$,

by condition (K-9(H)) substituting L for K, $= \mathrm{Cn}(((H - A) \cap L) \cup \{A\}) =$

,by monotony of Cn and de Morgan laws, $\subseteq \mathrm{Cn}(((H - A)) \cup \{A\}) =$

,by (K-5), $= H$, giving us a contradiction. QED

Although iterable AGM functions are binary and almost constant, they do not validate commutativity. In general $(K - A) - B$ is different from $(K - B) - A$. Just as AGM contraction and revision are inter-definable via the Levi and Harper identities, so are iterable AGM contractions and revisions. Specifically, the Levi identity lets us define iterable revision functions:

Levi. $K * A = (K - \neg A) + A$.

This is important since it allows for sequences of different kinds of changes, like for example $(\ldots((K + A) - B) * D \ldots * C)$.

In [Alchourrón and Makinson, 1985] it is shown that under appropriate conditions safe contractions are permutable ([Alchourrón and Makinson, 1985], Lemma 7.1). Given (K-9), an iterable revision function $*$ can be defined in terms of an iterable contraction function equivalently via **Levi** or **R-Levi**.

Observation 4.22 Iterable AGM contraction functions are permutable.

A direct proof of the above is immediate but the result also derives from Lemma 7.1 in [Alchourrón and Makinson, 1985].

As iterable AGM contractions induce safe contractions functions for each theory K, the results proved in [Alchourrón and Makinson, 1985] carry over:

Observation 4.23 Iterable AGM contractions validate:

If $A \in K_1 \cap K_2$ then $(K_1 \cap K_2) - A = (K_1 - A) \cap (K_2 - A)$
If $A \in K_1 \cap K_2$ then $(K_1 \cup K_2) - A = (K_1 - A) \cup (K_2 - A)$.
If $A \in K - B$ and $B \in K - A$ then $(K - A) - B = (K - A) \cap (K - B) = (K - B) - A$.

And similarly for iterable revisions:
(Weak Intersection) If $\neg A \in K_1 \cap K_2$ then $(K_1 \cap K_2) * A = (K_1 * A) \cap (K_2 * A)$
(Weak Union) If $\neg A \in K_1 \cap K_2$ then $(K_1 \cup K_2) * A = (K_1 * A) \cup (K_2 * A)$.
(Weak Commutativity) If $\neg A \in K * B$ and $\neg B \in K * A$ then
$(K * A) * B = (K * A) \cap (K * B) = (K * B) * A$.

In fact, these properties hold not just for two theories but also for indefinitely many. We shall now show that the D-Ventilation property, which is a strengthening of the property Weak Interscetion, is not generally valid for iterable AGM functions.

(D-Ventilation) $(K_1 \cap K_2) * A \in \{(K_1 * A) \cap (K_2 * A), K_1 * A, K_2 * A\}$.

Observation 4.24 There exist iterable AGM revisions that violate D-Ventilation.

PROOF. Let $L * A = \mathrm{Cn}(A \wedge \neg B \wedge \neg C)$. Assume $K_1 = \mathrm{Cn}(\{B, A \equiv \neg C\})$ and $K_2 = \mathrm{Cn}(\{C, A \equiv \neg B\})$. Then $(K_1 \cap K_2) * A = \mathrm{Cn}(\neg A \wedge B \wedge C)$ while $K_1 = \mathrm{Cn}(A \wedge B \wedge \neg C)$ and $K_2 = \mathrm{Cn}(A \wedge \neg B \wedge C)$. QED

Iterable AGM revisions satisfy just three of the six properties of Darwiche and Pearl [1997].

C1. If $A \in \mathrm{Cn}(B)$ then $(K * A) * B = K * B$.
C2. If $\neg A \in \mathrm{Cn}(B)$ then $(K * A) * B = K * B$.
C3. If $A \in K * B$ then $A \in (K * A) * B$.
C4. If $\neg A \notin K * B$ then $\neg A \notin (K * A) * B$.
C5. If $\neg B \in K * A$ and $A \notin K * B$ then $A \notin (K * A) * B$.
C6. If $\neg B \in K * A$ and $\neg A \in K * B$ then $\neg A \in (K * A) * B$.

Observation 4.25

i) All iterable AGM functions, satisfy C1, C3 and C4.

ii) There exist iterable AGM functions violating C2, C5 and C6.

Most noticeably, our iterable AGM functions validate the first six of the seven postulates of [Lehmann, 1995], the seventh requires the property of historic memory that iterable functions obviously lack.

I1. $K * A$ is a consistent theory.

I2. $A \in K * A$.

I3. If $B \in K * A$, then $A \supset B \in K$.

I4. If $A \in K$ then $K * B_1 * \ldots * B_n \equiv K * A * B_1 * \ldots * B_n$ for $n \geq 1$.

I5. If $A \in Cn(B)$, then $K * A * B * B_1 * \ldots * B_n \equiv K * B * B_1 * \ldots * B_n$.

I6. If $\neg B \notin K*A$ then $K*A*B*B_1*\ldots*B_n \equiv K*A*(A \wedge B)*B_1*\ldots*B_n$.

I7. $K * \neg B * B \subseteq Cn(K \cup \{B\})$.

Observation 4.26

i) All iterable AGM functions, satisfy I1-I6.

ii) There exist iterable AGM functions violating I7.

Finally, let's consider the properties in [Schlechta *et al.*, 1996].

(Or-Right) If $D \in (K*A)*C$ and $D \in (K*B)*C$ then $D \in (K*(A \vee B))*C$.

(Or-Left) If $D \in (K*(A \vee B))*C$ then $D \in (K*A)*C$ or $D \in (K*B)*C$.

Observation 4.27 Iterable AGM functions satisfy Or-Right and Or-Left.

We observe that iterable functions do not comply with the properties associated to iterative schemes.

(T) $K \circ A \circ B = K \circ B$.

(C) If $\neg B \in K \circ A$, then $K \circ A \circ B = K \circ B$.

(I) $K \circ A \circ B = K \circ (A \wedge B)$

(M)

$$K \circ A \circ B = \begin{cases} K \circ B & \text{, if } \neg B \in \text{Cn}(A) \\ K \circ (A \wedge B) & \text{, otherwise.} \end{cases}$$

Observation 4.28 There exist iterable functions violating (T),(M),(I) and (C).

Modest as they are, it is surprising that iterable AGM functions satisfy a good number of the standard properties put forward as relevant for iterated change.

Chapter 5

Update Functions

In this chapter we will concentrate on a distinctive binary function outside the AGM framework, Katsuno and Mendelzon's update [Katsuno and Mendelzon, 1992]. We will study various of its formal properties and we will complete and clarify previous results that extended the function for infinite languages.

Some years after the seminal paper of Alchourrón, Gärdenfors and Makinson, Katsuno and Mendelzon presented a new theory change operation which they called an *update*. In their paper, Katsuno and Mendelzon compared the update operation with the previous revision operation and following the work of [Keller and Winslett, 1985] provided some interesting remarks on the differences between the two approaches: while revision functions seemed well suited for modeling the change provoked by evolving knowledge about a static situation, update operations captured the change in knowledge provoked by an evolving situation. We quote [Katsuno and Mendelzon, 1992], page 387:

> "We make a fundamental distinction between two kinds of modifications to a knowledge base. The first one, *update* consists to bringing the knowledge base up to date, when the world described by it changes. ... The second kind of modification, *revision* is used when we are obtaining new information about a static world. ... We claim the AGM postulates describe only revisions."

The two forms of change can be illustrated with the following example, which is an adaptation of Katsuno and Mendelzon's original one.

Suppose that each day I either have no breakfast at all or I have coffee and toast. Suppose you are now informed that I had coffee at breakfast. How should you incorporate this information into your knowledge? You could take it as an indication that I had coffee and toast, with which moreover it is consistent, so you expand your knowledge. This is what AGM *revision* sanctions for the example. Another way to look at it, is to perform a case analysis over what you know. There are just two possibilities that are consistent with your knowledge: either (1) I have coffee and toast, or (2) I have neither coffee nor toast. Suppose (1). Finding that I had a coffee is perfectly reasonable with this case. Let us say that the outcome of case (1) is the scenario described by case (1) itself. Now suppose (2). Finding that I actually had a coffee conflicts with the case. You are obliged to jump to the "closest" scenario to case (2) that accommodates the information. For instance, it could be that I woke up late and left having no breakfast; but at the bus stop I bought just a coffee from a vendor. From this case analysis you conclude that definitely I had a coffee but that nothing can be said about me having toast. This is the type of change dictated by update.

They showed that the two operations have indeed different properties and, since then, AGM revision and update have been considered essentially different forms of theory change. The nature of their difference, though, is still an open question in the philosophical logic literature concerning theory change. For instance, are there other fundamentally different operations besides revision and update? A first formal difference between the two operations is that they do not stand on the same definitional ground. Katsuno and Mendelzon formalized their update operation as a binary connective between *formulae* in a logically finite language —specifically, a propositional language over a finite set of propositional variables—, Alchourrón, Gärdenfors and Makinson considered the general case of a possible infinite language and their revision operator takes *a theory and a formula* to the corresponding revised theory.

A number of formal comparisons between the two approaches have been already investigated. In particular Peppas and Williams [1995] have reformulated the update operation as a function over theories and extended Katsuno and Mendelzon's set of postulats so that an update operator may be used on first order languages. Implicitly, their article claims that Katsuno and Mendelzon's original postulates would be complete for general propositional languages, but not for first order. [Peppas and Williams, 1995], page 120:

> "Grove [5] used a syntactic representation based on maximal consistent extensions, or equivalently consistent complete theories, without the restriction of [6]. Katsuno and Mendelzon note in [6] that due to the one-to-one correspondence between consistent complete theories and interpretations in the finitary propositional case, their representation result is derivable from the work of Grove [5]. Furthermore, the one-to-one correspondence between consistent complete theories and interpretations does not require the finiteness property, and therefore, in the propositional case Grove's results have a semantic counterpart. However, this one-to-one correspondence does not hold for the more general first order case, and a model-theoretic characterization for this case has not hitherto been established. ... Katsuno and Mendelzon [7] introduced a set of postulates for an update operator on finitary propositional theories We extend their set of postulats so that an update operator may be used on arbitrary first order theories."

Their reference [5] stands [Grove, 1988], [6] for [Katsuno and Mendelzon, 1991] and [7] for [Katsuno and Mendelzon, 1992].

In this chapter we will give a clarification of Peppas and Williams' result. After presenting briefly the standard update operation in Section 5.1, we will define updates for infinite languages. Then we will prove an unexpected result: Katsuno and Mendelzon's original postulates characterizing finite updates are not sufficient for the infinite propositional case. In Section 5.4 we propose a strengthening of postulate (U8) which enables a representation theorem to be proved, obtaining the same postulate proposed in [Peppas and Williams, 1995]. Finally we will evaluate the update function against the general properties we studied in Chapter 3.

5.1 The Update Operation

Katsuno and Mendelzon define updates only for a classical propositional language based on a finite set of propositional variables P. This simplifying assumption has strong consequences as the set W of all possible valuations becomes finite. Two main properties result: every theory can be finitely axiomatized by a propositional formula; and every total order $\prec$ on W is free of infinite descending chains. These two properties let Katsuno and Mendelzon provide a simple definition of the update operator as a binary connective $\diamond$ in the propositional language: $A \diamond B$ is a well formed formula denoting the result of updating the theory $\mathrm{Cn}(A)$ with the formula B. The $\diamond$ operator is characterized through the following postulates:

(u1) $A \diamond B$ implies B.

(u2) If A implies B then $A \diamond B$ is equivalent to A.

(u3) If both A and B are satisfiable then $A \diamond B$ is also satisfiable.

(u4) If A_1 is equivalent to A_2 and B_1 is equivalent to B_2 then $A_1 \diamond B_1$ is equivalent to $A_2 \diamond B_2$.

(u5) $(A \diamond B) \wedge C$ implies $A \diamond (B \wedge C)$.

(u6) If $A \diamond B_1$ implies B_2 and $A \diamond B_2$ implies B_1 then $A \diamond B_1$ is equivalent to $A \diamond B_2$.

(u7) If $\mathrm{Cn}(A)$ is complete then $(A \diamond B_1) \wedge (A \diamond B_2)$ implies $A \diamond (B_1 \vee B_2)$.

(u8) $(A_1 \vee A_2) \diamond B$ is equivalent to $(A_1 \diamond B) \vee (A_2 \diamond B)$.

They furthermore consider an additional postulate:

(u9) If $\mathrm{Cn}(A)$ is complete and $(A \diamond B) \wedge C$ is satisfiable then $A \diamond (B \wedge C)$ implies $(A \diamond B) \wedge C$.

Katsuno and Mendelzon provide also a semantic characterization of the update operation through a notion of closeness between possible worlds. They consider a set of partial preorders on W, $\{\preceq_w : w \in W\}$. The intuitive meaning is that $v \preceq_w u$ if and only if v is at least as close to w as u is. The indexical preorders $\preceq_w$ are then used in the definition of the update operation: given that any theory K can be semantically represented as a set of possible worlds $[K] = \{w_i \in W : K \subseteq w_i\}$, we can update K by considering the most plausible changes (according to $\preceq_{w_i}$) to each w_i to accommodate the new information. The only requirement on $\preceq_w$ is a *centering condition*, establishing that for every w, no world is as close to w as w itself: if $v \preceq_w w$ then $v = w$.

The following characterization results hold for the update operation, see [Katsuno and Mendelzon, 1992] for the details.

Theorem 5.1 Let L be a finite propositional language. The update operator $\Diamond$ satisfies (u1)-(u8) iff there exists a model $\langle W, \{\preceq_w : w \in W\}\rangle$ where each $\preceq_w$ is a partial preorder over W that satisfies the centering condition, such that $[A\Diamond B] = \bigcup_{w\in[A]}\{v \in [B] : v \text{ is } \preceq_w\text{-minimal in } [B]\}$.

Theorem 5.2 Let L be a finite propositional language. The update operator $\Diamond$ satisfies (u1) - (u5),(u8) and (u9) iff there exists a model $\langle W, \{\preceq_w : w \in W\}\rangle$ where each $\preceq_w$ is a total preorder that satisfies the centering condition, such that $[A\Diamond B] = \bigcup_{w\in[A]}\{v \in [B] : v \text{ is } \preceq_w\text{-minimal in } [B]\}$. (Postulates (u6) and (u7) are superfluous in presence of the rest.)

5.2 Update for infinite languages

Following the notion of change advocated by Alchourrón, Gärden-fors and Makinson [Alchourrón *et al.*, 1985], we generalize the update operation to theories. We redefine the update operator $\Diamond$ as a function that takes a theory and a formula and returns a theory, $\Diamond : I\!\!K \times L \rightarrow I\!\!K$. Notice that in a finite propositional language this is just a notational variant of Katsuno and Mendelzon's original setting. We can straightforwardly recast the postulates governing the update function for possibly infinite languages as follows:

(U0) $K\Diamond A$ is a theory.

(U1) $A \in K\Diamond A$.

(U2) If $A \in K$ then $K\Diamond A = K$.

(U3) If $K \neq L$ and A is satisfiable then $K\Diamond A \neq L$.

(U4) If $\mathrm{Cn}(A) = \mathrm{Cn}(B)$ then $K\diamond A = K\diamond B$.

(U5) $K\diamond(A \wedge B) \subseteq \mathrm{Cn}(K\diamond A \cup \{B\})$.

(U6) If $B \in K\diamond A$ and $A \in K\diamond B$ then $K\diamond A = K\diamond B$.

(U7) If K is a complete theory then $K\diamond(A \vee B) \subseteq \mathrm{Cn}(K\diamond A \cup K\diamond B)$.

(U8) $(K_1 \cap K_2)\diamond A = (K_1\diamond A) \cap (K_2\diamond A)$.

The additional postulate becomes:

(U9) If K is complete and $\mathrm{Cn}((K\diamond A) \cup \{B\}) \neq L$ then $\mathrm{Cn}((K\diamond A) \cup \{B\}) \subseteq K\diamond(A \wedge B)$.

It is quite straightforward to extend the characteristic pointwise semantics of the standard update function to infinite languages. The notion of closeness between worlds requires some adjustment. In addition to the centering condition, each $\preceq_w$ should satisfy the limit assumption: let A be any formula in L, then there exists some non-empty set Y, $Y \subseteq [A]$ such that each element in Y is a $\preceq_w$-minimal element of $[A]$. Formally,

$$\forall w \in W, \forall A \in L, \exists Y \subseteq [A], Y \neq \emptyset \text{ such that } \forall y \in Y, \forall x \in [A], y \preceq_w x.$$

Notice that the limit assumption is trivially satisfied in finite propositional languages.

Definition 5.3 (Update function) Let L be a possibly infinite propositional language. Let $\langle W, \{\preceq_w : w \in W\}\rangle$ be such that each $\preceq_w$ is a total preorder over W satisfying the centering condition and the limit assumption. We define $\blacklozenge : \mathbb{K} \times L \rightarrow \mathbb{K}$ as

$$K\blacklozenge A = \mathrm{Th}(\bigcup_{w\in[K]} \{v \in [A] : v \text{ is } \preceq_w\text{-minimal in } [A]\}).$$

5.3 A Non-Representation Theorem

The generalized version of the update postulates (U0)-(U9) does not characterize the update operation in a language with an infinite number of propositional letters.

Theorem 5.4 If L is an infinite propositional language, postulates (U0)-(U9) do not fully characterize the $\blacklozenge$ operation.

PROOF. Given a propositional language L with an infinite but countable number of propositional letters we will exhibit a function $\circ : I\!K \times L \rightarrow I\!K$ satisfying (U0)-(U9) for which there is no model $\langle W, \{\preceq_w : w \in W\}\rangle$, satisfying that $\forall K \in I\!K$, $\forall A \in L, K \circ A = K \blacklozenge A$.

We semantically define $\circ$ as follows. Let us single out an (arbitrary) point v in W. For every $K \in I\!K$ and for every $A \in L$ define

$$[K \circ A] = \begin{cases} \emptyset & \text{if } [A] = \emptyset. \\ [K] & \text{if } [K] \subseteq [A]. \\ ([K] \cap [A]) \cup \{v\} & \text{if } A \in v \text{ and } [K] \cap [\neg A] \neq \emptyset \text{ is finite.} \\ [A] & \text{if } A \notin v \text{ or } [K] \cap [\neg A] \text{ is an infinite set.} \end{cases}$$

We first check that $\circ$ satisfies postulates (U0)-(U9). By definition $\circ$ trivially satisfies postulates (U0), (U1), (U2), (U3) and (U4).

(U5). We have to show that $K \circ (A \wedge B) \subseteq \text{Cn}(K \circ A \cup \{B\})$ holds. There are three cases.

(a) If $[K] \subseteq [A]$ then $K \circ A = K$. If $\neg B \in K$, then $\text{Cn}(K \circ A \cup \{B\}) = L$ and (U5) is verified. If $\neg B \notin K$, then $\text{Cn}(K \circ A \cup \{B\}) = \text{Cn}(K \cup \{B\})$. Since $A \in K$, $\text{Cn}(K \cup \{B\}) = \text{Cn}(K \cup \{A\} \cup \{B\}) = \text{Cn}(K \cup \{A \wedge B\}) = K \circ (A \wedge B)$. Thus, (U5) holds.

(b) Assume $[K] \cap [\neg A] \neq \emptyset$ is a finite set. If $[K] \cap [\neg A \vee \neg B]$ is an infinite set or $A \wedge B \notin v$ then $K \circ (A \wedge B) = \text{Cn}(A \wedge B)$ and (U5) holds. Suppose $[K] \cap [\neg A \vee \neg B]$ is finite and $A \wedge B \in v$. So $[K \circ (A \wedge B)] = ([K] \cap [A \wedge B]) \cup \{v\}$, while $[K \circ A] = ([K] \cap [A]) \cup \{v\}$. Since $B \in v$, $[K \circ A] \cap [B] = ((([K] \cap [A]) \cup \{v\}) \cap [B]) = ([K] \cap [A] \cap [B]) \cup (\{v\} \cap [B]) = ([K] \cap [A] \cap [B]) \cup \{v\} = [K \circ (A \wedge B)]$, thus (U5) is verified.

(c) If $[K] \cap [\neg A]$ is an infinite set then $[K] \cap [\neg A \vee \neg B]$ is also infinite. By definition $[K \circ (A \wedge B)] = [A \wedge B] = \text{Cn}([A] \cup [B]) = \text{Cn}(K \circ A \cup \{B\})$.

(U6). Suppose $B \in K \circ A$ and $A \in K \circ B$.

(a) If $[K] \subseteq [A]$ then $K \circ A = K$. Since $B \in K \circ A$, then $B \in K$, so $K \circ B = K = K \circ A$.

(b) Assume $[K] \cap [\neg A] \neq \emptyset$ is a finite set. If $A \in v$ then $[K \circ A] = ([K] \cap [A]) \cup \{v\}$. Since $B \in K \circ A$, then $([K] \cap [A]) \cup \{v\} \subseteq [B]$, and in particular, $B \in v$. Furthermore $[K] \cap [\neg B] \neq \emptyset$ is finite. Then, by definition, $[K \circ B] = ([K] \cap [B]) \cup \{v\}$. Since, in addition, $A \in K \circ B$, we obtain that $([K] \cap [B]) \cup \{v\} \subseteq [A]$. Therefore, $[K] \cap [A] = [K] \cap [B]$ and hence under the conditions in (b), $K \circ A = K \circ B$. Now suppose $A \notin v$. Then $[K \circ A] = [A]$. Since $B \in K \circ A$, $[A] \subseteq [B]$. As $A \in K \circ B$, $[K \circ B] \subseteq [A]$. Hence $[K \circ B] \neq ([K] \cap [B]) \cup \{v\}$, because we assumed $A \notin v$. Hence, it must be that $[K \circ B] = [B]$, so $[B] \subseteq [A]$. Therefore, $[A] = [B]$ and $K \circ A = K \circ B$.

(c) Assume $[K] \cap [\neg A]$ is an infinite set. Then, $[K \circ A] = [A]$. Since $B \in K \circ A$, then $[A] \subseteq [B]$. There are two possibilities for $K \circ B$. If $[K \circ B] = [B]$ then, using that $A \in K \circ B$, we obtain $[B] \subseteq [A]$ and $[K \circ A] = [K \circ B]$. If $[K \circ B] = ([K] \cap [B]) \cup \{v\}$ then $B \in v$ and $[K] \cap [\neg B]$ is a finite set. Because $A \in K \circ B$, $([K] \cap [B]) \cup \{v\} \subseteq [A]$, and $[K] \cap [B] \subseteq [K] \cap [A]$. Then, $[K] \cap [\neg A] \subseteq [K] \cap [\neg B]$; but this is impossible because we assumed $[K] \cap [\neg A]$ to be an infinite set and $[K] \cap [\neg B]$ to be finite.

(U7). We want to prove that if K is a complete theory then $K \circ (A \vee B) \subseteq \mathrm{Cn}(K \circ A \cup K \circ B)$. Assume K is complete.

If $A \in K$, $K \circ A = K$ and $K \circ (A \vee B) = K$. Thus, (U7) holds. If $\neg A \in K$, and $B \in K$, then $K \circ (A \vee B) = K \circ B = K$, so (U7) holds. If $\neg A \in K$, and $\neg B \in K$, if $A \in v$ or $B \in v$, then $K \circ (A \vee B) = v$, and either $K \circ B = v$ or $K \circ A = v$, so (U7) holds. If $\neg A \in v$ and $\neg B \in v$, then we obtain that $K \circ (A \vee B) = \mathrm{Cn}(A \vee B)$, $K \circ B = \mathrm{Cn}(B)$ and $K \circ A = \mathrm{Cn}(A)$. Hence, (U7) is verified.

(U8). We show that $(K_1 \cap K_2) \circ A = (K_1 \circ A) \cap (K_2 \circ A)$. Let $K = K_1 \cap K_2$.

(a) Assume $A \in K$. Then $K_1 \circ A = K_1$, $K_2 \circ A = K_2$ and $K \circ A = K$. Therefore (U8) is validated.

(b) Assume $[K] \cap [\neg A]$ is a finite non-empty set and $A \in v$. Then, $[K \circ A] = ([K] \cap [A]) \cup \{v\}$. If each $[K_i] \cap [\neg A]$, for $i = 1, 2$, is a finite set then $[K_i \circ A] = ([K_i] \cap [A]) \cup \{v\}$, $i = 1, 2$. So $[K \circ A] = ([K_1] \cap [A]) \cup ([K_2] \cap [A]) \cup \{v\} = [K_1 \circ A] \cup [K_2 \circ A]$. Otherwise, suppose $[K_1] \cap [\neg A]$ is an infinite set, and say $A \in K_2$. Then it also holds that $[K_1 \circ A] \cup [K_2 \circ A] = (([K_1] \cap [A]) \cup \{v\}) \cup [K_2] = (([K_1] \cap [A]) \cup \{v\}) \cup ([K_2] \cap [A]) = ([K_1] \cap [A]) \cup \{v\} \cup ([K_2] \cap [A]) = (([K_1] \cup [K_2]) \cap [A]) \cup \{v\} = ([K] \cap [A]) \cup \{v\} = [K \circ A]$.

(c) Assume $[K] \cap [\neg A]$ is an infinite set or $\neg A \in v$. If $\neg A \in v$ then $K \circ A = K_1 \circ A = K_2 \circ A = \text{Cn}(A)$, therefore, (U8) holds. Otherwise, either $[K_1] \cap [\neg A]$ or $[K_2] \cap [\neg A]$ or both are infinite sets. Clearly $[K \circ A] = [A]$ and, say, $[K_1] = [A]$. So $[K \circ A] = [K_1 \circ A]$, therefore, independently of the value of $[K_2 \circ A]$, we obtain that $[K \circ A] = [K_1 \circ A] \cup [K_2 \circ A]$.

(U9). Assume that K is complete and $[K \circ A] \cap [B] \neq \emptyset$. We prove that $[K \circ (A \wedge B)] \subseteq [K \circ A] \cap [B]$.

(a) If $A \in K$, $K \circ A = K$, by the hypotheses, $B \in K$. So $K \circ (A \wedge B) = K$. Thus, (U9) is verified.

(b) If $A \notin K$, then since K is complete $\neg A \in K$. If $A \in v$, $K \circ A = v$. By the hypothesis that $[K \circ A] \cap [B] \neq \emptyset$ we conclude $B \in v$. Thus, $[K \circ (A \wedge B)] \subseteq [K \circ A] \cap [B]$. In fact, $[K \circ (A \wedge B)] = [K \circ A] \cap [B] = \{v\}$. If $A \notin v$, $[K \circ A] = [A]$ and $[K \circ (A \wedge B)] = [A \wedge B]$. Thus, $[K \circ A] \cap [B] = [K \circ (A \wedge B)]$, hence (U9) is verified.

Now suppose for contradiction that there is a model $M = \langle W, \{\preceq_w : w \in W\}\rangle$, where each $\preceq_w$ is a total preorder on W satisfying the limit assumption and the centering condition, such that $\forall K \in \mathbb{K}$, $\forall A \in L, K \circ A = K \blacklozenge A$. Thus, for every theory K such that $[K]$ is a finite set, and for every formula A, if $\neg A \in K$ and $A \in v$, where v is the distinguished point appearing in the definition of $\circ$ above, $K \circ A = K \blacklozenge A = v$ must hold. This translates into the following condition on the model M.

$$\forall x \in [\neg A], \forall y \in [A], v \neq y, \quad v \prec_x y.$$

Now let K be a theory such that $[K]$ is an infinite set and let $A \in L$ be such that $A \in v$ and $\neg A \in K$. Then by definition of $\circ$, $[K \circ A] = [A]$. However, $[K \blacklozenge A] = \bigcup_{x \in [K]} \{y \in [A] : y \text{ is } \preceq_x\text{-minimal in } [A]\} = \{v\}$. Because the language is infinite $\{v\} \neq [A]$. QED

5.4 A Representation Theorem

In the previous section we proved that postulates (U0)-(U9) are insufficient to characterize the update operation in an infinite language. We propose the following postulate as a strengthening of Katsuno and Mendelzon's postulate (U8) to achieve the representation result.

(IU8) If $K = \bigcap H_i$ then $K \Diamond A = \bigcap (H_i \Diamond A)$.

(IU8) states that the update of an intersection is the intersection of the updates. Obviously (IU8) implies (U8). We now prove that postulates (U0)-(U9) plus (IU8) completely characterize the update operation when infinite languages are allowed. We will make use of the following lemma.

Lemma 5.5 Let $\preceq_w$ a preorder satisfying the limit assumption, an let X, Y be L-nameable subsets of W. If $\min_{\preceq_w}(X) \subseteq Y$ then $\min_{\preceq_w}(X \cap Y) = \min_{\preceq_w}(X)$.

PROOF. Assume $\min_{\preceq_w}(X) \subseteq Y$.

($\subseteq$). Suppose $v \in \min_{\preceq_w}(X \cap Y)$ but $v \notin \min_{\preceq_w}(X)$. Then for every $u \in X \cap Y$, $v \preceq_w u$, but there is some $z \in X$ such that $z \prec_w v$. By the assumption that $\min_{\preceq_w}(X) \subseteq Y$, $z \in Y$, so $z \in X \cap Y$. By transitivity of $\prec_w$, $z \prec_w v$, contradicting the minimality of v in $X \cap Y$.

($\supseteq$). Suppose $v \in \min_{\preceq_w}(X)$ but $v \notin \min_{\preceq_w}(X \cap Y)$.

Then there is some $z \in X \cap Y$ such that for all $u \in X \cap Y$, $z \prec_w u$.

By the assumption that $\min_{\preceq_w}(X) \subseteq Y$, $v \in Y$, moreover $v \in X \cap Y$. So $z \prec_w v$, contradicting the minimality of v in X. QED

Theorem 5.6 Let L be a possibly infinite propositional language, and let Cn be a classical consequence relation that is compact and satisfies the rule of introduction of disjunctions into the premisses. An operator $\diamondsuit$ satisfies postulates (U0)-(U7), (IU8), (U9) if and only if there exists a model $M = \langle W, \{\preceq_w : w \in W\}\rangle$, where each $\preceq_w$ is a total preorder over W centered in w that satisfies the limit assumption and for any $K \in \mathbb{K}$, $A \in L$, $K \diamondsuit A = K \blacklozenge A$.

PROOF.

[$\Leftarrow$]. We have to show that the operator $\blacklozenge$ satisfies postulates (U0)-(U7), (IU8) and (U9).

(U0) and (U1). Granted since, by Definition 5.3, $[K \blacklozenge A] \subseteq [A]$.

(U2). Follows as a consequence of the centering condition.

(U3). Follows by the definition of min on nonempty sets.

(U4). Obvious from the semantic definition of the update operation.

(U5). We have to show that $[K \blacklozenge A] \cap [B] \subseteq [K \blacklozenge (A \wedge B)]$. If $[K \blacklozenge A] \cap [B] = \emptyset$, the inclusion trivially holds. Assume $[K \blacklozenge A] \cap [B] \neq \emptyset$. Let u be any in $[K \blacklozenge A] \cap [B]$. Then $u \in \bigcup_{w \in [K]} \{v \in [A] : v \text{ is } \preceq_w\text{-minimal in } [A]\} \cap [B]$ $= \bigcup_{w \in [K]} \{v \in [A] \cap [B] : v \text{ is } \preceq_w\text{-minimal in } [A]\}$. Let $w_0 \in [K]$ be such that u is $\preceq_{w_0}$-minimal in $[A]$. That is $\forall v \in [A]$, $u \preceq_{w_0} v$. A fortiori, $u \in [A] \cap [B]$. Thus, there is no $v \in [A] \cap [B]$ such that $v \prec_{w_0} u$, so u is indeed $\preceq_w$-minimal in $[A] \cap [B]$.

(U6). Assume $B \in K \blacklozenge A$ and $A \in K \blacklozenge B$. We want to show $[K \blacklozenge A] = [K \blacklozenge B]$. $[K \blacklozenge A] = \bigcup_{w \in [K]} \{v \in [A] : v \text{ is } \preceq_w\text{-minimal in } [A]\}$. By the hypothesis that $B \in K \blacklozenge A$, $[K \blacklozenge A] \subseteq [B]$. So, $[K \blacklozenge A] = \bigcup_{w \in [K]} \{v \in [A] \cap [B] : v \text{ is } \preceq_w \text{-minimal in } [A] \cap [B]\}$. Similarly, $[K \blacklozenge B] = \bigcup_{w \in [K]} \{v \in [B] : v \text{ is } \preceq_w\text{-minimal in } [B]\}$, and by the hypothesis that $A \in K \blacklozenge B$, $[K \blacklozenge B] \subseteq [A]$. Hence, $[K \blacklozenge B] = \bigcup_{w \in [K]} \{v \in [A] \cap [B] : v \text{ is } \preceq_w\text{-minimal in } [A] \cap [B]\}$. Therefore, $[K \blacklozenge A] = [K \blacklozenge B]$, as required.

(U7). We have to prove that when $[K]$ is a singleton $[K \blacklozenge A] \cap [K \blacklozenge B] \subseteq [K \blacklozenge (A \vee B)]$. Assume $[K] = \{u\}$. Then, $[K \blacklozenge A] = \{v \in [A] : v$ is $\preceq_u$-

minimal in $[A]\}$, while $[K\blacklozenge B] = \{v \in [B] : v$ is $\preceq_u$-minimal in $[B]\}$. Furthermore $[K\blacklozenge(A \vee B)] = \{v \in [A \vee B] : v$ is $\preceq_u$-minimal in $[A \vee B]\} = \{v \in [A] \cup [B] : v$ is $\preceq_u$-minimal in $[A] \cup [B]\} = \{v \in [A] \cup [B] : v$ is $\preceq_u$-minimal in $[A]$ or v is $\preceq_u$-minimal in $[B]\}$. And finally, $[K\blacklozenge A] \cap [K\blacklozenge B] = \{v \in [A] \cap [B] : v$ is $\preceq_u$ -minimal in $[A]$ and v is $\preceq_u$ -minimal in $[B]\}$. Thus, $[K\blacklozenge A] \cap [K\blacklozenge B] \subseteq [K\blacklozenge(A \vee B)]$.

(IU8). Assume $[K] = \bigcup_{i\in I}[K_i]$ to show $[K\blacklozenge A] = \bigcup_{i\in I}[K_i\blacklozenge A]$. By definition, $[K\blacklozenge A] = \bigcup_{w\in\bigcup_{i\in I}[K_i]}\{v \in [A] : v$ is $\preceq_w$-minimal in $[A]\} = \bigcup_{i\in I}(\bigcup_{w\in[K_i]}\{v \in [A] : v$ is $\preceq_w$ -minimal in $[A]\}) = \bigcup_{i\in I}[K_i\blacklozenge A]$.

(U9). Assume $[K] = \{u\}$ and $([K\blacklozenge A]) \cap [B] \neq \emptyset$. We have to show $[K\blacklozenge(A \wedge B)] \subseteq [K\blacklozenge A] \cap [B]$. Suppose there is some $y \in [K\blacklozenge A \wedge B]$ but $y \notin [K\blacklozenge A] \cap [B]$. Then $[K\blacklozenge A] \subseteq [\neg B]$, contradicting $[K\blacklozenge A] \cap [B] \neq \emptyset$.

$[\Rightarrow]$. Let $\Diamond$ be a change function satisfying (U0)-(U7), (IU8) and (U9). We will construct a model $M = \langle W, \{\preceq_w : w \in W\}\rangle$ such that for every theory $K \in I\!K$ and formula $A \in L$, $K\Diamond A = K\blacklozenge A$.

We start by defining the model M. The domain W will be the set of all complete consistent theories in the language L. Assume $\{\preceq_w : w \in W\}$ is any set of total preorders satisfying the centering condition, the limit assumption and the following condition:

(i.) $v \preceq_w u$ iff there exists $A \in v \cap u$ such that $v \in [w\Diamond A]$ or there exists no satisfiable A such that $u \in [w\Diamond A]$.

In the limiting case when K is the inconsistent theory or A is unsatisfiable, $K\Diamond A$ and $K\blacklozenge A$ agree. We will now prove, for K and A satisfiable, that $u \in [K\Diamond A]$ iff $u \in [K\blacklozenge A]$ by analyzing the different cases.

Suppose $[K] = \{w\}$.

$[K\Diamond A] \subseteq [K\blacklozenge A]$. Let $v \in [K\Diamond A]$. By postulate (U1), $[K\Diamond A] \subseteq [A]$, so $v \in [A]$. By (i.), for every $u \in [A]$, $v \preceq_w u$. Hence, $v \in \{y \in [A] : y$ is $\preceq_w$-minimal in $[A]\} = [K\blacklozenge A]$.

$[K\blacklozenge A] \subseteq [K\Diamond A]$. Let $v \in [K\blacklozenge A]$. By definition of $\blacklozenge$, $v \in \{y \in [A]$: y is $\preceq_w$-minimal in $[A]\}$. So for all $u \in [A]$, $v \preceq_w u$; thus, by (i.), $v \in [w\Diamond A]$.

The general case, $[K] > 1$.

$[K\Diamond A] \subseteq [K\blacklozenge A]$. Let $v \in [K\Diamond A]$. By postulate (IU8), if $[K] = \bigcup_{i\in I}[K_i]$ then $[K\Diamond A] = \bigcup_{i\in I}[K_i\Diamond A]$.

In particular, $[K] = \bigcup_{i\in I}[T_i]$ for complete theories T_i. Thus, $v \in \bigcup_{i\in I}[T_i\Diamond A]$. Hence, v must be in, say, some $[T_j\Diamond A]$, $j \in I$. Then, by the previous case, $v \in [T_j\blacklozenge A]$. Therefore, $v \in \bigcup_{w\in[K]}\{y \in [A] : y$ is $\preceq_w$ -minimal in $[A]\}= [K\blacklozenge A]$.

$[K\blacklozenge A] \subseteq [K\Diamond A]$. Let $v \in [K\blacklozenge A]$. Then, $v \in \bigcup_{w\in[K]}\{y \in [A]$: y is $\preceq_w$-minimal in $[A]\}$. In particular, there exists some $w \in [K]$ such that $v \in \{y \in [A] : y$ is $\preceq_w$ -minimal in $[A]\}$. By the previous case, $v \in [w\Diamond A]$. But $[K] = \bigcup_{i\in I}[T_i]$ for complete theories T_i, such that $w = T_j$, for some $j \in I$. By postulate (IU8) we obtain that when $[K] = \bigcup_{i\in I}[K_i]$, $[K\Diamond A] = [T_j\Diamond A] \cup (\bigcup_{i\in I, i\neq j}[T_i\Diamond A])$. Hence $v \in [K\Diamond A]$. QED

Katsuno and Mendelzon's characterization results based on partial orders as opposed to partial pre-orders also lift to the infinite case, replacing postulate (U8) with postulate (IU8).

5.5 Properties

Keller and Winslett's [1985] first insightful distinction about two fundamentally different operations has been taken as a fundamental one in theory change. In what sense are revision and update so fundamentally different operations? AGM expansions, revisions and contractions are also different operations from one another, but not fundamentally so. Revisions and contractions are interdefinable, and that expansions are a limiting case of revisions. Most importantly, all AGM functions can be understood in the same semantic framework.

However, revision and update have been given different types of semantics. Update has a characteristic pointwise semantics that appears in no representation of the AGM functions. Namely, AGM revision has been recast as some ordering over maximal non-implying sets, entrenchment orderings, plain systems of spheres and safe hierarchies; all single global orderings. In Chapter 6 we will show that it is possible to define an AGM operation based on the update semantic apparatus, our analytic revision function.

Among the formal distinctions between revision and update, we have observed that updates have been defined over propositional languages while revisions are for general languages. But, as we have seen, nothing crucial relies on this difference being possible to characterize the update operation for infinite languages. Another difference is that update is truly a binary function, while, as we already remarked, general AGM functions are unary. Moreover, there is a fundamental property that relates the update of two theories: the update function is monotone with respect to its first argument, the second held fixed. Monotony is a direct consequence of postulate (IU8) and the fact that $\Diamond$ always returns a theory, i.e., $K\Diamond A = \mathrm{Cn}(K\Diamond A)$.

Observation 5.7 $\Diamond$ satisfies Monotony.

PROOF. Assume $K \subseteq H$. Then $K = K \cap H$. By (IU8) $K\Diamond A = K\Diamond A \cap H\Diamond A \subseteq H\Diamond A$. QED

Let's concentrate now on the properties we presented in section 3.3. Since $\Diamond$ validates (U8) it also validates Weak Intersection and D-ventilation. Katsuno and Mendelzon have defined a notion of *erasure* associated to that of update, which is defined via the Harper identity. Namely, $K - A = K \cap K\Diamond A$. Using the Levi identity erasure and update are interdefinable; hence, $K\Diamond A = (K - \neg A) + A$, where $+$ as the standard expansion function. Our first observation indicates that the update operator is not permutable, since the update operation does not overcome the inconsistent theory.

Observation 5.8 $\Diamond$ is not not permutable.

PROOF. For K a consistent theory such that $\neg A$ in K, $K - \neg A = K \cap K\Diamond A$ is a consistent theory, such that $\neg A \notin K - \neg A$.

However, $(K + A)$ equals L, the inconsistent theory, and applying the Harper identity, $L - \neg A = L \cap L\Diamond\neg A = L$.

Therefore, $(K - \neg A) + A \neq (K + A) - \neg A$. QED

The update function does not validate Union, Weak Union, Elimination, nor Commutativity. For ease of presentation we will write again the the different properties, putting a generic function $\circ : \mathbb{K} \times L \to \mathbb{K}$.

(Union) $\text{Cn}(K_1 \cup K_2) \circ A = \text{Cn}((K_1 \circ A) \cup (K_2 \circ A))$.

(Weak Union) If $\neg A \in K_1 \cap K_2$, then $\text{Cn}(K_1 \cup K_2) \circ A = \text{Cn}((K_1 \circ A) \cup (K_2 \circ A))$.

(Elimination) $(K \circ A) \circ B = K \circ (A \wedge B)$.

(Commutativity) $(K \circ A) \circ B = (K \circ B) \circ A$.

Observation 5.9 There exist update functions violating each of Union, Weak Union, Commutativity and Elimination.

PROOF. Let L be a propositional language based on just two propositional letters A and B. Let $K = \text{Cn}(\neg A \wedge \neg B)$. Suppose $[A] = \{w_1, w_2\}$ and $[B] = \{w_1, w_3\}$

Let's first prove that the update function fails Union. Let $\diamondsuit$ be any one satisfying the following pairs of the respective centered preorder relations; let $w_1 \prec_{w_4} w_2$, $w_2 \prec_{w_3} w_1$ and $w_2 \prec_{w_4} w_1$.

Therefore, $[K] = \{w_4\}$ and $[K \diamondsuit A] = [K] = \{w_1\}$. But $[\text{Cn}(\neg A) \diamondsuit A] = \{w_1, w_2\}$ and $[\text{Cn}(\neg B) \diamondsuit A] = \{w_2\}$. Thus, $\diamondsuit$ does not satisfy Union since, $\text{Cn}(K_1 \cup K_2) \diamondsuit A \neq \text{Cn}((K_1 \diamondsuit A) \cup (K_2 \diamondsuit A))$.

The same example shows that $\diamondsuit$ does not satisfy Weak Union. To see that the function does not validate Commutativity, let $\diamondsuit$ be an update function such that $w_3 \prec_{w_2} w_4$, $w_1 \prec_{w_3} w_2$ and let $[K] = \{w_2\}$.

Then $[K \diamondsuit A] = \{w_2\}$, $[K \diamondsuit A \diamondsuit B] = \{w_3\}$, $[K \diamondsuit B] = \{w_3\}$ and $[K \diamondsuit B \diamondsuit A] = \{w1\}$. Thus, $K \diamondsuit A \diamondsuit B \neq K \diamondsuit B \diamondsuit A$. This example also shows that $K \diamondsuit A \diamondsuit B \neq K \diamondsuit (A \wedge B)$, since $[K \diamondsuit (A \wedge B)] = \{w_1\}$ and $[K \diamondsuit A \diamondsuit B] = \{w_3\}$. QED

Interestingly $\diamondsuit$ fails (Or-Left) but satisfies (Or-Right).

(Or-Right) If $D \in (K \circ A) \circ C$ and $D \in (K \circ B) \circ C$ then $D \in (K \circ (A \vee B)) \circ C$.

(Or-Left) If $D \in (K \circ (A \vee B)) \circ C$ then $D \in (K \circ A) \circ C$ or $D \in (K \circ B) \circ C$.

Observation 5.10 $\diamondsuit$ validates Or-Right but fails Or-Left.

PROOF. (Or-Right). Assume (1) $D \in (K\Diamond A)\Diamond C$ and (2) $D \in (K\Diamond B)\Diamond C$.

$[K\Diamond A \vee B] = \bigcup_{w\in[K]}\{v \in [A \vee B] : v \in \min_{\preceq_w}([A \vee B])\} =$

$\bigcup_{w\in[K]}\{v \in [A] \cup [B] : v \in \min_{\preceq_w}([A] \cup [B])\} =$

$\bigcup_{w\in[K]}\{v \in [A] : v \in \min_{\preceq_w}([A] \cup [B])\} \cup \bigcup_{w\in[K]}\{v \in [B] : v \in \min_{\preceq_w}([A] \cup [B])\}$.

Since $\bigcup_{w\in[K]}\{v \in [A] : v \in \min_{\preceq_w}([A] \cup [B])\} \subseteq [K\Diamond A]$, and $\bigcup_{w\in[K]}\{v \in [B] : v \in \min_{\preceq_w}([A] \cup [B])\} \subseteq [K\Diamond B]$, we obtain that $[K\Diamond A \vee B] \subseteq [K\Diamond A] \cup [K\Diamond B]$.

$[(K\Diamond A \vee B)\Diamond C] = \bigcup_{w\in[K\Diamond A\vee B]}\{v \in [C] : v \in \min_{\preceq_w}([C])\} \subseteq \bigcup_{w\in[K\Diamond A]}\{v \in [C] : v \in \min_{\preceq_w}([C])\} \cup \bigcup_{w\in[K\Diamond B]}\{v \in [C] : v \in \min_{\preceq_w}([C])\}$.

By (1) $\bigcup_{w\in[K\Diamond A]}\{v \in [C] : v \in \min_{\preceq_w}([C])\} \subseteq [D]$,

and by(2) $\bigcup_{w\in[K\Diamond B]}\{v \in [C] : v \in \min_{\preceq_w}([C])\} \subseteq [D]$.

Therefore, $[(K\Diamond A \vee B)\Diamond C] \subseteq [D]$.

(Or-Left). Let L be a propositional language based on four propositional letters A, B, C and D. Let $K = \text{Cn}(\neg A \wedge \neg B \wedge \neg C) = w_5 \cap w_6$. $w_2 = \text{Cn}(A\wedge\neg B\wedge C\wedge D)$, $w_9 = \text{Cn}(\neg A\wedge B\wedge C\wedge D)$, and $w_3 = \text{Cn}(A\wedge B\wedge C\wedge\neg D)$,

Let's prove that the update function fails Or-Left. Let $\Diamond$ be any one satisfying the following pairs of the respective centered preorder relations; let $w_2 \prec_{w_5} w_i$ for all $i \neq 5, 6$, $w_9 \prec_{w_6} w_i$ for all $i < 6$, and $w_3 \prec_{w_5} w_i$ for all $i \neq 2, 5$, $w_3 \prec_{w_6} w_i$ for all $i \neq 9, 6$.

Therefore, $[K] = \{w_5, w_6\}$ and $[K\Diamond(A \vee B)] = \{w_2, w_9\}$. Since $w_2, w_9 \in [C]$ and $w_2, w_9 \in [D]$ then $[K\Diamond(A \vee B)\Diamond C] = \{w_2, w_9\}$ and $D \in K\Diamond(A \vee B)\Diamond C$.

On the one hand $[K\Diamond A] = \{w_2, w_3\}$. Since $w_2, w_3 \in [C]$, $[K\Diamond B\Diamond C] = \{w_2, w_3\}$. Since $w_3 \notin [D]$, $D \notin K\Diamond A\Diamond C$.

On the other hand $[K\Diamond B] = \{w_3, w_9\}$. Since $w_3, w_9 \in [C]$, $[K\Diamond B\Diamond C] = \{w_3, w_9\}$. Since $w_3 \notin [D]$, $D \notin K\Diamond B\Diamond C$. QED

Let's turn our attention to Lehmann's postulates for iterated change. The postulates that deal with iteration are not validated by the update operation.

(I1) $K \circ A$ is a consistent theory.

(I2) $A \in K \circ A$.

(I3) If $B \in K \circ A$, then $A \supset B \in K$.

(I4) If $A \in K$ then $K \circ B_1 \circ \ldots \circ B_n = K \circ A \circ B_1 \circ \ldots \circ B_n$ for $n \geq 1$.

(I5) If $A \in \mathrm{Cn}(B)$, then $K \circ A \circ B \circ B_1 \circ \ldots \circ B_n = K \circ B \circ B_1 \circ \ldots \circ B_n$.

(I6) If $\neg B \notin K \circ A$ then $K \circ A \circ B \circ B_1 \circ \ldots \circ B_n = K \circ A \circ (A \wedge B) \circ B_1 \circ \ldots \circ B_n$.

(I7) $K \circ \neg B \circ B \subseteq \mathrm{Cn}(K \cup \{B\})$.

Observation 5.11

i) All update functions satisfy (I1), (I2), (I3), (I4).

ii) There exist update functions violating (I5), (I6) and (I7).

PROOF. The violation of (I5), (I6) and (I7) can be proved by constructing a counterexample.

(I1), (I2), (I3), (I4) follow from postulates (U0)-(U5). QED

Updates do not validate any of Darwiche and Pearl's postulates [1997].

(C1) If $A \in \mathrm{Cn}(B)$ then $(K \circ A) \circ B = K \circ B$.

(C2) If $\neg A \in \mathrm{Cn}(B)$ then $(K \circ A) \circ B = K \circ B$.

(C3) If $A \in K \circ B$ then $A \in (K \circ A) \circ B$.

(C4) If $\neg A \notin K \circ B$ then $\neg A \notin (K \circ A) \circ B$.

(C5) If $\neg B \in K \circ A$ and $A \notin K \circ B$ then $A \notin (K \circ A) \circ B$.

(C6) If $\neg B \in K \circ A$ and $\neg A \in K \circ B$ then $\neg A \in (K \circ A) \circ B$.

Observation 5.12 There exist update functions violating (C1)-(C6).

Finally, we observe that the update function validates none of the postulates for iterated change associated to iterative schemes.

(T) $K \circ A \circ B = K \circ B$.

(C) If $\neg B \in K \circ A$, then $K \circ A \circ B = K \circ B$.

(I) $K \circ A \circ B = K \circ (A \wedge B)$

(M)

$$K \circ A \circ B = \begin{cases} K \circ B & \text{, if } \neg B \in \mathrm{Cn}(A) \\ K \circ (A \wedge B) & \text{, otherwise.} \end{cases}$$

Observation 5.13 There exist update functions violating (T),(M),(I) and (C).

Chapter 6

Analytic AGM Functions

In this chapter we will resume the discussion we initiated in Chapter 5 about the difference between AGM revision and update and we will end up establishing a bridge between the two kinds of change. We will provide a new presentation of AGM revision based on the update semantic apparatus, a pointwise semantics for revision. Our strategy will be to define a new semantic operation as a variant of the update operation that we will dub *analytic revision*. The key idea is that for each theory K we will define a preorder relation obtained from the indexed relations of an update model. We will show that our analytic revision is indeed a binary AGM function, that is defined for every theory and every formula.

Theorem 6.15 is the main theorem of this chapter, and provides a characterization of analytic functions as those AGM functions satisfying (K*1)-(K*8) plus two new postulates, (K∗∃) and (K∗∀), governing the revision of different theories.

This study builds on our initial work in [Becher, 1995b]. In that paper our current analytic function was called a "lazy update" reflecting that it was semantically defined as a variant of the standard update operation. Lazy update were just defined for finite languages and we proved they satisfy all AGM revision postulates.

The independent work of Schlechta, Lehmann and Magidor's "Distance Semantics for Belief Revision" [Schlechta *et al.*, 1996] turned out to be related to ours. Notably, their revision function based on distances and our analytic revision function are definitionally equivalent, modulo some considerations over the formal structures they are based on. Our work extends and continues theirs in several respects. We consider an infinite language and we obtain characterization results for functions built over non symmetric distances —a question left open in [Schlechta *et al.*, 1996]. Also novel in our work is the definition of AGM revision in the update semantic structure, which allows us to connect these two seemingly incomparable forms of theory change.

6.1 Analytic revision functions

Our aim is to define the AGM revision function in the semantic framework of update. We start by noticing some important particularities of the update function which are not shared with revision. As we already remarked in Chapter 5 update is an "authentic" binary function, but general AGM revisions are not. Another difference is how they deal with the inconsistent theory. In the update setup the inconsistent theory is a sink from which the change function cannot escape. In contrast, the revision of even the inconsistent theory should be consistent as far as the new information is (by postulate (K*5)); i.e. revision can recover from inconsistency, update cannot overcome it. Finally, following the ideal of minimal change, a consistent revision always coincides with an expansion (by (K*3) and (K*4)), which does not hold for update. And a crucial consideration is that update possesses a pointwise semantics, that appears in no presentation of the AGM theory.

Consider a theory as a set of possible scenarios. Katsuno and Mendelzon's operation can be calculated by means of a case analysis over the set of complete scenarios compatible with the original theory. First, for each case find out its closest outcome that accommodates the new information; then take as the overall result what is common to all outcomes. Even though for each case the closest outcome entailing the new information is selected, some outcomes could be relatively implausible. Could we have a measure to determine when one outcome is more plausible than another? What is a sensible notion to compare outcomes? We suggest that one outcome is more plausible than another when it is at a *closer distance* from the theory under change. We will first formalize a notion of distance we will be concerned with, and then define a new operation that picks as a result of the change just the outcomes that are minimally distant. We will call this operation an *analytic revision.*

A *distance* is a binary function $f : X \times X \to Y$, such that X is a set and Y is a totally ordered set with minimal element, satisfying that $f(x, y) = \min(Y)$ iff $x = y$ (centering) and $f(x, y) = f(y, x)$ (symmetry). But there are weaker notions. Ultrametric distances satisfy the centering and the triangular inequality and pseudo distances just satisfy the centering condition. Since we seek the connection between revision and update, we are interested in a notion of distance that corresponds to the proerders of update models. Thus, we shall be concerned with pseudo distances only, and making some language abuse we will refer to them just as distances.

Assume L is a possibly infinite propositional language, and W is the set of all it maximal consistent sets. Let's recall Definition 5.3. An update model is a structure $M = \langle W, \{\preceq_w : w \in W\}\rangle$, such that each $\preceq_w$ is a total preorder over W satisfying the centering condition and the limit assumption.

It is possible to recast an update model into a model based on functions having as range any totally ordered set with smallest element. We will consider the set $\mathbb{R}^+$ of real numbers greater or equal to 0, but any other (not well founded) totally ordered set with smallest element would do. It is clear how each total preorder in the update model induces a function d_w such that all the information encoded in $\preceq_w$ is placed in $d_w : W \to \mathbb{R}$.

$$v \preceq_w u \text{ iff } d_w(v) \leq d_w(u).$$

The centering condition establishes a restriction on the possible values of the functions.

(centering) $d_w(w) = 0$ and for every $v \in W$ such that $v \neq w$, $d_w(v) > 0$.

For each indexical total preorder $\preceq_w$ the limit assumption requires that for each L-nameable set (by a single formula) $[A]$ there exists some $\preceq_w$-minimal elements of $[A]$.

(limit assumption) For each $x \in W$, for each $[A] \subseteq W$, there are $y \in [A]$ such that $\forall y' \in [A]$, $d_x(y) \leq d_x(y')$.

The update of K by A is defined as:

$$K \Diamond A = \mathrm{Th}\Big(\bigcup_{w \in [K]} \{v \in [A] : d_w(v) = d_w([A])\}\Big).$$

Just for convenience we give the following

Definition 6.1 (Distance between two points) Let an update model $M = \langle W, \{d_w : w \in W\}\rangle$ be given. We define the *distance* function $d : W \times W \longrightarrow I\!R$ between pairs of worlds v, w as the value of w in d_v: $d(v, w) = d_v(w)$.

Since functions d_w obey the centering condition, the distance from a point to itself is 0 and the distance from a point to every other point is greater than 0. We will require no further properties on d for the moment. Notice in particular that this conception of distance is not symmetric since $d(w, v)$ may differ from $d(v, w)$. Boutilier (personal communication) has provided a good rational for it: "The lack of symmetry seems certainly appropriate when the ordering mirrors exogenous change; for instance, it is quite easy to break an egg while it is hopeless to put it back together."

We shall extend the above definition to distance between sets, as the result of a double minimization. The definition of $d : \mathcal{P}(W) \times \mathcal{P}(W) \longrightarrow I\!R$ covers the limiting case of the empty proposition in a way that will be convenient.

Definition 6.2 (Distance from a set to a set) Let d be a distance function obtained from an update model $\langle W, \{d_w : w \in W\}\rangle$. Let X, Y be subsets of W. Let $f : W \longrightarrow I\!R$ be any positive (greater that 0) function. We define.

$$d(X, Y) = \begin{cases} \min_{x \in X} \min_{y \in Y} \{d(x, y)\} & \text{, if } X, Y \neq \emptyset. \\ \min_{y \in Y} \{f(y)\} & \text{, if } X = \emptyset, Y \neq \emptyset. \\ 0 & \text{, if } Y = \emptyset. \end{cases}$$

From now on we assume the extended distance function and, abusing notation, we will write singleton sets without braces, i.e. we will write $d(u, v)$ instead of $d(\{u\}, \{v\})$. As before, notice the lack of symmetry: in general $d(X, Y)$ is different from $d(Y, X)$. Furthermore we will directly consider models $M = \langle W, d\rangle$ instead of the indexical models as we can straightforwardly move from one to the other. We are ready now to give the formal semantic definition of analytic revision.

Definition 6.3 (Analytic revision) Let $M = \langle W, d\rangle$ be a model and $X, Y \subseteq W$, then the analytic revision $\bullet : \mathcal{P}(W) \times \mathcal{P}(W) \to \mathcal{P}(W)$ is defined as

$$X \bullet Y = \{y \in Y : d(X, y) = d(X, Y)\}.$$

The syntactic counterpart taking as arguments a theory and a formula, $\bar{\bullet} : I\!K \times L \to I\!K$ is simply

$$K \bar{\bullet} A = \mathrm{Th}([K] \bullet [A]).$$

6.2 Connections

6.2.1 Analytic revision and update

The crucial semantic difference between analytic revision and update is that analytic operation relies on two minimizations while the update just one. As a direct consequence an analytic revision ignores some of the possible outcomes that an update would consider. Then the theory resulting from an analytic revision is at least as informed as that of an update.

Observation 6.4 If K is consistent, $K \Diamond A \subseteq K \bar{\bullet} A$.

PROOF. We want to show that $X \bullet Y$ is included in $X \Diamond Y$. Suppose $y \in X \bullet Y$. Then $\min_{x \in X}\{d(x, y)\} = \min_{x \in X, y' \in Y}\{d(x, y')\}$. Fix a value x_0 of $x \in X$ such that $d(x_0, y) = \min_{x \in X}\{d(x, y)\}$. Then $d(x_0, y) = \min_{x \in X, y' \in Y}\{d(x, y'\}$. Hence $d(x_0, y) = \min_{y' \in Y}\{d(x_0, y'\}$. Hence $y \in X \Diamond Y$. QED

The reason for this observation being relative to the consistency of K is that the update function of the inconsistent theory results in the inconsistent theory. In contrast, analytic revision overcomes inconsistency. The following result asserts that when the theory is also complete the two operations coincide.

Observation 6.5 If K is consistent and complete then $K \bar{\bullet} A = K \Diamond A$.

PROOF. The proof is quite trivial. Let K be consistent and complete, so its proposition is a singleton $[K] = \{u\}$.

Then, $[K \Diamond A] = \bigcup_{i \in [K]}\{w \in [A] : d_i(w) = d_i([A])\} = \{w \in [A] : d_u(w) = d_u([A])\} = \{w \in [A] : d([K], w) = d([K], [A])\} = [K] \bullet [A]$. QED

We establish precisely the connection between analytic revision and update, generalizing the two results above.

Observation 6.6 Let K be a consistent theory and $\langle W, d\rangle$ a structure for the update operation $\diamond$. Then for every formula A there exists a consistent theory $K' \supseteq K$ such that $K \bar{\bullet} A = K' \diamond A$. In particular, K' may be chosen as $\mathrm{Th}(\{w \in [K] : d(w,[A]) = d([K],[A])\})$. (Notice that K' depends on A.)

PROOF. By Observation 6.4 we know that taking $K' = K$ provides us with a theory that is too weak to satisfy the observation. Let's study this in detail.

If A is a satisfiable formula, $[A] \neq \emptyset$, so $[K \diamond A]$ is not empty.

By definition $[K \diamond A] = \bigcup_{w \in [K]} \{v \in [A] : d_w(v) = d_w([A])\} = \bigcup_{w \in [K]} \{v \in [A] : d(w,v) = d(w,[A])\} =$

$\bigcup \{\{v \in [A] : d(w,v) = d(w,[A])\} : w \in [K] \text{ and } d(w,[A]) = d([K],[A])\}$

$\cup$

$\bigcup \{\{v \in [A] : d(w,v) = d(w,[A])\} : w \in [K] \text{ and } d(w,[A]) > d([K],[A])\}$.

Thus, $[K']$ should be chosen as $[K'] = \{w \in [K] : d(w,[A]) = d([K],[A])\}$ in which case $K' \diamond A = K \bar{\bullet} A$. QED

The next lemma states that when a formula is consistent with the theory, the analytic revision operation is just the addition of the formula to the theory.

Lemma 6.7 If A is consistent with K, then $K \bar{\bullet} A = \mathrm{Cn}(K \cup \{A\})$.

PROOF. Assume A is consistent with K. Then $[K] \cap [A] \neq \emptyset$. By the centering condition $d([K],[A]) = 0$ and for any $v \notin [K], d([K],v) > 0$. Then by Definition 6.3, $[K] \bullet [A] = \{w \in [A] : d([K],w) = 0\}$. Thus, $[K] \bullet [A] = [K] \cap [A]$. QED

In spite of the technical connection it is not surprising to find that the analytic revision is not an update operator.

Observation 6.8 $\bar{\bullet}$ satisfies (U0)-(U7) and (U9), fails (IU8) and fails monotony.

PROOF. Let's see first that $\bar{\bullet}$ satisfies (U1)-(U7) and (U9).

(U0) and (U1) are granted since by Definition 6.3, $[K] \bullet [A] \subseteq [A]$.

(U2) follows as a direct consequence of Lemma 6.7.

(U3) is a consequence of the limit assumption of d.

(U4) is obvious from the semantic definition of analytic revision.

(U5)[1] We have to show that $([K] \bullet [A]) \cap [B] \subseteq [K] \bullet [A \wedge B]$. If $([K] \bullet [A]) \cap [B] = \emptyset$, the inclusion trivially holds.

Assume $([K]\bullet[A])\cap[B] \neq \emptyset$. By Definition 6.3, $([K]\bullet[A])\cap[B] = \{w \in [A] : d([K], w) = d([K], [A])\} \cap [B] = \{w \in [A] \cap [B] : d([K], w) = d([K], [A])\}$. Also, $[K] \bullet [(A \wedge B)] = \{w \in [A] \cap [B] : d([K], w) = d([K], [A] \cap [B])\}$.

Suppose for contradiction that (1) $u \in ([K] \bullet [A]) \cap [B]$, and (2) $u \notin [K] \bullet [(A \wedge B)]$. From (1) we obtain (3) $u \in [A] \cap [B]$, while (2) can be rewritten as (2') $u \notin \{w \in [A] \cap [B] : d([K], w) = d([K], [A] \cap [B])\}$.

Then by (2') and (3) we obtain (4) $d([K], u) > d([K], [A] \cap [B])$. By (1) we have that $d([K], u) = d([K], [A])$, and (3) assures that $u \in [A] \cap [B]$. Hence we obtain $d([K], u) = d([K], [A] \cap [B])$, contradicting (4).

(U6) Assume $B \in K\bar{\bullet}A$ and $A \in K\bar{\bullet}B$. Since $d([K], [A]) = \min_{x\in[K]} \min_{y\in[A]}\{d(x, y)\}$, there exists $v \in [K] \bullet [A]$ such that $d([K], v) = d([K], [A])$. Similarly, there exists $w \in [K] \bullet [B]$ such that $d([K], w) = d([K], [B])$.

Since $[K] \bullet [A] \subseteq [B]$, then $d([K], [A]) = d([K], v) \geq d([K], [B])$. Also since $[K] \bullet [B] \subseteq [A]$ $d([K], [B]) = d([K], w) \geq d([K], [A])$. We obtain $d([K], [A]) \geq d([K], [B]) \geq d([K], [A])$, thus, $d([K], [A]) = d([K], [B])$. We conclude, $[K] \bullet [A] = [K] \bullet [B]$, as required.

(U7) Assume $[K] = \{u\}$, then distance from $[K]$ is exactly distance from u and $d(u, [A\vee B]) = d(u, [A]\cup[B]) = \min\{d(u, [A]), d(u, [B])\}$. Without loss of generality assume $d(u, [A]) \leq d(u, [B])$. Then $[K] \bullet [A \vee B] = [K] \bullet [A]$; hence, $([K] \bullet [A]) \cap ([K] \bullet [B]) \subseteq [K] \bullet [A \vee B]$.

(U9)[2] We have to show that if $[K]$ is a singleton and $([K]\bullet[A])\cap[B]$ is not empty then $([K]\bullet[A\wedge B]) \subseteq ([K]\bullet[A])\cap[B]$. Assume (1) $([K]\bullet[A])\cap[B] \neq$

[1]Notice that this postulate corresponds to the AGM revision postulate (K*7).

[2]Notice that this postulate is a particular case of the AGM revision postulate (K*8).

$\emptyset$. Then there is some $x \in [A] \cap [B]$ such that $d([K],[A]) = d([K],x)$.

Suppose (2) $[K] \bullet [A \wedge B] \not\subseteq ([K] \bullet [A]) \cap [B]$. Then there is some $u \in [K] \bullet [A \wedge B)])$ but $u \notin ([K] \bullet [A]) \cap [B]$. By (1) and (2) we obtain (3) $d([K],[A \wedge B]) = d([K],u) > d([K],[A])$. By Definition 6.3 and (3), for every $w \in [A]$, if $d([K],w) = ([K],[A])$ then $w \in [A] \cup [\neg B]$, contradicting (1). Notice for later use that for this proof we have not made use of the hypothesis that $[K]$ is a singleton.

To prove that $\bar{\bullet}$ fails postulate (IU8) suffices to to provide witnesses to $(X \cup Y) \bullet Z \neq (X \bullet Z) \cup (Y \bullet Z)$. Let $X, Y, Z \subseteq W$ non-empty, such that $X \cap Z = \emptyset$ and $Y \cap Z \neq \emptyset$. Hence $(X \cup Y) \cap Z = Y \cap Z \neq \emptyset$.

By Lemma 6.7,$Y \bullet Z = Y \cap Z$ and $(X \cup Y) \bullet Z = (X \cup Y) \cap Z = Y \cap Z$. Therefore, $(X \cup Y) \bullet Z = Y \bullet Z$. From postulate (U3) proved above, $X \bullet Z \neq \emptyset$.

Since $X \bullet Z$ may not be included in $Y \bullet Z$, (U8) may not be satisfied. For instance let $X = \{x\} \subseteq [A \wedge \neg B]$, $Y = \{y\} \subseteq [A \wedge B \wedge C]$, $Z = [B]$, and let $v \in [B \wedge \neg C]$. Let d_x, d_y satisfy the centering condition such that, $d_x(v) = 1$. Then, $v \in X \bullet Z$ and $Y \bullet Z = \{y\}$. Thus, $(X \cup Y) \bullet Z$ is different from $(X \bullet Z) \cup (Y \bullet Z)$.

That $\bar{\bullet}$ fails monotony can be proved using the same strategy of Observation 2.1. QED

The analytic revision operation relies only on those possible worlds that regard the change as minimally distant from the theory under change. Then, if possible, the analytic revision will understand new information as having caused no change at all, a mere confirmation of what already was a possibility in our picture of the world. This behaviour has been stated as Lemma 6.7 and is shared with AGM revision. In the next section we will show that AGM revisions and analytic revisions are indeed connected.

6.2.2 Analytic revision and AGM revision

First we will note that the analytic revision function $\bar{\bullet}$ satisfies the AGM postulates (K*1)-(K*8).

Theorem 6.9 $\bar{\circ}$ is a revision operator satisfying (K*1)-(K*8).

PROOF. Most postulates follow directly from Definition 6.3 or from Lemma 6.7. (K*7) and (K*8) have been proved as postulates (U5) and (U9) respectively, in Observation 6.8. QED

It is important to remark that the key idea behind an analytic revision is to define a meaningful distance relation between sets in terms of the functions d_w (which in turn were obtained from the ternary relations $\preceq_w$). For example, a candidate distance from a theory K could have been any arbitrary d_v. But it is evident that the change operation this approach would induce does not satisfy the complete set of AGM revision postulates.

Observation 6.10 Assume L a language with at least two propositional letters, K an incomplete theory of L, $v \in [K]$ a single element of W and d_v an real function for v satisfying the centering condition. Let $\circ$ be a change operation for K defined as $K \circ A = \mathrm{Th}(\{y \in A : d_v(y) = d_v([A])\})$. Then $\circ$ satisfies (K*1),(K*2),(K*5)-(K*8) but in general fails (K*3)(K*4).

PROOF. (K*1),(K*2),(K*5)-(K*8) have identical proofs as those in Theorem 6.9.

(K*3). Since we assume K is not complete then there is a formula A such that $A, \neg A \notin K$. Then, either $v \in [A]$ or $v \in [\neg A]$. Without loss of generality, suppose $v \in [\neg A]$. Then, there is some $x \in [K] \cap [A]$. We show a counterexample to (K*3) such that $x \notin [K \circ A]$. Since L has at least two propositional letters, there is some $u \in [A]$, $u \neq x$. Let $d_v(u) < d_v(x)$. Then, $x \notin [K \circ A] = \{y \in [A] : d_v(y) = d_v([A])\}$, as x is not a minimal element in d_v satisfying A.

If we add to the the previous counterexample that $u \notin [K]$ and $d_v(u) = d_v([A])$, then postulate (K*4) also fails as $u \in [K \circ A]$ but $u \notin [K] \cap [A]$. QED

Distance from theory K becomes the standard ordering used in the semantic presentations of AGM revision (a world w is as close as v from theory K if and only if the distance from $[K]$ to w is not greater than the distance from $[K]$ to v).

Theorem 6.9 shows that every analytic revision function is an AGM revision function. However, what is most interesting about analytic revision is that *a transitively relational AGM revision function is an analytic revision.* Only after this result we can speak of a true connection between AGM revision and the semantic structure of update.

Theorem 6.11 (Makinson, personal communication) Every revision function $*$ for K satisfying the extended set of AGM postulates (K*1)-(K*8) is an analytic revision function for K.

PROOF. Let $*$ be an AGM revision function for K satisfying (K*1)-(K*8). By Grove's result, there is a system of spheres S^K for K that represents $*$.

By Observation 2.14 S^K induces an real function d_K on W into the reals greater or equal 0, satisfying (centering) and (Limit Assumption).

The proof of the theorem just consists in showing that *any* real function $d : W \rightarrow \mathbb{R}^+$ satisfying (centering) and (Limit Assumption) can be extended to a distance function, obtaining the semantic structure of analytic revision. We define $d : W \times W \rightarrow \mathbb{R}^+$ as follows.

i. $\forall w, v \in [K], w \neq v$, $d(w, v) = 1$,

ii. $\forall w \in W$, $d(w, w) = 0$,

iii. $\forall w \in [K], \forall v \in W \setminus [K]$, $d(w, v) = d_K(v)$,

iv. $\forall w \in W \setminus [K], \forall v \in W, d(w, v) = g_w(v)$,

where $g_w : W \rightarrow \mathbb{R}^+$ is any function at all assigning values greater than 0. We extend d as a function on sets as usual, taking $d(\emptyset, v) = d_K(v)$, for the empty set. We have to check that the function d is of the kind needed to generate a analytic revision operation. We just check that the induced relations $\preceq_w$ over W defined by setting

$$u \preceq_w v \text{ iff } d(w, u) \leq d(w, v), \text{ for all } u, v, \in W$$

satisfy (1) $\preceq_w$ is a a total preorder on W, and (2) $\preceq_w$ is centered at w; i.e. if $v \preceq_w w$ then $v = w$.

Now (1) is immediate. To prove (2), let $u \in W$ with $u \neq w$. We want to show that $w \prec_w u$; i.e. that if $u \neq w$ then $d(w,w) < d(w,u)$. By the second case of our definition of d, $d(w,w) = 0$ for all $w \in W$, hence we have to show that for $w \neq u$, $d(w,u) > 0$. If u, w are both in $[K]$, then by the first case $d(w,u) = 1 > 0$. If w is not in $[K]$, it follows from the fourth case that $d(w,u) > 0$. If w in $[K]$ and u is not, then $d(w,u) = r(u) > 0$ since r is itself centered in $[K]$. Thus in all cases $d(w,u) > 0$ and we are done.

It is immediate from the definition of d that (3) for all $u, v \in W \setminus [K]$, for any $w \in [K]$, $d(w,v) \leq d(w,u)$ iff $d_K(v) \leq d_K(u)$ iff $c^K(v) \subseteq c^K(u)$, and (4) for any $u, w \in [K]$ and for all $v \in W \setminus [K]$, $d(w,v) = d(u,v) = d([K],v) = d_K(v)$.

Now let $\bullet$ be the analytic revision function determined by the structure $\langle W, d\rangle$. We have to show that for all A, $[K * A] = [K] \bullet [A]$. If $[K] \cap [A] \neq \emptyset$, by (K*4) in Lemma 6.7 we have $[K * A] = [K] \cap [A] = [K] \bullet [A]$. Suppose $[K] \cap [A] = \emptyset$. By definition of analytic revision and (4) $[K] \bullet [A] = \{v \in [A] : d([K],v) = d([K],[A])\} = \{v \in [A] : d_K(v)$ is $\leq$-minimal in $\{d_K(w) : w \in [A]\}\} = \{v \in [A] :$ v is in the $\subseteq$-minimal sphere in S that intersects $[A]\}$ = (by (3) above) $[K * A]$. QED

We observe in the proof above and also in Definition 6.2 that we have considerable freedom when defining the behaviour for the revision of the inconsistent theory. For example we could require what Makinson called the Overkilling property (O). It says that the analytic revision of an inconsistent theory should result in plain acceptance of the new information. Coincidentally, this property defines the revision of the inconsistent theory in [Schlechta *et al.*, 1996].

$$\text{(O) If } K \text{ is inconsistent then } K \bullet A = \mathrm{Cn}(A).$$

The analytic revision function that comply with (O) can be characterized by the function $f : W \to I\!R^+$ involved in the definition of d (see Definition 6.2).

Observation 6.12 (Makinson, personal communication) $\bullet$ satisfies (K*1)-(K*8) and (O) if and only if f is a constant function.

PROOF. $\bar{\bullet}$ satisfies (K*1)-(K*8) and (O) iff, by Theorem 6.9 and 6.11, $\bar{\bullet}$ is a analytic revision in $\langle W, d\rangle$ s.t. if K is inconsistent then $K\bar{\bullet}A = \mathrm{Cn}(A)$ iff $\bar{\bullet}$ is a analytic revision in $\langle W, d\rangle$ and for any A, $\{w \in [A] : d(\emptyset, w) = d(\emptyset, [A])\} = [A]$. Now, $\{w \in [A] : d(\emptyset, w) = d(\emptyset, [A])\} = [A]$ iff for any v, w in $[A]$, $f(w)= f(v)$ iff f is a constant function. QED

6.3 Representation Theorems

Theorem 6.11 proved the correspondence between analytic revision and AGM transitively relational partial meet revisions of a given theory. However, analytic revisions are defined for every theory, not just for a given theory, and the analytic revisions of different theories are not independent. Thus the question whether we can characterize the family of AGM unary functions corresponding to a given analytic operation remains. We are looking for the postulates that link the behaviour of revision of different theories. In the case of a finite propositional language the needed postulate the D-Ventilation condition that we introduced in Chapter 3 as the dual to the Ventilation, which we now name

(K*fin) $\quad (K_1 \cap K_2) * A \in \{K_1 * A, K_2 * A, (K_1 * A) \cap (K_2 * A)\}.$

(K*fin) forces a constraint between the revision of a theory and the revision of theories in which it is included. We can indeed show that in a finite language, (K*1)-(K*8) and (K*fin) completely characterize analytic revision functions.

Theorem 6.13 Given a finite propositional language L, an operator $*$ satisfies postulates (K*1)-(K*8) and (K*fin) if and only if there exists an analytic revision function $\bar{\bullet}$ such that for any $K \in \mathbb{K}$, $A \in L$, $K * A = K\bar{\bullet}A$.

PROOF. By Theorem 6.9 we know that $\bar{\bullet}$ validates (K*1)-(K*8). We shall verify that $\bar{\bullet}$ also validates (K*fin).

Let M be any model for $\bar{\bullet}$ $M = \langle W, d\rangle$, A any formula of L and K any theory of L such that $K = K_1 \cap K_2$ for theories K_1, K_2.

We have to show that in model M, $[K\bar{\bullet}A] \in \{[K_1\bar{\bullet}A], [K_2\bar{\bullet}A], [(K_1\bar{\bullet}A)] \cup [(K_2\bar{\bullet}A)]\}$.

By Definition 6.3 $[K\bar{\bullet}A] = \{v \in [A] : d([K], v) = d([K], [A])\}$. Also by definition, $d([K], v) = d([K_1] \cup [K_2], v) = \min\{d([K_1], v), d([K_2], v)\}$ and $d([K], [A]) = d([K_1] \cup [K_2], [A]) = \min\{d([K1], [A]), d([K_2], [A])\}$.

Then either $d([K_1], [A]) < d([K_2], [A])$ and $K\bar{\bullet}A = K_1\bar{\bullet}A$, or $d([K_2], [A]) < d([K_1], [A])$ and $K\bar{\bullet}A = K_2\bar{\bullet}A$, or $d([K_1], [A]) = d([K_2], [A])$ and then $K\bar{\bullet}A = K_1\bar{\bullet}A \cap K_2\bar{\bullet}A$.

By Theorem 6.11, given a fixed theory K, $*$ restricted to K is a analytic revision function, but a priori, with respect to different models M_K, one for each theory. We want to prove that this family of functions can actually be obtained from a single update model. i.e. that when considered as a binary function, $*$ can be obtained in the semantic framework of analytic revision.

Take the following model, $M = \langle W, d \rangle$ where W is the set of complete, consistent theories of the language and d is defined as $d(w, v) = d_w(v)$, for d_w a function characterizing the behaviour of $*$ when taking w fixed as first parameter. Also $d(\emptyset, v) = d_\emptyset(v)$. We extend d to a function on sets as we did before, by means of the min function. We now proceed by induction on the size of K.

Clearly, if K is empty or a singleton, $K * A = K\bar{\bullet}A$, by definition of d. Suppose K is not a singleton.

$[K\bar{\bullet}A \subseteq K * A]$. We want to show that if $w \in [K * A]$ then $w \in [K] \bullet [A]$. Clearly, $K\bar{\bullet}A = K * A$ for $[K]$ a singleton or $[K]$ the empty set.

Assume $[K] = \{x_1, \ldots, x_n\}$, $v \in [K * A]$ and $v \notin [K\bar{\bullet}A]$. Since $*$ validates (K*fin) and K is finite, then there must be some x in $[K]$ such that $v \in [x * A]$. Let $\mathrm{IN} = \{x \in [K] : v \in [x * A]\}$. Also, by Definition 6.3 there must exist some $y \in [K]$ such that $d(y, [A]) = d([K], [A])$. Then $v \notin \{y\} \bullet [A]$. Hence $v \notin [y * A]$. Let $\mathrm{OUT} = \{y \in [K] : v \notin [y * A]\}$.

Consider the following sets of two elements, $\{y_1, y_2\} \subseteq \mathrm{OUT}$, then trivially, by an application of (K*fin) $v \notin \{y_1, y_2\} * A$. Take now $\{x, y\}$ such that $x \in \mathrm{IN}$ and $y \in \mathrm{OUT}$, then either (1) $d(x, [A]) < d(y, [A])$ or (2) $d(x, [A]) = d(y, [A])$ or (3) $d(x, [A]) > d(y, [A])$. But (1) is impossible since $x, y \in [K]$

and $d(y,[A]) = d([K],[A])$. If (2) holds then, (using that $v \in [x * A]$), $d(y,[A]) = d(x,v)$. Therefore, $v \in [K] \bullet [A]$, contrary to our assumption. Then (3) should be the case for any pair x, y. According to our definition of d, $c^{\{y\}}(A) = c^{\{x,y\}}(A)$ and $c^{\{x\}}(A) \neq c^{\{x,y\}}(A)$. Hence $\{x,y\} * A = y * A$, therefore, $v \notin \{x,y\} * A$.

Now we are almost done. Notice that by pairing elements of IN with elements of OUT we can "delete" the elements of IN from $[K]$. I.e. let $x \in$ IN, $y \in$ OUT and write $[K]$ as $\{x,y\} \cup ([K] \setminus \{x\})$, then applying (K*fin) $v \in [\mathrm{Th}([K] \setminus \{x\}) * A]$. Because IN is finite, we will finally have $v \in [\mathrm{Th}(\mathrm{OUT}) * A]$. A contradiction.

$[K * A \subseteq K \bar{\bullet} A]$. Let $u \in [K \bar{\bullet} A]$ and let $x \in [K]$ such that $d([K],[A]) = d(x,u)$. Then $u \in [x * A]$. Also, because K is finite, by repeatedly applying (K*fin) we have $[K * A] = \bigcup [T_i * A]$ for some complete theories T_i extending K. If $x = T_i$ for some i we are done. Suppose $u \notin [T_i * A]$ for any T_i. We now use again (K*fin) and comparison of pairs to arrive to a contradiction (write $[K] = \{x, T_i\} \cup ([K] \setminus \{T_i\})$ and consider $K * A \subseteq \mathrm{Th}(\{x, T_i\}) * A$ must hold for each T_i). Full details are given for the case of infinite languages in Theorem 6.15. QED

Postulate (K*fin) appears in [Schlechta *et al.*, 1996] as a property that revisions based on pseudo distances satisfy.

The general case is slightly harder. Postulates (K*1)-(K*8) and (K*fin) do not fully characterize the $\bar{\bullet}$ operation in a language with an infinite number of propositional letters.

Observation 6.14 Consider an infinite propositional language L. Postulates (K*1)-(K*8) and (K*fin) do not fully characterize the $\bar{\bullet}$ operation.

PROOF. Given a propositional language L with an infinite but countable number of propositional letters we will exhibit a function $*$ satisfying postulates (K*1)-(K*8) and (K*fin) for which there is no model $M = \langle W, d \rangle$, satisfying that $\forall K \in \mathbb{K}, \forall A \in L, K * A = K \bar{\bullet} A$. We define $*$ semantically as follows. Let $K \in \mathbb{K}$, $A \in L$ and $v \in [A]$, then

$$[K * A] = \begin{cases} [K] \cap [A] & \text{, if } [K] \cap [A] \neq \emptyset. \\ \{v\} & \text{, if } [K] \cap [A] = \emptyset \text{ and } [K] \text{ is finite.} \\ [A] & \text{, if } [K] \cap [A] = \emptyset \text{ and } [K] \text{ is infinite.} \end{cases}$$

For each incomplete theory $K \in \mathbb{K}$ such that $[K]$ has a finite number of elements (i.e., there are only a finite number of maximal consistent sets extending K), then let $*_K$ be a *fixed* AGM maxichoice revision function for K always returning one and the same maximal consistent set of A. And for each incomplete theory $K \in \mathbb{K}$ such that $[K]$ has an infinite number of elements then let $*_K$ be the full meet revision function for K, namely $K * A = \text{Cn}(A)$.

Clearly $*$ validates (K*fin). If $[K]$ is finite it is easily verified. If $[K]$ is infinite, for any theories K_1, K_2 such that $K = K_1 \cap K_2$, either $[K_1]$ or $[K_2]$ are infinite. Then either $K_1 * A = \text{Cn}(A)$ or $K_2 * A = \text{Cn}(A)$, as required.

Suppose for contradiction that there is a model $M = \langle W, d \rangle$ such that for every $K \in \mathbb{K}$, for every $A \in L$, $K * A = \text{Th}(\{y \in [A] : d([K], A) = d([K], y)\})$.

According to our definition of $*$, for every theory K such that $[K]$ is finite, if $[K] \cap [A] = \emptyset$ then $[K * A] = \{v\}$. Therefore d must verify that $\forall x \in W, d(x,x) = 0$; $\forall x, w \in W, w \neq v$, $d(x,v) < d(x,w)$.

For any $[K]$ such that $[K] \cap [A] = \emptyset$, Then $0 < d([K],[A]) = d([K], v)$, since for each $x \in [K]$, $d(x,v) = d(x,[A])$. Then $[K * A] = \{y \in [A] : d([K], A) = d([K], y)\} = \{v\}$. This contradicts the case when $[K]$ is infinite, because according to our definition $[K * A] = [A]$. QED

(K*fin) gives us the following insight: when performing the analytic revision of K by A, we should hear the opinions of the theories to which K can be extended. If we now turn to the way $\bullet$ is defined given K and A, we see that we can always identify an element w of $[K]$ which is responsible for defining $d([K],[A])$. Then $[K] \bullet [A]$ is obtained as the subset of $[A]$ standing at the same distance from $[K]$ as w is. These complete theories are clearly the ones we should pay attention to. Following this intuition we propose:

(K$*\exists$) $K * A = \bigcap(T_i * A)$, for some complete theories T_i extending K.

(K$*\forall$) If $K \subseteq K' \subseteq T$, for T a complete theory then, for all A, $K * A \subseteq T * A$ implies $K * A \subseteq K' * A \subseteq T * A$.

(K$*\exists$) claims there are some complete theories — "the intended interpretations" of our theory — that determine the result of the revision. (K$*\forall$) expresses the primacy of these complete theories and establishes a restricted form of monotony for the $*$ operator. In particular, if our theory K is regarded as an intersection of two larger theories K_1 and K_2, then (K$*\exists$) and (K$*\forall$) constrain the revision of K in terms of the other two. By (K$*\exists$) the revision of each K is guided by some complete theories. These complete theories either extend K_1 or K_2 or both. Then, by (K$*\forall$) the revision of K is included in the revision of K_1 or in the revision of K_2, or both. Notice that, in the presence of (K*1)-(K*8), the postulates (K$*\exists$) and (K$*\forall$) imply (K*fin).

We now prove that the eight AGM postulates plus (K$*\exists$) and (K$*\forall$) completely characterize the analytic revision operation.

Theorem 6.15 (Representation Theorem, general case) An operator $*$ satisfies postulates (K*1)-(K*8), (K$*\exists$) and (K$*\forall$) if and only if there exists a model $M = \langle W, d\rangle$, where d is a distance function and for any $K \in \mathbb{K}$, $A \in L$ $K * A = K \bar{\bullet} A$.

PROOF. We have proved in Theorem 6.9 that $\bar{\bullet}$ satisfies postulates (K*1)-(K*8). That $\bar{\bullet}$ validates (K$*\exists$) follows immediately from Definition 6.2, since min requires the existence of elements in $[K]$ such that their distance to $[A]$ is minimal. $\bar{\bullet}$ also validates (K$*\forall$) since for any Y if $x \in [K]$ and $d([K], Y) = d(x, Y)$ then $d(x, Y) = \min_{z \in [K]}\{d(z, Y)\}$. Therefore, for all $X \subseteq [K]$, if $x \in X$ then $d(X, Y) = d(x, Y)$ and $d(X, Y) = d([K], Y)$ as required. This proves the right to left implication.

Let's see the left to right part. Let $*$ be a change function satisfying (K*1)-(K*8),(K$*\exists$) and (K$*\forall$). We will construct a analytic revision model $M = \langle W, d\rangle$ which corresponds to $\bullet$.

We have to show that $\forall K \in \mathbb{K}, \forall A \in L, K * A = K \bar{\bullet} A$.

We start by defining the model M. The domain W will be the set of all complete theories in the language L. To define the distance function d, let $\{S^K\}$ be the family of systems of spheres corresponding to $*$. If S^K is a given system of sphere we note as S_i^K a particular element of it, and for a given formula A, $c^K(A)$ is the minimal sphere in S^K with nonempty intersection with $[A]$.

As before, we start by determining the value of d for elements in W and then extend the function to subsets of W as in Definition 6.2. Any function $d : \mathcal{P}(W) \times \mathcal{P}(W) \to \mathbb{R}^+$ satisfying the following restrictions is appropriate.

i. $\forall v \in W,\ d(v,v) = 0$.

ii. $\forall v,u,m, d(v,u) < d(v,m)$ iff $\exists S_1^v, S_2^v \in S^v u \in S_1^v, m \in S_2^v \& S_1^v \subset S_2^v$.

iii. $\forall v,u,m, d(v,u) = d(v,m)$ iff $\forall S_i^v \in S^v u \in S_i^v \Leftrightarrow m \in S_i^v$.

iv. $d(\{x,y\},X) = d(x,X)$ iff $c^{\{x\}}(X) = c^{\{x,y\}}$.

v. $d(x,X) < d(y,X)$ iff $c^{\{x\}}(X) = c^{\{x,y\}}(X)$ and $c^{\{y\}}(X) \neq c^{\{x,y\}}(X)$.

vi. $d(x,X) = d(y,X)$ iff $c^{\{x\}}(X) \cup c^{\{y\}}(X) = c^{\{x,y\}}(X)$.

vii. $\forall v,u,m, d(\emptyset,u) < d(\emptyset,m)$ iff $\exists S_1^\emptyset, S_2^\emptyset \in S^\emptyset u \in S_1^\emptyset, m \in S_2^\emptyset \& S_1^\emptyset \subset S_2^\emptyset$.

To verify that there are indeed distance functions satisfying i) to vii) above is easy. It is also clear that by case vii), when K is the inconsistent theory $K * A$ and $K \bar{\bullet} A$ agree. Furthermore if $[A] = \emptyset$, by $(K*5)$, $K * A = L$, and also $K \bar{\bullet} A = L$ by definition. We will now prove, for K and A consistent, that $u \in [K * A]$ iff $u \in [K \bar{\bullet} A]$ by analyzing the different cases.

Suppose $[K] = \{v\}$.

$[K \bar{\bullet} A \subseteq K * A]$. Let $u \in [K * A]$, to prove (1) $u \in \{w \in [A] : d(w,[A]) = d(v,w)\}$. Let $m \in [A]$ be such that $d(v,[A]) = d(v,m)$, then (1) is equivalent to (2) $d(v,m) = d(v,u)$. By iii) we have to prove that for all $S_i^{\{v\}} \in S^{\{v\}}, u \in S_i^{\{v\}} \Leftrightarrow m \in S_i^{\{v\}}$. As $d(v,[A]) = d(v,m)$ then $m \in c^{\{v\}}(A)$. Let $S_i^{\{v\}}$ be any. If $c^{\{v\}}(A) \subseteq S_i^{\{v\}}$ then both m and u are in $S_i^{\{v\}}$. If $S_i^{\{v\}} \subset c^{\{v\}}(A)$, then $u \notin S_i^{\{v\}}$. Suppose $m \in S_i^{\{v\}}$, but then $d(v,[A]) > d(v,m)$, a contradiction.

$[K * A \subseteq K\bar{\bullet}A]$. To prove the other inclusion, let $u \in [A]$ and suppose $d(v,m) = d(v,u)$ for $m \in [A]$ such that $d(v,[A]) = d(v,m)$. Suppose $u \notin c^{\{v\}}(A)$. Then by iii) $m \notin c^{\{v\}}(A)$. Let $S_i^{\{v\}}$ be the $\subseteq$-smallest such that $m \in S_i^{\{v\}}$, $c^{\{v\}}(A) \subset S_i^{\{v\}}$. By the limit assumption $c^{\{v\}}(A)$ is defined and let $m' \in c^{\{v\}}(A) \cap [A]$. But then by i) $d(v,m') < d(v,m)$ contradicting the selection of m.

The general case, $[K] > 1$.

$[K * A \subseteq K\bar{\bullet}A]$. Let $u \in [K\bar{\bullet}A]$ and let $x \in [K]$ be such that $d([K],[A]) = d(x,u)$ (notice that then, $u \in [x\bar{\bullet}A]$ and by the previous case $u \in [x * A]$). By (K$*\exists$), $[K * A] = \bigcup[T_i * A]$ for some complete theories extending K. If for some i, $u \in [T_i * A]$ we are done, so assume $u \notin [T_i * A]$ for all i.

Consider for any i the proposition $\{x, T_i\} \subseteq [K]$. Then by (K$*\forall$), $[T_i * A] \subseteq [\text{Th}(\{x,T_i\}) * A] \subseteq [K * A]$. Apply (K$*\exists$) to $\text{Th}(\{x,T_i\}) * A$ now. If $[x * A] \subseteq [\text{Th}(\{x,T_i\}) * A]$ we are done. Rests to consider the case when $[\text{Th}(\{x,T_i\}) * A] = [T_i * A]$, and furthermore $[\text{Th}(\{x,T_i\}) * A] \neq [x * A]$. But then by condition v), $d(T_i,[A]) < d(x,[A])$, contradicting the choice of x.

$[K\bar{\bullet}A \subseteq K * A]$. For this inclusion, we should further prove the case for $[K] = \{v,w\}$ separately. Suppose $u \in [K * A]$, then by (K$*\exists$), $u \in \bigcup[T_i * A]$ for some T_i complete theories extending K, either

a. $K * A = v * A$. Then by iii), $d(\{v,w\},[A]) = d(v,[A])$. As $u \in c^{\{v,w\}}(A)$, by definition of d, i) and ii) we have that $d(v,u) = d(v,[A]) = d([K],[A])$. Hence $u \in [K] \bullet [A]$.

b. $K * A = w * A$. Similar to a.

c. $K * A = v * A \cap w * A$. By iv), $d(v,[A]) = d(w,[A])$. Also, either $u \in c^{\{v\}}(A)$ or $u \in c^{\{w\}}(A)$. Hence, as above, either $d(v,u) = d(v,[A])$ or $d(w,u) = d(v,[A])$. In both cases, $u \in [K] \bullet [A]$.

$[K] > 2$. Suppose $u \in [K * A]$, then by (K$*\exists$), $u \in \bigcup[T_i * A]$ for some T_i complete theories extending K. In particular, let $T_i \in W$ be such that $u \in [T_i * A]$.

Let x be any in $[K]$, by (K$*\forall$), $K * A \subseteq T_i * A$ implies $(T_i \cap x) * A \subseteq T_i * A$.

Hence, $(T_i \cap x) * A \subseteq T_i * A$. We are now in the previous cases, of revising theories whose proposition has cardinality one or two. Therefore we can claim that $(T_i \cap x)\bar{\bullet}A \subseteq T_i\bar{\bullet}A$. I.e., by definition for all $w \in [A]$, $d(\{T_i, x\}, [A]) = d(\{T_i, x\}, w)$ then $d(T_i, [A]) = d(\{T_i, x\}, w)$, iff for all $w \in [A]$, $\min\{d(T_i, [A]), d(x, [A])\} = \min\{d(T_i, w), d(x, w)\}$ then $d(T_i, [A]) = d(T_i, w)$.

Therefore $d(T_i, [A]) = d(\{T_i, x\}, [A])$. As this is true for all $x \in [K]$, $d(T_i, [A]) = d([K], [A])$. Because $u \in [T_i\bar{\bullet}A]$, $d(T_i, u) = d(T_i, [A])$ and $u \in K\bar{\bullet}A$. QED

Hence, the analytic revision function is indeed a binary AGM function.

Theorems 6.13 and 6.15 are interesting because they give general characterization results for AGM revisions based on pseudo-distances, for both, the finite and the general cases.

We now turn our attention to two natural constraints on the distance functions which give rise to proper subclasses of analytic AGM revisions. One is to consider a distance function $d : W \times W \longrightarrow I\!R^+$ is such that no two points are at the same distance from a given point, if $d(v, u) = d(v, w)$ then $v = w$. This is to take d_v, the the projection of the distance function over its first argument, to be injective. It is quite strightforward to prove that such a distance function gives rise to an analytic AGM revision that takes consistent complete theories to consistent complete theories. For complete theories this analytic function behaves as a maxichoice AGM revision. For this reason we name it *maxi-analytic AGM functions*, and we show that they are characterized by the following postulate.

(K*M) If K is consistent and complete then, for any A, $K * A$ is complete.

Observation 6.16 (maxi-analytic AGM functions) An operator $*$ satisfies postulates (K*1)-(K*8), (K*∃) (K*∀) and (K*M) if and only if there exists a distance model $M = \langle W, d\rangle$, such that for each $v \in W$, $d_v = d(v, w)$ is injective, and for any $K \in I\!K$, $A \in L$ $K * A = K\bar{\bullet}A$.

PROOF. The characterization result follows directly for the fact that for every nameable $Y \subseteq W$, $\{x | d_v(Y) = d_v(x)\}$ is a singleton. QED

Another interesting consideration is the case of well founded distances, that is distances that are definable over the ordinals, $d : W \times W \rightarrow \mathcal{O}$. Applying Observation 2.14, a well founded system of spheres centered in $[K]$ can be represented by ordinal function $d_K : W \rightarrow \mathcal{O}$. In this setting actual values of the function $d(w, v)$ can be obtained by counting the number of ancestors of the argument along the well founded system of spheres centered in $\{w\}$. As a result we can precise an actual mapping of well founded update models $\langle W, \{\leq_w : w \in W\}\rangle$ to well founded distance models $\langle W, d\rangle$. In section 2.3.4 we reported that [Peppas, 1993] characterized the class of AGM revision functions, that are definable over well founded system of spheres. Peppas called them *well behaved revision functions* and showed that they are characterized by postulates (K*1)-(K*8) plus

(K*WB) For every nonempty set X of consistent formulae of L there exists a formula $A \in X$ such that $\neg A \notin K * (A \vee B)$, for every $B \in X$.

Of course, this characterization carries over analytic functions and update functions. Well behaved analytic AGM functions satisfy (K*1)-(K*8),(K$*\exists$),(K$*\forall$) and (K*WB), and are a proper subclass of general analytic functions that can be characterized semantically by a distance function d over the ordinals.

It is apparent from the proofs of Theorems 6.13 and 6.15 that the distance function that we use is just a convenient means to express the comparative relations relative to sets, that are induced from the comparative relations relative to single points. In fact the analytic operation can be regarded as a particular case of a more general framework. Consider a model with two ordering relations, $\langle W, \{\preceq^1_w : w \in W\}, \{\preceq^2_X : X \in \mathcal{P}(W)\}\rangle$, being $\preceq^1$, $\preceq^2$ possibly independent (total) preorders on W. Then the $\bullet$ operation would be a double minimization over the two relations, defined as

$$\min_{\preceq^2_X} \bigcup_{x \in X} \min_{\preceq^1_x}(Y)$$

where $\min_{\preceq}(V) = \{v \in V : \forall z \in V, v \preceq z\}$. Our definition of analytic revision in terms of distances obtains in this general framework, by considering $\preceq^1$ as an ordering encoding $d : W \times W \rightarrow \mathbb{R}^+$ and $\preceq^2$ as one encoding the extension $d : \mathcal{P}(W) \times \mathcal{P}(W) \rightarrow \mathbb{R}^+$. We believe it is interesting to study characterization results for the double minimization operation on the general framework. This seems to be the proper setup to investigate which are the needed properties connecting the two orderings as well as the particular properties of each of them.

6.4 Properties

We turn now into properties of analytic revisions. Trivially, it is possible to define a binary AGM function as an accidental collection of unary AGM functions, one for *each* theory. But, if there are no properties linking the revisions of the different theories the result obtained can be erratic. For example, as we reported in Chapter 4 the postulate (K*9) counts as simple way of linking the revisions of all different theories. It is apparent that the link between the analytic revisions of different theories is more subtle than the link provided by (K*9).

Observation 6.17 The following properties are not validated by the analytic revision operation.

(Commutativity) $(K * A) * B = (K * B) * A$.
(Weak Intersection) If $\neg A \in K_1 \cap K_2$ then $(K_1 \cap K_2) * A = (K_1 * A) \cap (K_2 * A)$.
(Union) $(K_1 \cup K_2) * A = (K_1 * A) \cup (K_2 * A)$.
(Weak Union) If $\neg A \in K_1 \cap K_2$ then $(K_1 \cup K_2) * A = (K_1 * A) \cup (K_2 * A)$.
(K*9) If $\neg A \in K$, $K * A = L * A$.

PROOF. It is not difficult to find analytic revision functions violating each of these conditions. We prove Commutativity. Let L be a propositional language over $\{A, B\}$. Let $[A] = \{w_2, w_4\}$ and $[B] = \{w_3, w_4\}$. Assume $[K] = \{w_1\}$ and $d(w_1, w_2) < d(w_1, w_4)$, $d(w_1, w_3) < d(w_1, w_4)$, $d(w_2, w_3) < d(w_2, w_4)$, and $d(w_3, w_4) < d(w3, w_2)$.

$[K \bar{\bullet} A] = \{w_2\}$, $[K \bar{\bullet} B] = \{w_3\}$, $[K \bar{\bullet} A \bar{\bullet} B] = \{w_3\}$ but $[K \bar{\bullet} B \bar{\bullet} A] = \{w_4\}$.
QED

Being binary AGM functions, analytic revisions can freely perform iterated change, inheriting the form of iteration of the standard update operation. The formal structure $M = \langle W, d\rangle$ determines the distance from every $[K]$. Since the analytic revision of K by A is a theory $K\bar{\bullet}A$, also a proposition in the same model M, distance from $K\bar{\bullet}A$ is also defined in the structure. However, as we have already seen, there is a crucial difference between the iterating capabilities of the two forms of update: the standard update can not recover from inconsistency, while analytic revision can. Moreover, analytic revisions satisfy some natural conditions of iterated change. For any pair of theories K_1, K_2 and sentences A, B, C, D,

(Or-Left) If $D \in (K*(A \vee B))*C$ then $D \in (K*A)*C$ or $D \in (K*B)*C$.

(Or-Right) If $D \in (K*A)*C$ and $D \in (K*B)*C$ then $D \in (K*(A\vee B))*C$.

Observation 6.18 Analytic revision functions satisfy Or-Left and Or-Right.

Proof. Let's name $X = [K] \bullet [A]$, $Y = [K] \bullet [B]$.

(Or-Left). $[K] \bullet [A \vee B]) \bullet [C] = \{w \in [C]$ $\min_{\{x\in[K]\bullet[A\vee B]\}} \min_{\{y\in[C]\}}\{d(x,y)\}\} =$

(by (K*7) and (K*8)) $[K]\bullet[A\vee B] = [K]\bullet[A]$, or $[K]\bullet[A\vee B] = [K]\bullet[B]$, or $[K]\bullet[A\vee B] = ([K]\bullet[A]) \cup ([K]\bullet[B])$.

Then, either

(1) $\{w \in [C] : \min_{\{x\in X\}} \min_{\{y\in[C]\}}\{d(x,y)\}\} = [K]\bullet[A]$; or

(2) $\{w \in [C] : \min_{\{x\in Y\}} \min_{\{y\in[C]\}}\{d(x,y)\}\} = [K]\bullet[B]$; or

(3) $\{w \in [C] : \min_{\{x\in X\cup Y\}} \min_{\{y\in[C]\}}\{d(x,y)\}\} = \{w \in [C] : \min\{\min_{\{x\in X\}} \min_{\{y\in[C]\}}\{d(x,y)\}, \min_{\{x\in Y\}} \min_{\{y\in[C]\}}\{d(x,y)\}\}\}$

is either equal to $[K]\bullet[A]$ or it is equal to $[K]\bullet[B]$.

(Or-Right). Assume (1) $D \in (K\bar{\bullet}A)\bar{\bullet}C$ and (2) $D \in (K\bar{\bullet}B)\bar{\bullet}C$.

By (1) $\{w \in [C] : \min_{\{x\in X\}} \min_{\{z\in[C]\}}\{d(x,z)\}\} \subseteq [D]$.

By (2) $\{w \in [C] : \min_{\{y\in Y\}} \min_{\{z\in[C]\}}\{d(y,z)\}\} \subseteq [D]$.

And $[K]\bullet[A\vee B]\bullet[C] =$

$\{w \in [C] : \min_{\{x\in X\cup Y\}} \min_{\{z\in[C]\}}\{d(x,z)\}\} = \{w \in [C] :$
$\min(\min_{\{x\in X\}} \min_{\{z\in[C]\}}\{d(x,z)\}, \min_{\{x\in Y\}} \min_{\{z\in[C]\}}\{d(x,z)\})\}$

is either equal to $[K] \bullet [A]$ or is equal to $[K] \bullet [B]$.

Then $[K] \bullet [A \vee B] \bullet [C] \subseteq [D]$. QED

The analytic revision function validates five out of these seven postulates of [Lehmann, 1995].

(I1) $K * A$ is a consistent theory.

(I2) $A \in K * A$.

(I3) If $B \in K * A$, then $A \supset B \in K$.

(I4) If $A \in K$ then $K * B_1 * \ldots * B_n = K * A * B_1 * \ldots * B_n$ for $n \geq 1$.

(I5) If $A \in \mathrm{Cn}(B)$, then $K * A * B * B_1 * \ldots * B_n = K * B * B_1 * \ldots * B_n$.

(I6) If $\neg B \notin K * A$ then $K * A * B * B_1 * \ldots * B_n = K * A * (A \wedge B) * B_1 * \ldots * B_n$.

(I7) $K * \neg B * B \subseteq \mathrm{Cn}(K \cup \{B\})$.

Condition (I7) implies dependency between two revision steps and consequently enforces (at least to some extent) the property of "historical memory" which analytic revisions lack. As remarked by Lehmann, the standard update operation fails postulates (I4), (I5) and (I7), and satisfies the rest. It is then expected that the analytic revision operation violates (I5) and (I7) and validates the rest.

Observation 6.19

i) All analytic revision functions satisfy (I1), (I2), (I3), (I4) and (I6).

ii) There exist analytic revision functions violating (I5) and (I7).

PROOF. The violation of (I5) and (I7) can be proved by constructing a counterexample.

(I1), (I2), (I3), (I4) follow from the AGM postulates (K*1)-(K*4).

For (I6) we should prove that if $\neg B \notin K \bar{\bullet} A$ then $K \bar{\bullet} A \bar{\bullet} B = K \bar{\bullet} A \bar{\bullet} (A \wedge B)$. But this is obvious since $K \bar{\bullet} A \bar{\bullet} B = \mathrm{Cn}(K \bar{\bullet} A \cup \{B\}) = \mathrm{Cn}(K \bar{\bullet} A \cup \{A \wedge B\})$.

QED

Analytic revisions do not validate any of Darwiche and Pearl's postulates [1997].

(C1) If $A \in \mathrm{Cn}(B)$ then $(K * A) * B = K * B$.

(C2) If $\neg A \in \mathrm{Cn}(B)$ then $(K * A) * B = K * B$.

(C3) If $A \in K * B$ then $A \in (K * A) * B$.

(C4) If $\neg A \notin K * B$ then $\neg A \notin (K * A) * B$.

(C5) If $\neg B \in K * A$ and $A \notin K * B$ then $A \notin (K * A) * B$.

(C6) If $\neg B \in K * A$ and $\neg A \in K * B$ then $\neg A \in (K * A) * B$.

Observation 6.20

There exist analytic revision functions violating each of (C1)-(C6).

PROOF. C1 is just postulate I5 above.

C2. Assume a propositional language with three variables A, B and C. Let $w \in [\neg A] \cap [\neg B]$, $z \in [A] \cap [\neg B]$, $v \in [\neg A] \cap [B] \cap [C]$ and $u \in [\neg A] \cap [B] \cap [\neg C]$. Suppose $d(z, v) < d(z, u)$ and $d(w, u) < d(w, v)$. Let $[K] = \{w\}$. Then $[K \bar{\bullet} A] = \{z\}$, $[K \bar{\bullet} A \bar{\bullet} B] = \{v\}$ but $[K \bar{\bullet} B] = \{u\}$.

C3 and C4. Let $w \in [\neg A] \cap [\neg B]$, $z \in [A] \cap [\neg B]$, $v \in [\neg A] \cap [B] \cap [C]$ and $x \in [A] \cap [B] \cap [C]$. Suppose $d(w, z) < d(w, x) < d(w, v)$ and $d(z, v) < d(z, x)$. Let $[K] = \{w\}$. Then $[K \bar{\bullet} B] = \{x\} \subseteq [A]$, $[K \bar{\bullet} A] = \{z\}$ and $[K \bar{\bullet} A \bar{\bullet} B] = \{v\} \not\subseteq [A]$.

C5 and C6. Let $w \in [\neg A] \cap [\neg B]$, $z \in [A] \cap [\neg B]$, and $x \in [A] \cap [B]$ and $u \in [\neg A] \cap [B]$. Suppose $d(w, z) < d(w, u) < d(w, x)$ and $d(z, x) < d(z, u)$. Let $[K] = \{w\}$. Then $[K \bar{\bullet} B] = \{u\} \subseteq [\neg A]$, $[K \bar{\bullet} A] = \{z\} \subseteq [\neg B]$ but $[K \bar{\bullet} A \bar{\bullet} B] = \{x\} \subseteq [A]$.

QED

As expected, analytic AGM revisions do not validate the postulates of iterative schemes.

(T) $K \circ A \circ B = K \circ B$.

(C) If $\neg B \in K \circ A$, then $K \circ A \circ B = K \circ B$.

(I) $K \circ A \circ B = K \circ (A \wedge B)$.

(M)

$$K \circ A \circ B = \begin{cases} K \circ B & \text{, if } \neg B \in \mathrm{Cn}(A) \\ K \circ (A \wedge B) & \text{, otherwise.} \end{cases}$$

Observation 6.21 There exist analytic revisions violating (T),(M),(I) and (C).

Chapter 7

Logical Calculi for Theory Change

Alchourrón's logic DFT [Alchourrón, 1995] and Boutilier's CO [Boutilier, 1992a] are conditional logics that provide a logical calculus for the AGM theory. In a very natural way they can be used to calculate changes in different theories, by appealing to the consequence operation in each logic. Both logics share the special characteristics with respect to the conditional connective common to most logics for defeasible inference. Namely, they defeat the rules of Modus Ponens, Strengthening the antecedent, Transitivity, and Contraposition. But the two logics are clearly different. Although both are modal conditional logics with possible worlds semantics, CO has a relational semantics requiring a preorder over possible worlds, while DFT possesses a non-relational semantics based on a selection function Ch defined over the logical language. They also differ in their expressive power and have quite different axiomatic presentations. Specially, the respective definitions of the conditional connectives stand on different grounds.

In this chapter we compare the two logics and investigate their connection. After considering some general results of [Rott, 1993] showing links between selection functions and binary relations, we will briefly present the two logics assuming basic knowledge of the standard modal systems. (For a thorough presentation of standard modal systems see [Chellas, 1980;

Hughes and Cresswell, 1968; Hughes and Cresswell, 1984]). In particular we will consider the modal systems S5 and S4.3 (S5 is the extension of classical propositional logic with the Necessitation rule and the characteristic axioms K, 4 and 5; while S4.3 possesses the characteristic axioms K, 4 and 3). We will then reveal the connection between DFT and CO by two main results. One is that there is a one to one correspondence between the finite models of the two logics. The other is that the respective definitions of the conditional connectives are semantically equivalent. These two results will allow us to prove that satisfiable sentences in the respective finite propositional languages augmented solely by the respective conditional connectives are in a one to one correspondence. Since the conditional connectives have the same interpretation we will conclude that, in the restricted language, the two logics validate the *same* conditional sentences.

As of notation, the symbol $\vdash$ will be used to indicate derivability in different systems, using a subscript to specify the system. Semantic entailment will be denoted with the symbol $\models$. To denote satisfiability in a point w of a model M we will use $M \models_w$. In addition we will refer to the set of models for a set of sentences X as: $Mods(X)$=$\{M : M \models A,$ for each $A \in X\}$.

7.1 Selection functions and Binary relations

Let X be a set and $\mathcal{X}$ be a non-empty subset of $\mathcal{P}(X) \setminus \emptyset$. A selection function, or choice function over $\mathcal{X}$ is a function $s : \mathcal{X} \to \mathcal{P}(X)$ such that $s(Y)$ is a non empty subset of $Y \in \mathcal{P}(X)$. Intuitively, selection functions are supposed to give us the "best" elements of each $Y \in \mathcal{P}(X)$. The requirement that $s(Y)$ be non-empty means that the selection function is effective.

A set $\mathcal{X}$ of subsets of X is called *n-covering* ($n = 1, 2, 3...$) if it contains all subsets of X with exactly n elements, $\mathcal{X}$ is called n_1n_2-covering if it is n_1-covering and n_2-covering. $\mathcal{X}$ is called *ω-covering* if it is n-covering for all natural numbers $n = 1, 2, 3, \ldots$. A set $\mathcal{X}$ of subsets of X is called *additive* if it is closed under arbitrary unions, and it is called *finitely additive* if it is closed under finite unions. $\mathcal{X}$ is *substractive* if for every X and X' in $\mathcal{X}$ such that $X \not\subseteq X'$, $X \setminus X'$ is also in $\mathcal{X}$. (If X is 1-covering and finitely additive then $\mathcal{X}$ is ω-covering.)

Finally, $\mathcal{X}$ is *compact* if for every X and X_i, $i \in I$, if $X \subseteq \bigcup\{X_i : i \in I\}$ then $X \subseteq \bigcup\{X_i : i \in I_0\}$ for some finite $I_0 \subseteq I$.

For example, let L be an arbitrary infinite language, and W the set of maximal consistent extensions of L. For any language sentence A, $[A] = \{w \in W : A \in w\}$. Let $\mathcal{X}=\{[A] \subseteq W : A \in L\} \subseteq \mathcal{P}(W)$; be the set of nameable subsets of W. By cardinality considerations, $\mathcal{X}$ is a proper subset of $\mathcal{P}(W)$. Moreover, $\mathcal{X}$ is not additive nor finitely additive nor 1-covering nor compact nor substractive. However, if we take L a propositional language over a *finite* set of propositional variables P, and we take W as the set of all maximal consistent extensions of L, then $\mathcal{X}$ is finitely additive, n-covering, substractive and compact.

A selection function with domain $\mathcal{X}$ is said to be n-covering, (finitely) additive, substractive, etc., if $\mathcal{X}$ is n-covering, (finitely) additive, substractive, etc. Rott shows that under certain conditions it is possible to recover the relations underlying choice functions. And conversely, under appropriate conditions a relation induces a selection function. Generically, choice sets are taken to be sets of "best" elements in some relation $\leq$. A selection function is *relational* with respect to $\leq$ over X, and we write $s = \mathcal{S}(\leq)$, when for every $Y \in \mathcal{P}(X)$

$$s(Y) = \{y \in Y : y \leq y' \text{ for all } y' \in Y\}.$$

Samuelson preferences are a classical way to recover a relation underlying a selection function:

$$\leq_s = \{\langle x, x'\rangle \in X \times X : \exists Y \in \mathcal{P}(X) \text{ such that } \langle x, x'\rangle \subseteq Y \text{ and } x' \in s(Y)\}$$

$\leq_s$ is not guaranteed to be reflexive unless s is 1-covering.

In order to show the correspondence of properties of selection functions and binary relations Rott [1993] formulates the following postulates.

I . For all $Y, Y' \in \mathcal{X}$ such that $Y \cup Y' \in \mathcal{X}$ $s(Y \cup Y') \subseteq s(Y) \cup s(Y')$.

II . For all $Y, Y' \in \mathcal{X}$ such that $Y \cup Y' \in \mathcal{X}$ $s(Y) \cap s(Y') \subseteq s(Y \cup Y')$.

III . For all $Y \in \mathcal{X}$ and Y' such that $Y \cup Y' \in \mathcal{X}$ if $s(Y \cup Y') \cap Y' \neq \emptyset$ then $s(Y) \subseteq s(Y \cup Y')$.

IV . For all $Y \in \mathcal{X}$ and Y' such that $Y \cup Y' \in \mathcal{X}$, if $s(Y \cup Y') \cap Y \neq \emptyset$ then $s(Y) \subseteq s(Y \cup Y')$.

The following lemmas show the connection between selection functions and preference relations.

Lemma 7.1 ([Rott, 1993],Lemma 1)
If s satisfies I and II and is 12-covering or additive then $s = \mathcal{S}(\leq_s)$.

Lemma 7.2 ([Rott, 1993],Lemma 2) (Notation adapted).

(a) If s is $12n$-covering and satisfies I then the complement of $\leq_s$ is n-acyclic. If s is ω-covering and satisfies I then the complement of $\leq_s$ is acyclic.

(b) If s is 123-covering and satisfies I and III then $\leq_s$ is transitive.

(c) If s is finitely additive and satisfies IV, then the complement of $\leq_s$ is transitive.

Lemma 7.3 ([Rott, 1993],Lemma 3) (Notation adapted).

(a) If the strict part of $\leq$ is well-founded with respect to $\mathcal{X}$ then $\mathcal{S}(\leq)$ is a selection function over $\mathcal{X}$ which satisfies (I) and (II).

(b) If $\leq$ is transitive then $\mathcal{S}(\leq)$ is a selection function over $\mathcal{X}$ which satisfies (III).

(c) If the complement of $\leq$ is transitive then $\mathcal{S}(\leq)$ is a selection function over $\mathcal{X}$ which satisfies (IV).

7.2 The Logic DFT

Alchourrón's modal conditional logic is based on a propositional language L augmented with an S5-necessity operator $\Box$ and a revision operator f, which is in fact another modality. We will refer to this modal language with L_{DFT}. Alchourrón bases his construction on the very idea that in a *defeasible conditional* the antecedent is a *contributory* condition of its consequent, as opposed to be a sufficient condition for the consequent. Hence, he defines a defeasible conditional $A >_{DFT} B$ meaning that the antecedent A jointly with the set of assumptions that comes with it is a sufficient condition for the consequent B. In order to represent in the object language the joint assertion of the proposition expressed by a sentence A and the set of assumptions (or presuppositions) that comes with it he uses a *revision operator* f. For example, if $A_1, \ldots A_n$ are the assumptions associated with A then fA stands for the joint assertion (conjunction) of A with all the A_i (for all $1 \leq i \leq n$), where A is always one of the conjuncts of fA. Although Alchourrón does not explicitly refer to the cardinality of the set of assumptions for a given proposition, this set may well be infinite and fA stands for a nominal of the infinite conjunction.

Since L_{DFT} is the standard modal language of S5 augmented with f, the $S5$-possibility operator $\Diamond$ and the strict conditional $\Rightarrow$ are defined in terms of $\Box$ as usual:

$$\Diamond A \equiv_{\mathrm{df}} \neg\Box\neg A \text{ and } A \Rightarrow B \equiv_{\mathrm{df}} \Box(A \supset B).$$

Definition 7.4 (logic DFT, [Alchourrón, 1995]) The conditional logic DFT is the smallest set $S \subseteq L_{DFT}$ such that S contains classical propositional logic and the following axiom schemata, and is closed under the following rules of inference:

- **K** $\Box(A \supset B) \supset (\Box A \supset \Box B)$.
- **T** $\Box A \supset A$.
- **4** $\Box A \supset \Box\Box A$.
- **5** $A \supset \Box\Diamond A$.
- **f.1** $(fA \supset A)$. (Expansion)
- **f.2** $(A \equiv B) \supset (fA \equiv fB)$. (Extensionality)
- **f.3** $\Diamond A \supset \Diamond fA$. (Limit Expansion)
- **f.4** $(f(A \vee B) \leftrightarrow fA) \vee (f(A \vee B) \leftrightarrow fB) \vee (f(A \vee B) \leftrightarrow (fA \vee fB))$ (Hierarchical Ordering)
- **Nes** From A infer $\Box A$.
- **MP** From $A \supset B$ and A infer B.

Axioms **K**, **T**, **4** and **5** giver rise to S5, and **f.1**-**f.4** are constraints imposed on the revision operator f. Condition **f.1** is in fact the characteristic axiom T of standard modal systems. As an axiom constraining f it is quite natural since it states that fA stands for the conjunction of A and its presuppositions. **f.2** asserts that equivalent sentences have equivalent presuppositions. **f.3** links the two modalities. It ensures the existence of consistent presuppositions for any sentence that is not a contradiction. We will see below that condition **f.3** carries some consequences that we will analyze in semantic terms. **f.4** asserts that the presuppositions of a disjunction are either the presuppositions of one of the disjuncts, or else the disjunction of the presuppositions of each of the disjuncts. In a forward reading it implies that f is a normal modality, in the sense that it satisfies the characteristic axiom K (notice that $\vdash_{DFT} f(\neg A) \supset \neg(fA)$).

Alchourrón gives a formal semantic interpretation of the language L_{DFT} based on standard non-relational S5-models.

Definition 7.5 (DFT model) A model for L_{DFT} is $M_{DFT} = \langle W, Ch, [\,]\rangle$ where W is a non-empty set of possible worlds, the valuation function $[\,]$ maps P into $\mathcal{P}(W)$, and $Ch : L \to \mathcal{P}(W)$ is a selection function such that for each sentence A, B of L_{DFT}

Ch.1 $Ch(A) \subseteq [A]$.

Ch.2 If $[A] = [B]$ then $Ch(A) = Ch(B)$.

Ch.3 If $[A] \neq \emptyset$ then $Ch(A) \neq \emptyset$.

Ch.4 $Ch(A \vee B) \in \{Ch(A), Ch(B), Ch(A) \cup Ch(B)\}$.

We shall mention that [Alchourrón, 1995] defines the selection function as Ch^{α} meaning that the selection is indexed by the particular preferences of an individual α (as opposed to be a universal selection function for every individual). For the purposes of this note this is an irrelevant restriction. The selection function Ch is proposed as the semantic counterpart of the syntactic revision operator. $Ch(A)$ is the proposition of the joint assertion of A and its assumptions, i.e., the worlds in which fA are true.

$$[fA] = Ch(A).$$

The four constraints on Ch are in exact correspondence with the four on f. In particular, **Ch.3** reflects that every consistent proposition must contain some chosen elements.

A DFT frame $\langle W, Ch\rangle$ is the set of all DFT models having W and Ch. Satisfaction of a modal formula at world w in a model $M_{DFT} = \langle W, Ch, [\,]\rangle$ is given by:

$M_{DFT} \models_w A$ iff $w \in [A]$ for atomic sentence A.

$M_{DFT} \models_w \neg A$ iff not $M_{DFT} \models_w A$.

$M_{DFT} \models_w A \wedge B$ iff $M_{DFT} \models_w A$ and $M_{DFT} \models_w B$.

$M_{DFT} \models_w \Box A$ iff $[A] = W$.

$M_{DFT} \models_w fA$ iff $w \in Ch(A)$.

The derived satisfaction conditions for the connectives $\Diamond$ and $\Rightarrow$ are:

$M_{DFT} \models_w A \Rightarrow B$ iff $[A] \subseteq [B]$.

$M_{DFT} \models_w \Diamond A$ iff there is some $v \in W$ such that $v \in [A]$.

Truth in a model $M_{DFT} = \langle W, Ch, [\ ]\rangle$ is truth at every point:

$M_{DFT} \models A$ iff $M_{DFT} \models_w A$ for every $w \in W$.

Truth in a frame $\langle W, Ch\rangle$ is truth at every model $\langle W, Ch, [\ ]\rangle$.

$\langle W, Ch\rangle \models A$ iff $\langle W, Ch, [\ ]\rangle \models A$ for all valuation functions $[\ ]$.

Alchourrón proves that his semantic and axiomatic presentations coincide.

Observation 7.6 ([Alchourrón, 1995], Theorem cm-DFT)
For any $A \in L_{DFT}$, $\vdash_{\mathrm{DFT}} A$ iff $\models_{\mathrm{DFT}} A$.

We are ready for the definition of the conditional $A >_{DFT} B$. Alchourrón wants to capture the idea that the antecedent A jointly with the set of assumptions that comes with it is a sufficient condition for the consequent B. To reflect this intuition, Alchourrón adopts the following definition due to Lennart Åquist [1973]:

Definition 7.7 (DFT conditional connective) $A >_{DFT} B \equiv_{\mathrm{df}} \Box(fA \supset B)$.

Satisfaction of a conditional sentence at world w in a model $M_{DFT} = \langle W, Ch, [\ ]\rangle$ is given by: $M_{DFT} \models_w A >_{DFT} B$ iff $Ch(A) \subseteq [B]$ iff $M_{DFT} \models A >_{DFT} B$. As a result, $M_{DFT} \models A >_{DFT} B$ iff $[A >_{DFT} B] = W$. Conversely, $M_{DFT} \not\models (A >_{DFT} B)$ iff $[A >_{DFT} B] = \emptyset$ iff $M_{DFT} \models \neg(A >_{DFT} B)$. This means that Alchourrón conditionals are true at every point in a model, or at none.

Observation 7.8 ([Alchourrón, 1995])
$\vdash_{DFT} (A >_{DFT} B) \supset \Box(A >_{DFT} B)$ and $\vdash_{DFT} \neg(A >_{DFT} B) \supset \Box\neg(A >_{DFT} B)$.

In DFT $>_{DFT}$ is in general different from $\Rightarrow$.

Observation 7.9 $\vdash_{DFT} A \Rightarrow B \supset A >_{DFT} B$ but $\not\vdash_{DFT} A >_{DFT} B \supset A \Rightarrow B$.

PROOF. To prove that $\vdash_{DFT} A \Rightarrow B \supset A >_{DFT} B$, assume $\vdash_{DFT} A \Rightarrow B$. Then for every DFT-model $[A] \subseteq [B]$. Since $Ch(A) \subseteq [A]$, then $Ch(A) \subseteq [B]$, hence, $\vdash_{DFT} A >_{DFT} B$.

To prove $\not\vdash_{DFT} A >_{DFT} B \supset A \Rightarrow B$, suppose $\vdash_{DFT} A >_{DFT} B$. Hence, for every DFT model $Ch(A) \subseteq [B]$. In particular for $M_{DFT} = \langle, W, Ch, [\,]\rangle$ such that W is the set of valuations of the language based on two propositional variables, A, B. Suppose $\{w_1, w_2\} = [A]$ and $\{w1, w3\} = [B]$ and $Ch(A) = \{w_1\}$ provides a model where $M_{DFT} \models A >_{DFT} B$ and $M_{DFT} \not\models A \Rightarrow B$. QED

This proof also shows that $>_{DFT}$ in DFT does not validate Modus Ponens nor Contraposition. And similarly, with three propositional letters can be shown that $>_{DFT}$ does not validate Strengthening the antecedent nor Transitivity.

Modus Ponens From $A > B$ and A infer B.

Strengthening From $A > B$ infer $A \wedge C > B$.

Transitivity From $A > B$ and $B > C$ infer $A > C$.

Contraposition From $A > B$ infer $\neg B > \neg A$.

As a corollary of the observation above we obtain that in a limiting case $>_{DFT}$ and $\Rightarrow$ are equivalent. In the particular case where the Choice function sanctions $Ch(A) = [A]$ for every $A \in L_{DFT}$, $>_{DFT}$ collapses with $\Rightarrow$. In this case the Choice function induces an ignorant revision function f, where every sentence becomes its own presupposition. Then, the conditional $>_{DFT}$ looses all its defeating properties.

Alchourrón also gives a gives an axiomatic presentation of his logic DFT, in a purely conditional language, having the conditional connective $>$ added to those of classical propositional logic. Let's denote this language by $L^{>}$. The following abbreviations are used in Alchourrón's axiomatisation. A notion of necessity N, a notion of possibility M and a notion of comparativeness $\succeq$.

$$NA \equiv_{\rm df} \neg A > \bot;\; MA \equiv_{df} \neg N \neg A;$$
$$(A \succeq B) \equiv_{\rm df} (N(\neg A \wedge \neg B)) \;\vee\; \neg((A \vee B) > \neg A)).$$

Definition 7.10 ([Alchourrón, 1995]) The conditional logic DFT$_{>}$ is the smallest set $S \subseteq L^{>}$ such that S contains classical propositional logic and the following axiom schemata, and is closed under the following rules of inference:

DFT1 $\vdash_{>} (A > A)$.

DFT2 $\vdash_{>} (A > (B \wedge C)] \equiv [A > B) \wedge (A > C))$.

DFT3.1 $\vdash_{>} ((A > C) \wedge (B > C)) \supset ((A \vee B) > C)$.

DFT3.2 $\vdash_{>} (A \succeq B) \supset ((A \vee B) > C) \supset (A > C)$.

DFT4 $\vdash_{>} (A > B) \supset N(A > B)$.

DFT5 $\vdash_{>} \neg(A > B) \supset N\neg(A > B)$.

DFT6 $\vdash_{>} NA \supset A$.

Ext If $\vdash_{>} A \equiv B$ then $\vdash_{>} (A > C) \equiv (B > C)$ and $\vdash_{>} (C > A) \equiv (C > B)$.

In this purely conditional axiomatization it is also apparent that a conditional sentence is always impossible or necessary (this is directly entailed by DFT.4 and DFT.5). Alchourrón shows the following correspondence between DFT and DFT$_{>}$. Let Ψ be a translation function from $L^{>}$ to L_{DFT}.

$\Psi(A) = A$, if A is a propositional variable.

$\Psi(\top) = \top$ and $\Psi(\bot) = \bot$.

$\Psi(\neg A)) = \neg\Psi(A)$.

$\Psi(A \wedge B) = \Psi(A) \wedge \Psi(B)$.

$\Psi(A > B) = \Box(f\Psi(A) \supset \Psi(B))$.

Alchourrón proves that the logic DFT$_{>}$ is properly embedded in DFT.

Observation 7.11 ([Alchourrón, 1995], Corr.3) For every $A \in L^{>}$, $\vdash_{>} A$ iff $\vdash_{DFT} \Psi(A)$.

Since the translation Ψ is not surjective on L_{DFT}, that is, there are formulae of L_{DFT} which are not equivalent to the image of any formula of $L^{>}$, then the expressive power of DFT exceeds that of DFT$_{>}$.

We end up this section with a final remark. Some (infinite) sets of conditional sentences in $L^{>}$ define single DFT models. Let $\Gamma \subseteq L^{>}$ such that for every purely propositional $A, B \in L$, either $A > B \in \Gamma$ or $\neg(A > B) \in \Gamma$ but not both. Such a Γ characterises a single DFT-model. We will return to this idea when we study how DFT provides a logical calculus for theory change. We shall now turn our attention into Boutilier's logic CO.

7.3 The Logic CO

The logic CO is one in Boutilier's family of conditional logics for theory change and default reasoning [Boutilier, 1992a]. He bases his logics on Humberstone's bimodal logic [Humberstone, 1983], which provides a modality that denotes truth along an accessibility relation and another modality that denotes truth along the complement of the accessibility relation. The expressive power of this bimodal logic exceeds that of standard mono modal systems. For instance, it can express a number of relational properties that are inexpressible in standard modal logics, like total connectedness, asymmetry and irreflexivity. Humberstone's logic is closely related to temporal logics, which are also based on two modalities. In temporal logics the modality for the "future" coincides with Humberstone's modality for denoting truth along the accessibility relation R. However, the temporal operator for the "past" denotes truth along the *inverse* of relation R, while in Humberstone's logic the second modality denotes truth along the *complement* of R.

Humberstone presented his logic as an enumerable set of axioms, and left open the question of whether a finite axiomatization existed [Humberstone, 1983]. Boutilier [Boutilier, 1992a] provided the sought finite axiomatization.

The language L_{CO} is defined as a propositional language L augmented with two modal operators. $\overrightarrow{\Box}$ is the modality for accessibility along a relation R and $\overleftarrow{\Box}$ is the modality for inaccessibility, denoting truth along the complement of relation R. Since Boutilier's conditional connective is only an abbreviation of an involved formula in the bimodal language, the expressive power of CO is precisely that of Humberstone's. Boutilier defines several connectives in terms of the primitive $\overrightarrow{\Box}$ and $\overleftarrow{\Box}$ as follows:

$\overrightarrow{\Diamond}A \equiv_{\mathrm{df}} \neg\overrightarrow{\Box}\neg A$;

$\overleftarrow{\Diamond}A \equiv_{\mathrm{df}} \neg\overleftarrow{\Box}\neg A$;

$\overleftrightarrow{\Box}A \equiv_{\mathrm{df}} \overrightarrow{\Box}A \wedge \overleftarrow{\Box}A$; and

$\overleftrightarrow{\Diamond}A \equiv_{\mathrm{df}} \neg\overleftrightarrow{\Box}\neg A$.

Definition 7.12 (logic CO [Boutilier, 1992a]) The conditional logic CO is the smallest set $S \subseteq L_{CO}$ such that S contains classical propositional logic and the following axiom schemata, and is closed under the following rules of inference:

K $\overrightarrow{\Box}(A \supset B) \supset (\overrightarrow{\Box}A \supset \overrightarrow{\Box}B)$.

K' $\overleftarrow{\Box}(A \supset B) \supset (\overleftarrow{\Box}A \supset \overleftarrow{\Box}B)$.

4 $\overrightarrow{\Box}A \supset \overrightarrow{\Box}\overrightarrow{\Box}A$.

S $A \supset \overleftarrow{\Box}\overrightarrow{\Diamond}A$.

H $\overleftrightarrow{\Diamond}(\overrightarrow{\Box}A \wedge \overleftarrow{\Box}B) \supset \overleftrightarrow{\Box}(A \vee B)$.

Nes From A infer $\overleftrightarrow{\Box}A$.

MP From $A \supset B$ and A infer B.

Axioms **K** and **K'** indicate that the two modalities are normal. Axiom **4** ensures transitivity of the accessibility relation and axiom **S**, which is only expressible in a bimodal language, ensures total connectedness. Axiom **H** gives the relationship between the two modalities.

CO is sound and complete with respect to S4.3 structures, the structures whose relations are total preorders.

Definition 7.13 (CO-model, [Boutilier, 1992a]) A CO -model is a triple $\mathcal{M}_{CO} = \langle W, R, [\ \ ] \rangle$ where W is a set of worlds, with valuation function $[\,] : P \rightarrow \mathcal{P}(W)$, and R is a total preorder on W.

Satisfaction at world w in a model $\mathcal{M}_{CO} = \langle W, R, [\ \ ] \rangle$ is given by:

$\mathcal{M}_{CO} \models_w A$ iff $w \in [A]$ for atomic sentence A.

$\mathcal{M}_{CO} \models_w \neg A$ iff $\mathcal{M}_{CO} \not\models_w A$.

$\mathcal{M}_{CO} \models_w \overrightarrow{\Box}A$ iff for each v such that wRv, $\mathcal{M}_{CO} \models_v A$.

$\mathcal{M}_{CO} \models_w \overleftarrow{\Box}A$ iff for each v such that not wRv, $\mathcal{M}_{CO} \models_v A$.

The derived connectives have the following truth conditions: $\overrightarrow{\Diamond}$ ($\overleftarrow{\Diamond}A$) is true at a world if A holds at some accessible (inaccessible) world; $\overleftrightarrow{\Box}A$ ($\overleftrightarrow{\Diamond}A$) holds iff A holds at all (some) worlds. Therefore, the $\overleftrightarrow{\Box}$ and $\overleftrightarrow{\Diamond}$ modalities behave as S5 modalities.

Truth in a model $\mathcal{M}_{CO} = \langle W, R, [\ \ ] \rangle$ and in a frame $\langle W, R \rangle$ are defined as usual.

$\mathcal{M}_{CO} \models A$ iff $\mathcal{M}_{CO} \models_w A$ for every $w \in W$.

$\langle W, R \rangle \models A$ iff $\langle W, R, [\,] \rangle \models A$, for every $[\,]$.

The system CO is characterized by the class of CO-models.

Theorem 7.14 ([Boutilier, 1992a]) $\vdash_{CO} A$ iff $\models_{CO} A$.

The conditional connective is defined in the bimodal language as follows.

Definition 7.15 (conditional connective in CO , [Boutilier, 1992a]) $(A >_{CO} B) \equiv_{\mathrm{df}} \overleftarrow{\Box}(A \supset \overleftarrow{\Diamond}(A \wedge \overleftarrow{\Box}(A \supset B)))$

The conditional $A >_{CO} B$ holds in a model when either there are no A worlds at all, or, when every A-world has access to some point where every R-accessible world satisfying A also satisfies B. The conditional $A >_{CO} B$ states that the (possibly infinite) chain of R-minimal A-worlds must satisfy B. Boutilier does not assume the existence of *the minimal* A-worlds. In the case where such worlds do exist, obviously $A >_{CO} B$ holds just when B holds at all such worlds. In contrast, suppose there is some unending chain of R-minimal A-worlds. If some B-world lies in this chain having the property that B-holds whenever A does, at all farther accessible worlds in the infinite descending chain, then $A >_{CO} B$ ought to be considered true. B would hold at the *hypothetical limit* of A-worlds in this chain. This is the same truth conditions that Lewis' [Lewis, 1973] has imposed to his counterfactuals conditionals in models that do not comply the limit assumption.

Boutilier argues against the limit assumption. He explains that without the limit assumption a selection function fails and, vacuously, makes all conditionals true. But certainly some conditionals should remain true and some others false. Since CO makes no commitment to the limit assumption this is a point in which the Boutilier's and Alchourrón's formalisms differ. A proper subclass of CO-models is that of models whose accessibility relation satisfies the limit assumption. Since the limit assumption is not expressible in CO, this class cannot be syntactically characterized in the bimodal language. In models that satisfy the limit assumption it is possible to define the set of R-minimal A-worlds.

Definition 7.16 (min) Let $\mathcal{M}_{CO}$ a CO-model satisfying the limit assumption. We define $\min : \mathcal{L}_{CO} \to \mathcal{P}(W)$ as:
$\min(A) = \{w \in W : \mathcal{M}_{CO} \models_w A \text{ and } \mathcal{M}_{CO} \models_v \overleftarrow{\Diamond} A \text{ implies } wRv \text{ for all } v \in W\}.$

When dealing with CO-models that comply the limit assumption, $A >_{CO} B$ is true in a model $\mathcal{M}_{CO}$ just when B is true at each of the R-minimal A-worlds. The definition of a conditional can be expressed semantically as follows:

$$\mathcal{M}_{CO} \models A >_{CO} B \text{ iff } \min(A) \subseteq [B].$$

We are ready to compare logic DFT and CO and reveal their connection.

7.4 The Connection between DFT and CO

Let's first prove the correspondence between DFT models and CO models using the results in section 7.1. Let's start specifying how a CO-model $\mathcal{M}_{CO} = \langle W, R, [\ ]\rangle$ satisfying the limit assumption induces a Choice function Ch_R.

Definition 7.17 (Ch_R) Let a CO-model $\mathcal{M}_{CO} = \langle W, R, [\ ]\rangle$ that satisfies the limit assumption, and let $A \in L$. The Choice function Ch_R induced by the accessibility relation R is defined as:

$$Ch_R(A) = \{w \in [A] : wRw',\ \forall w' \in [A]\}.$$

To discover the properties of Ch_R we want to apply lemma 7.3. As R is a total preorder satisfying the limit assumption, then strict part of R is well founded, R is transitive and the complement of R is also transitive. (To see this last property suppose not xRy and not yRz but xRz. Since R is connected, then it must be zRy. Thus, by the transitivity of R we obtain xRy contrary to our assumption.) Hence, by lemma 7.3 Ch_R satisfies (I), (II), (III), and (IV). We have to check now that Ch_R validates Ch.1-Ch.4, the characteristic properties of Alchourrón's choice functions.

Proposition 7.18 Ch_R satisfies the following properties:

(Ch.1) $Ch_R(A) \subseteq [A]$.

(Ch.2) If $[A] = [B]$ then $Ch_R(A) = Ch_R(B)$.

(Ch.3) If $[A] \neq \emptyset$ then $Ch_R(A) \neq \emptyset$.

(Ch.4) $Ch_R(A \vee B) \in \{Ch_R(A), Ch_R(B), Ch_R(A) \cup Ch_R(B)\}$.

PROOF. That Ch_R satisfies Ch.1 and Ch.2 is obvious by definition 7.17.

To see Ch.3 suppose $Ch_R(A) = \emptyset$. Then, there is no $w \in [A]$ such that wRw' for all $w' \in [A]$. Since R satisfies the limit assumption, $[A] = \emptyset$.

Let's see Ch.4. Let $X = [A]$ and $Y = [B]$. There are four cases.

(1) If $Ch_R(X \cup Y) \cap X = \emptyset$ and $Ch_R(X \cup Y) \cap Y = \emptyset$ then, by case Ch.1 above, $X \cup Y = \emptyset$, and Ch.4 trivially holds.

(2) Assume $Ch_R(X \cup Y) \cap X \neq \emptyset$ and $Ch_R(X \cup Y) \cap Y \neq \emptyset$. By postulate (I) $Ch_R(X \cup Y) \subseteq Ch_R(X) \cup Ch_R(Y)$. By postulate (IV) $Ch_R(X) \subseteq Ch_R(X \cup Y)$ and $Ch_R(X) \subseteq Ch_R(X \cup Y)$. By de Morgan laws, $Ch_R(X) \cup Ch_R(Y) \subseteq Ch_R(X \cup Y)$. Thus $Ch_R(X) \cup Ch_R(Y) = Ch_R(X \cup Y)$, and Ch.4 holds.

(3) Assume $Ch_R(X \cup Y) \cap X \neq \emptyset$ and $Ch_R(X \cup Y) \cap Y = \emptyset$. By postulates (III) and (IV) $Ch_R(X) \subseteq Ch_R(X \cup Y)$. By postulate (II) $Ch_R(X) \cap Ch_R(X \cup Y) \subseteq Ch_R(X \cup X \cup Y) = Ch_R(X \cup Y)$. And by postulate (I) $Ch_R(X \cup Y) \subseteq Ch_R(X) \cup Ch_R(Y)$. Since by Ch.1 $Ch(Y) \subseteq Y$, and by assumption of Ch.3 $Ch_R(X \cup Y) \cap Y = \emptyset$, then $Ch_R(X \cup Y) \subseteq Ch_R(X)$. Hence $Ch_R(X) \subseteq Ch_R(X \cup Y) \subseteq Ch_R(X)$; namely, $Ch_R(X) = Ch_R(X \cup Y)$, and Ch.4 is verified.

(4) The case $Ch_R(X \cup Y) \cap X = \emptyset$ and $Ch_R(X \cup Y) \cap Y \neq \emptyset$ is analogue to case (3) above. QED

Now let's see how a DFT-model $\mathcal{M}_{DFT} = \langle W, Ch, [\,] \rangle$ induces a total preorder R_{Ch} on W and gives rise to a CO-model $\mathcal{M}_{CO} = \langle W, R_{Ch}, [\,] \rangle$.

Definition 7.19 (R_{Ch}) Let $\mathcal{M}_{DFT} = \langle W, Ch, [\,] \rangle$ with $Ch : L \rightarrow \mathcal{P}(W)$. The relation R_{Ch} induced by Ch is defined as follows.

$$R_{Ch} = \{(w, v) \in W \times W : \exists Y \in \mathcal{P}(W) \text{ such that } w, v \in Y \text{ and } w \in Ch(Y)\}$$

We have to check R_{Ch} is a total preorder on W. Lemma 7.1 states that if Ch is additive or 12-covering and satisfies (I) and (II) then there exists some some relation $\leq$ on W such that the selection function induced by $\leq$ coincides with Ch. But Ch over an infinite set of propositional variables is not additive nor 12-covering, so the we can't apply the lemma. As suggested in section 7.1 this is the problem we face when dealing with infinite languages. Let's consider L an infinite propositional language, W the set of all its maximal consistent extensions and $\mathcal{X} \subseteq \mathcal{P}(W)$ the set of all the L-nameable subsets of W.

A preorder $R \subseteq W \times W$ automatically determines a preorder relation over *every* subset of W, that is, $\forall X \subseteq W, R \cap X \times X$ is a relation on X. In contrast, Alchourrón's choice function is intrinsically linguistic, that is, it is defined from L to subsets of W. Hence, Ch provides a selection just for nameable subsets of W. The cardinality of L is less than the cardinality of $\mathcal{P}(W)$ so it is impossible to provide a one to one correspondence between binary relations on W and linguistic selection functions. In order to establish a one to one correspondence we have to be able to name all subsets of W. With this objective we will restrict to propositional languages based on finite sets of propositional variables. Thus, Ch becomes additive and 12-covering, and we can apply lemma 7.2.

Proposition 7.20 R_{Ch} is a total preorder on W.

PROOF. We apply lemma 7.2. Since Ch is 12-covering and satisfies (I) the complement of R_{Ch} is acyclic. Since Ch is 123-covering and satisfies (I) and (II) R_{Ch} is transitive. Since Ch is finitely additive and satisfies (IV) the complement of R_{Ch} is transitive.

That R_{Ch} is totally connected follows from acyclicity of the complement of R_{Ch}, (if not $xR_{Ch}y$ and not $yR_{Ch}x$ then the complement R_{Ch} would not be acyclic). QED

Let's check that in the finite case $Ch_{R_{Ch}} = Ch$ and $R_{Ch_R} = R$.

Observation 7.21 Given a finite propositional language L, $Ch_{R_{Ch}} = Ch$ and $R_{Ch_R} = R$.

PROOF. Assume $R \subseteq W \times W$, a total preorder. Let's define $Ch_R(A) = \{w \in [A] : wRw',\ \forall w' \in [A]\}$. This is additive n-covering choice function satisfying (I)-(IV).

$R_{Ch_R} = \{(w, v) \in W \times W : \exists Y \in \mathcal{P}(W)$ such that $w, v \in Y$ and $w \in Ch_R(Y)\}$. By lemma 7.1, directly $Ch_{R_{Ch}} = Ch$.

Let's see that $R = R_{Ch_R}$. Suppose wRv and not vRw. Then, there is some $A \in w \cap v$ and some $B \in w$ $B \notin v$ such that $w \in Ch_R(A \vee B)$ and $v \notin Ch_R(A \vee B)$. Hence $wR_{Ch_R}v$ and not vRw.

Suppose wRv and vRw. Then, for every $A \in L$ such that $A \in w \cap v$, $w, v \in Ch_R(A)$. Hence, $wR_{Ch_R}v$ and $vR_{Ch_R}w$.

Therefore, $\forall w, v, wRv$ iff $wR_{Ch_R}v$. QED

Consequently finite DFT models and finite CO models (with the same universe set W and the same valuation function $[\,]$) are in a one to one correspondence. Our next result is that the semantic definitions of the conditional connectives in CO and DFT coincide.

Observation 7.22 Let A, B propositional formulae of L, then,
$\langle W, R, [\,]\rangle \models A >_{CO} B$ iff $\langle W, Ch_R, [\,]\rangle \models A >_{DFT} B$.

PROOF. Let $CrM = \langle W, R, [\,]\rangle$.

By the definition of the conditional connective in CO,

$\langle W, R, [\,]\rangle \models A >_{CO} B$ iff $\min_R(A) \subseteq [B]$ iff,

by observation 7.21 $Ch_R(A) \subseteq [B]$ iff,

by definition of the conditional connective in DFT,

$\langle W, Ch_R, [\,]\rangle \models A >_{DFT} B$. QED

We are now able to state our main result, which reveals the connection between the two logics: the two logics validate the *same* conditional sentences in a restricted language. Let's define $L^{>}_{DFT}$ and $L^{>}_{CO}$ as the propositional languages formed from a finite set P of propositional variables together with the connectives $\neg, \wedge$ augmented solely with the respective conditional connective $>_{DFT}$ and $>_{CO}$ (the connectives $\supset, \vee \equiv$ are defined in terms of $\neg, \wedge$ as usual).

We will define a bijective translation function taking a sentence in $L^{>}_{DFT}$ and returning a sentence in $L^{>}_{CO}$. We will then prove that this bijective translation preserves satisfiability in the two logics. As a result will be able to assert that there is a one to one correspondence of valid sentences in the respective restricted languages in the two logics, with exactly the same interpretation. Since the translation just interchanges the respective conditional connectives the two logics validate the *same* conditional sentences. Let Ψ be a translation function from $L^{>}_{DFT}$ to $L^{>}_{CO}$.

$\Psi(A) = A$, if A is a propositional variable.

$\Psi(\top) = \top$ and $\Psi(\bot) = \bot$.

$\Psi(\neg A)) = \neg\Psi(A)$.

$\Psi(A \wedge B) = \Psi(A) \wedge \Psi(B)$.

$\Psi(A >_{DFT} B) = \Psi(A) >_{CO} \Psi(B)$.

Let's remark that Ψ is a bijective translation function.

Theorem 7.23 $\vdash_{DFT} A$ iff $\vdash_{CO} \Psi(A)$.

PROOF. Suppose $A \in L^{>}_{DFT}$ such that not $\vdash_{DFT} A$.

Given that CO and DFT are sound and complete with respect to their respective classes, there is a DFT model $\langle W, Ch, [\,]\rangle$ where A is not true.

By observation 7.22 there is a CO model $\langle W, R_{Ch}, [\,]\rangle$ where $\Psi(A)$ is not true. QED

We obtain the following corollary. For any set of sentences $X \subseteq L^{>}_{DFT}$, let's define the translation of X as $\Psi(X) = \{B \in L^{>}_{CO} : B = \Psi(A) : A \in X\}$. Then, $Mods(X) \models A$ iff $X \vdash_{DFT} A$ iff $\Psi(X) \vdash_{CO} \Psi(A)$ iff $Mods(\Psi(X)) \models \Psi(A)$.

We have proved that in the respective restricted languages the theorems of CO and DFT are in a one to one correspondence, and have the same interpretation. But this correspondence only holds in the restricted languages, that is, the two logics are not equivalent as a whole. For instance, in DFT there is no counterpart of the CO modalities for accessibility and inaccessibility. A question still to be answered in this direction is whether the revision operator f of DFT is expressible in CO. It is clear that the expressive power of DFT extends that of $S5$ without being exactly clear what is the expressivity added by the "revision function" f. The study of the formal properties that become expressible in DFT that are inexpressible in standard systems is an interesting issue that remains to be investigated.

7.5 A Logical Calculus for Theory Change

Conditional logics were initially developed for modeling "if ... then" statements in natural language. Robert Stalnaker [1968] gives a possible worlds semantics for his logic for "subjunctive conditionals". A conditional $A > B$, read as "if A were true B would be true". Stalnaker argues that the conditional connective $>$ should not validate transitivity, nor the strengthening rule, nor contraposition. For instance, we accept the conditional "If this match were struck, it would light", while we deny that "If this match were wet and struck, it would light".

Stalnaker gives the following "recipe" based on the *Ramsey test* to evaluate a conditional in a given theory or state of belief:

> "First, add the antecedent (hypothetically) to your stock of beliefs; second, make whatever adjustments are required to maintain consistency (without modifying the hypothetical belief in the antecedent); finally, consider whether or not the consequent is then true."(Stalnaker 1968, page 44)

Stalnaker's formulation of the Ramsey test has been used to provide a formal connection between theory change and conditional logic.

A conditional $A > B$ is true iff B belongs to the revision of K by A.

Based on this formulation Boutilier provides a logical calculus for AGM revision [1992a].

$$M_{CO} \models A >_{CO} B \text{ is equated with } B \in K * A.$$

Given the Ramsey test, $>_{CO}$ is nothing more than a subjunctive conditional, interpreted as "If K were revised by A, then B would be accepted". For any propositional A, the theory resulting from revision of K by A is:

$$K * A = \{B \in L : M_{CO} \models A >_{CO} B\}.$$

Since total preorders on W satisfying the limit assumption are isomorphic to Grove's systems of spheres with no empty center CO-models are appropriate for AGM revision, when the theory K being revised is assumed to be a propositional theory. By appealing to Grove's result [Grove, 1988] for representing revision functions, each CO-model satisfying the limit assumption represents a revision function. Those worlds consistent with K should be *exactly* those minimal in R. The interpretation of R is as follows: wRv iff v is as close to theory K as w.

$$\forall w \in [K], \forall v \in W, \quad wRv.$$

CO models that satisfy this constraint are called *revision models* for K.

Definition 7.24 ([Boutilier, 1992a]) A revision model for theory K is any structure $M_{CO} =< W, R, [\,] >$ such that R satisfies the limit assumption, R is transitive and totally connected on W and $v \in \{w : W \models_w A$ for all $A \in K\}$ iff v is R-minimal in $\mathcal{M}_{CO}$.

Full models are those where all propositional valuations are represented. They have to to be considered in order to allow every consistent sentence be capable of generating a consistent revision. Boutilier proves that the revision function determined by a full revision model for K satisfies the eight AGM postulates for revision (K*1)-(K*8).

Observation 7.25 ([Boutilier, 1992a], Theorem 6.7) Let $\mathcal{M}_{CO}$ be a full revision model for K and $*^M$ the revision function determined by M. $*^M$ is defined for each $A \in L_{CPL}$ by $K *^M A = \{B \in L : M \models A >_{CO} B\}$. Then, $*^M$ satisfies postulates (K*1)-(K*8).

Boutilier defines a modality Bel_{CO} to refer to the sentences in K. $\mathsf{Bel}_{CO} A$ is read as A is accepted in K. He calls it a modality for belief. The modality Bel_{CO} is defined as follows.

Definition 7.26 ([Boutilier, 1992a]) $\mathsf{Bel}_{CO} A \equiv_{\mathsf{df}} \overleftarrow{\Box}\overrightarrow{\Diamond}\overrightarrow{\Box} A$.

The sentence $\mathsf{Bel}_{CO}(A)$ holds in a revision model when A is true at each minimal worlds: $\mathcal{M}_{CO} \models \mathsf{Bel}_{CO}(A)$ iff $\min(\top) \subseteq [A]$ iff $\mathcal{M}_{CO} \models \top >_{CO} A$.

By appealing to the derivability in the logic CO it is possible to calculate the results of revising a theory K. Each set of conditional sentences $\Gamma \subseteq L^{>}_{CO}$ such that $Mods(\Gamma)$ is a singleton represents a theory K and the AGM revision function $*$ for K. For instance, $K = \{A \in L : \top >_{CO} A \in \Gamma\}$. Then, if Γ is conditionally complete then $Mods(\Gamma) = \{\mathcal{M}_{CO}\}$. So, we obtain the following chain of equivalences:

$$\Gamma \vdash_{CO} A >_{DFT} B \text{ iff } \mathcal{M}_{CO} \models A >_{CO} B \text{ iff } \min(A) \subseteq [B] \text{ in } \mathcal{M}_{CO} \text{ iff } B \in K * A.$$

In this way the logic CO provides a logical calculus for change in different theories, by appealing to derivability from different sets Γ_1 and Γ_2. Given the correspondence we have proved between CO and DFT, all the considerations about CO as a logical calculus for theory change directly apply to DFT.

One could wonder about calculating iterated change in logic CO. It is possible to use the Ramsey test to relate iterated changes and acceptance of nested conditionals [Levi, 1988; Boutilier, 1992b; Lindström and Rabinowicz, 1992].

Different intuitions correspond to whether the nesting of the conditional connective appears in the antecedent or in the consequent of a conditional construction. One could inspect whether $(A > (B > C)$ can be taken to mean that $C \in K * A * B$. This would require two applications of the Ramsey test.

$$A > (B > C) \text{ is true iff } (B > C) \in K * A \text{ iff } C \in K * A * B.$$

But the nested occurrences of the conditional connective collapse into the flat portions of CO and DFT, as follows. Given a revision model $\mathcal{M}_{CO}$, we have that $\mathcal{M}_{CO} \models A > (B > C)$ is identical to either $\mathcal{M}_{CO} \models A > \top$ or $\mathcal{M}_{CO} \models A > \bot$, depending whether $\mathcal{M}_{CO} \models B > C$ or not. On the one hand $\mathcal{M}_{CO} \models A > \top$ is always true, because for all A, $min(A) \subseteq W$ and if A is not satisfiable then $\min(A) = \emptyset$. On the other hand $\mathcal{M}_{CO} \models A > \bot$ is always false unless $\mathcal{M}_{CO} \models \neg A$. In full revision models this means that A is not satisfiable. Consequently, $A > B > C$ is true iff $(B > C)$ is true or A is itself inconsistent. This is equated via the Ramsey test as $C \in K * A * B$ iff A is inconsistent or $C \in K * B$. Hence, for consistent formulae A $A > (B > C)$ says that the set $K * A * B$, is just $K * B$. The notion of iteration it yields validates the following postulate that we advanced in Chapter 3 for a trivial revision function.

(T)

$$K \circ A \circ B = \begin{cases} K \circ B & \text{, if } A_1, \ldots A_{n-1} \text{ are satisfiable.} \\ L & \text{, otherwise.} \end{cases}$$

We conclude that nested conditionals in CO or DFT do not provide an interesting logical calculus for iterated change. This conclusion can also be reached from our interpretation of CO models as models for revision. If $\mathcal{M}_{CO}$ is a revision model for K, then all the worlds consistent with K are R-minimal in $\mathcal{M}_{CO}$. Therefore, $\mathcal{M}_{CO}$ is just a revision model for K and in general it is not a revision model for $K * A$.

Let's analyze no the case when the nested conditional connective appears in the antecedent of a conditional construction. Again following the Ramsey test, the conditional $(A > B) > C$ is equated with C being accepted in the theory resulting by the revision of K by the conditional sentence $(A > B)$. But such a revision would collapse with our initial assumptions about how CO provides a calculus for revision. A CO model induces an AGM revision function $*^M$

for a propositional theory K, such that $*^M$ satisfies (K*1)-(K*8). Since (K*2) requires $(A > B) \in K$ then K should contain conditional sentences, contrary to our assumption that K is propositional. The so called "Gärdenfor's triviality theorem" or "impossibility theorem" [Gärdenfors, 1986; Rott, 1989] shows that the AGM revision operation becomes trivial when it is applied to conditional theories whose conditional sentences are interpreted with the Ramsey test. No sound conclusions about iterated change can be derived from CO nor DFT from this interpretation.

Chapter 8

Conclusions

In this thesis we have argued that although AGM functions provide coherent change operations for single theories separately, these change operations are not necessarily jointly coherent. We have regarded this as a serious limitation of the AGM formalism and the work in this thesis has been devoted to overcome this limitation.

According to the AGM theory, the change of one theory may be unrelated to the change of another. For this reason, we have posed that standard AGM functions are better regarded as unary functions relative to an underlying theory, which take a formula and return an updated theory.

In this thesis we have defined authentic binary functions for theory change and we have argued that they solve the problem of change in different theories. Being definitionally simple they also solve, to some extent, the problem of iterated change. Since binary functions are defined for every theory, the result of one application of a change function is a theory that can yet be put as an argument of the same change function. Consequently, the scheme for iterated change induced by binary functions is deterministic with respect to their arguments. This behaviour has been interpreted as a lack of historic memory, which is not always desirable in a model of iterated change.

We have started our study of binary functions with two exceptions in the AGM theory, which satisfy a number of elegant properties, AGM expansions and full meet functions. We have continued with a distinctive binary operation for theory change outside the AGM framework: Katsuno and Mendelzon's update. In contrast to the AGM tradition, Katsuno and Mendelzon formalized their update operation as a binary connective in a finite language. We have shown that nothing crucial relies on this formal difference, as it is possible to reformulate the update operator as a binary function that takes a theory and a formula and returns a theory. However, we have exhibited an unexpected result. Katsuno and Mendezon's postulates are incomplete to characterize the update function for infinite propositional languages. We have provided an appropriate set of postulates, strengthening theirs, and proved the corresponding representation theorem for possibly infinite propositional languages. In this way we have extended Katsuno and Mendelzon's original work just defined for the finite case. Our results complete and clarify those of [Peppas and Williams, 1995], who realized that Katsuno and Mendlezon's framework was incomplete for first order languages. In addition, we have put the AGM revision and update in an even definitional basis that may allow for a better comparison or understanding, when the nature of their difference is still an open question in the philosophical logic literature.

We have given two different formulations extending the AGM framework, iterable AGM functions and analytic AGM functions. We have proposed them as plausible candidates for changing multiple theories, and we have also shown that they satisfy significant properties of iterated change. We have defined both functions for possibly infinite languages and in both cases we have provided postulates extending AGM's and given representation theorems for different formal structures.

We have defined iterable AGM functions with the peculiar property of being almost constant on their first argument, the second argument held fixed. In spite of their quite simple definition they provide a strong notion of coherence with respect to the change in different theories. According to iterable functions, the change in one theory depends on the change of the largest theory, the whole language. We have shown that they satisfy a number of significant properties that have been presented in the literature.

Analytic AGM functions have been defined as almost monotone functions on their first argument (the other held fixed) without being almost constant. As

AGM functions for changing multiple theories we have shown that they possess a significant property. The analytic change operation can be calculated by means of a case analysis, such that if one theory is an extension of another the cases considered for the first can be lift to the cases for the second. In addition, we have defined and characterized maxi-analytic AGM functions that when applied to a consistent complete theory they also return a consistent complete theory.

But analytic AGM revisions have also another interest. We have shown that they provide a formal connection with the update function of Katsuno and Mendelzon. Analytic functions provide a new presentation of AGM revision based on the update semantic apparatus establishing in such a way a bridge between the two seemingly incomparable frameworks.

Finally, we have studied and compared two conditional logics that provide a logical calculus for theory change, Alchourrón's logic DFT and Boutilier's logic CO. By appealing to the notion of consequence, the two logics can be used to calculate changes in different different theories. We have revealed the connection between the two logics showing that in a restricted language, the two logics validate the same conditional sentences. Hence, under appropriate restricting conditions the two logics are equivalent. In addition we have identified the scheme of iterated change induced by the nested occurrences of the conditional connective in the two logics and we have shown that it yields a trivial notion of iterated change.

8.1 Further Work

Iterable functions and analytic functions are just two instances of binary AGM functions, and there is possibly a whole landscape of binary functions that remains to be considered. Iterable and analytic functions can be regarded as two extreme poles. The result of an iterable revision is either an expansion or just the result of revising to the largest theory, the whole language. In contrast, the result of an analytic revision of some theory is always dependent on the revision of each of its maximal consistent extensions. It may be possible to define binary functions that stay in between the two.

In a different perspective, we believe that our analytic AGM functions can be a definitional basis for merge operators [Fuhrmann and Hansson, 1994]. In their most general form they are n-ary functions taking n theories and returning a theory, $\circ : \mathbb{K}^n \rightarrow \mathbb{K}$, and this path of investigation has hardly been addressed.

Given the link existing between conditionals and theory change, as pursued for example by [Grahne, 1991] and [Boutilier, 1996], it seems interesting to investigate conditional logics for our frameworks. In such logics our binary functions would become connectives in the object language and only finitely axiomatizable theories would be considered. The iteration of our functions would be reflected as logical formulae with nested occurrences of the change operators. Presumably this logic would provide further light on new properties of binary functions and establish a closer link between theory change and the field of conditional logics.

In this thesis we have not addressed the problem of change functions of conditional theories. A conclusive result, known as Gärdenfors impossibility theorem [1986] has showed that AGM revisions operating on a conditional language are incompatible with the Ramsey test for interpreting conditionals. There is considerable work in the literature on how to deal with the impossibility theorem, proposing either to weaken the Ramsey test or alter the properties of revisions [Gärdenfors, 1987; Gärdenfors *et al.*, 1991; Rott, 1989; Levi, 1988; Boutilier and Goldszmidt, 1993; Hansson, 1992] . But, whatever be the solution to this dilemma, the notion of change in a conditional theory seems to be best modeled via binary functions. In their most general form they would take a conditional theory and a conditional formula and they would return a conditional theory. In the context of conditional theories the property of historic memory seems to play no role, for what binary functions would provide the appropriate notion of change in multiple theories and iterated change.

Bibliography

[Alchourrón and Makinson, 1981] Carlos Alchourrón and David Makinson. Hierarchies of regulations and their logic. In Risto Hilpinen (ed.), editor, *New Studies in Deontic Logic: Norms, Actions, and the Foundations of Ethics*, pages 125–148, 1981.

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

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

[Alchourrón and Makinson, 1986] Carlos Alchourrón and David Makinson. Maps between some differents kinds of contraction functions: The finite case. *Studia Logica*, 45:187–198, 1986.

[Alchourrón *et al.*, 1985] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

[Alchourrón, 1995] Carlos Alchourrón. Defeasible logics: Demarcation and affinities. In L. Fariñas del Cerro G. Crocco and A. Herzig, editors, *Conditionals: from Philosophy to Computer Science*, Studies in logic and Computation, Series editor D.M. Gabbay, pages 67–102. Oxford Science Publications, 1995. Also appeared as Defeasible, Strict and Counterfactual Conditionals. Nonmonotonic Inference Relations, in the proceedings of the Workshop on *Conditionals and Artificial Intelligence*, 1993. Toulouse, France.

[Åquist, 1973] Lennart Åquist. Modal logic with subjunctive conditionals and dispositional predicates. *Journal of Philosophical Logic*, 2:1–76, 1973.

[Areces and Becher, 1999] Carlos Areces and Verónica Becher. Iterable AGM functions. In Rott H. and Williams M., editors, *Frontiers of Belief Revision.* Kluwer Applied Logic Series, 1999. Forthcoming.

[Areces and Rott, 1999] Carlos Areces and Hans Rott. Revising by sequences. manuscript, April 1999.

[Becher *et al.*, 1998] V. Becher, E. Fermé, S. Lazzer, R. Rodriguez, C. Oller, and G. Palau. Some observations on carlos alchourrón's theory of defeasible conditionals. In Mc Namara P. Prakken H. (eds), editor, *Norms Logics and Information Systems, New Studies on Deontic Logic and Computer Science.* IOS Press, 1998.

[Becher, 1995a] Verónica Becher. Two conditional logics for defeasible inference: A comparison. In J. Wainer and A. Carvalho, editors, *Advances in Artificial Intelligence*, volume 991 of *Lecture Notes in Artificial Intelligence*, pages 49–58. Springer Verlag, 1995.

[Becher, 1995b] Verónica Becher. Unified semantics for revision and update, or the theory of lazy update. In *Anales de las 24 Jornadas Argentinas de Informatica e Investigación Operativa, 7-9 de Agosto 1995*, pages 641–650, Buenos Aires, Argentina, 1995. SADIO.

[Boutilier and Goldszmidt, 1993] Craig Boutilier and Moisés Goldszmidt. Revision by conditional beliefs. In *Proceedings of the 11th National Conference on Artificial Intelligence (AAAI)*, pages 649-654. Morgan Kaufmann, Washington, D.C., july 1993.

[Boutilier, 1992a] Craig Boutilier. *Conditional Logics for Default Reasoning and Belief Revision.* PhD thesis, University of Toronto, Canada, 1992.

[Boutilier, 1992b] Craig Boutilier. Normative, subjunctive, and autoepistemic defaults: Adopting the Ramsey test. In William Nebel, Bernhard; Rich, Charles; Swartout, editor, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*, pages 685–696, Cambridge, MA, October 1992. Morgan Kaufmann.

[Boutilier, 1992c] Craig Boutilier. Sequences of revisions: On the semantics of nested conditionals. Technical Report TR-92-23, Department of Computer Science, University of BritishColumbia, 1992.

[Boutilier, 1993] Craig Boutilier. Revision sequences and nested conditionals. In Ruzena Bajcsy, editor, *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, pages 519–525. Morgan Kaufmann, Chambéry, Savoie, France, august 1993.

[Boutilier, 1996] Craig Boutilier. Iterated revision and minimal revision of conditional beliefs. *Journal of of Philosophical Logic*, 25:263–305, 1996.

[Cantwell, 1997] John Cantwell. Some logics of iterated belief change. *Studia Logica*, 1997. to appear.

[Chellas, 1980] Brian F. Chellas. *Modal Logic: An Introduction.* Cambridge University Press, Cambridge, 1980.

[Darwiche and Pearl, 1997] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1-2):1–29, 1997. Also appeared in Ronald Fagin, ed. TARK'94 - Proceedings of the Fifth Confererence on Theoretical Aspects of Reasoning About Knowledge, Morgan Kaufmann, Pacific Grove, Cal. pp5-23.

[Fermé, 1999] Eduardo Fermé. Irrevocable belief revision and epistemic entrenchment. In *Workshop on Logic, Language and Information 1999, Itatiaia, Brasil*, 1999.

[Freund and Lehmann, 1994] Michael Freund and Daniel Lehmann. Belief revision and rational inference. Technical Report TR 94-16, The Leibniz Center for Research in Computer Science, Institute of Computer Science,Hebrew University, Jerusalem, Israel, july 1994.

[Friedman and Halpern, 1996] Nir Friedman and Joseph Halpern. A qualitative Markov assumption and its implications for belief change. In *Uncertainty in artificial intelligence*, pages 263–273, Portland, OR, 1996. Morgan Kaufmann.

[Fuhrmann and Hansson, 1994] André Fuhrmann and Sven Ove Hansson. A survey of multiple contraction. *Journal of Logic, Language and Information*, 3:39–74, 1994.

[Gärdenfors and Makinson, 1990] Peter Gärdenfors and David Makinson. Belief revision and nonmonotonic logic: Two sides of the same coin? In *European Conference on Artificial Intelligence*, pages 768–773, 1990.

[Gärdenfors *et al.*, 1991] Peter Gärdenfors, Sten Lindström, Michael Morreau, and Wlodzimierz Rabinowicz. The negative Ramsey test: Another triviality result. In Fuhrmann and Morreau, editors, *The Logic of Theory Change*, pages 93–126, Berlin, 1991. Springer-Verlag.

[Gärdenfors, 1982] Peter Gärdenfors. Rules for rational changes of belief. In Tom Pauli, editor, *Philosophical Essays dedicated to Lennart Åquist on his fiftieth birthday*, volume 34, pages 88–101. 1982.

[Gärdenfors, 1984] Peter Gärdenfors. Epistemic importance and minimal changes of belief. *Australasian Journal of Philosophy*, 62:136–157, 1984.

[Gärdenfors, 1986] Peter Gärdenfors. Belief revisions and the Ramsey test for conditionals. *Philosophical Review*, 95:81–93, 1986.

[Gärdenfors, 1987] Peter Gärdenfors. Variations on the Ramsey test: More triviality results. *Studia Logica*, 46(4):321–327, 1987.

[Gärdenfors, 1988] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books/MIT Press, Cambridge (Mass.), 1988.

[Gärdenfors, 1992] Peter Gärdenfors. Belief revision: An introduction. In Peter Gärdenfors, editor, *Belief Revision*, number 29 in Cambridge Tracts in Theoretical Computer Science, pages 1–28. Cambridge University Press, 1992.

[Grahne, 1991] Gösta Grahne. Updates and counterfactuals. In *Proceedings of the International Conference of Principles of Knowledge Representation and Reasoning*, pages 269–276, Cambridge, 1991.

[Grove, 1988] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

[Halpern and Friedman, 1996] Joseph Halpern and Nir Friedman. Belief revision: A critique. In Jon Doyle Luigia Carlucci Aiello and Stuart Shapiro (eds.), editors, *Proceedings of the Fifth International Conference of Principles of Knowledge Representation and Reasoning (KR'96)*, pages 421-431, San Mateo, Cal., 1996. Morgan Kaufmann. Extended version to appear in the Journal of Logic Language and Information.

[Hansson, 1992] Sven Ove Hansson. In defense of the Ramsey test. *The Journal of Philosophy*, 89:522-540, 1992.

[Hansson, 1994] Sven Ove Hansson. Kernel contraction. *Journal of Symbolic Logic*, 59:845-859, 1994.

[Hansson, 1998] Sven Ove Hansson. *A Textbook of Belief Dynamics*. Kluwer Academic Publisher, Dordrecht, 1998. In Press.

[Hayes, 1985] P. Hayes. Some problems and non-problems in representation theory. In R. Brachman and R. Levesque, editors, *Readings in Knowledge Representation*, pages 4-22. Morgan Kaufmann, Los Altos, 1985.

[Hughes and Cresswell, 1968] G. E. Hughes and M. J. Cresswell. *An Introduction to Modal Logic*. Methuen, London, 1968.

[Hughes and Cresswell, 1984] G. E. Hughes and M. J. Cresswell. *A Companion to Modal Logic*. Methuen, London, 1984.

[Humberstone, 1983] I. L. Humberstone. Inaccessible worlds. *Notre Dame Journal of Formal Logic*, 24(3):346-352, 1983.

[Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto Mendelzon. Propositional knowledge base revision and minimal change. *Journal of Artificial Intelligence*, 52:263-294, 1991.

[Katsuno and Mendelzon, 1992] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In Peter Gärdenfors, editor, *Belief Revision*, number 29 in Cambridge Tracts in Theoretical Computer Science, pages 183-203. Cambridge University Press, 1992. Also appeared in Proceedings of First Conference on Knowledge Representation and Reasoning KR91,Cambridge, pp. 387-394, 1991.

[Keller and Winslett, 1985] A. Keller and M. Winslett. On the use of an extended relational model to handle changing incomplete information. *IEEE Transactions on Software Engineering*, SE-11(7):620–633, july 1985.

[Lehmann, 1995] Daniel Lehmann. Belief revision, revised. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI95)*, pages 1534-1540, Montreal, Canada, August 1995. Morgan Kaufmann Publishers.

[Levi, 1988] Isaac Levi. Iteration of conditionals and the ramsey test. *Synthese*, 76(1):49–81, 1988.

[Lewis, 1973] David Lewis. *Counterfactuals.* Blackwell, Oxford, 1973.

[Lindström and Rabinowicz, 1992] Sten Lindström and Wlodzimierz Rabinowicz. Belief revision, epistemic conditionals and the Ramsey test. *Synthese*, 91:195–237, 1992.

[Makinson and Gärdenfors, 1988] David Makinson and Peter Gärdenfors. Revisions of knowledge systems using epistemic entrenchment. In Moshe Y. Vardi, editor, *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge (TARK88)*, pages 83–95, Los Altos, 1988. Morgan Kaufmann.

[Makinson and Gärdenfors, 1991] David Makinson and Peter Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In Andre Fuhrmann and Michael Morreau, editors, *The Logic of Theory Change*, pages 185–205. Springer-Verlag, Berlin, 1991.

[Makinson, 1993] David Makinson. Five faces of minimality. *Studia Logica*, 52(3):339–379, 1993.

[Moore, 1982] R. Moore. The role of logic in knowledge representation and commonsese reasoning. In *Proceedings of the Second International Conference of Principles of Knowledge Representation and Reasoning*, pages 428–433, Pittsburgh, 1982.

[Nayak, 1994] Abhaya Nayak. Iterated belief change based in epistemic entrenchment. *Erkentniss*, 41(3):353–390, 1994.

[Peppas and Williams, 1995] Pavlos Peppas and Mery-Anne Williams. Constructive modellings for theory change. *Notre Dame Journal of Formal Logic*, 36(1):120–133, 1995.

[Peppas, 1993] Pavlos Peppas. *Belief Change and Reasoning about Action.* PhD thesis, Basser Department of Computer Science, University of Sydney, Australia, December 1993.

[Putnam, 1970] H. Putnam. Is semantics possible? In H.E. Kiefer snd M.K. Munitz (eds), editor, *Language Belief and Metaphisics*, pages 50–63. SUNY Press, Albany, 1970.

[Rott, 1989] Hans Rott. Conditionals and theory change: Revisions, expansions, and additions. *Synthese*, 81(1):91–113, 1989.

[Rott, 1991] Hans Rott. Two methods of constructing contractions and revisions of knowledge systems. *Journal of Philosophical Logic*, 20:149–173, 1991.

[Rott, 1992a] Hans Rott. On the logic of theory change: More maps between different kinds of contraction functions. In Peter Gärdenfors, editor, *Belief Revision*, number 29 in Cambridge Tracts in Theoretical Computer Science, pages 122–141. Cambridge University Press, 1992.

[Rott, 1992b] Hans Rott. Preferential belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information*, 1:45–78, 1992.

[Rott, 1993] Hans Rott. Belief contraction in the context of the general theory of rational choice. *Journal of Symbolic Logic*, 58(4):1426–1450, 1993.

[Rott, 1998] Hans Rott. Coherence and conservatism in the dynamics of belief. Part II: Iterated belief change without dispositional coherence. Technical report, Department of Philosophy and ILLC, University of Amsterdam, may 1998. submitted to special issue of Studia Logica, edited by Karl Schlechta.

[Rott, 1999] Hans Rott. Coherence and conservatism in the dynamics of belief. Part I: Finding the right framework. *Erkenntnis*, 1999. to appear.

[Schlechta *et al.*, 1996] Karl Schlechta, Daniel Lehmann, and Menachem Magidor. Distance semantics for belief revision. In Yoav Shoham, editor, *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and*

Knowledge, pages 137–146, San Francisco, 1996. Morgan Kaufmann Publishers.

[Segerberg, 1995] Krister Segerberg. Some questions about hypertheories. In Sven Ove Hansson and Wlodzimierz Rabinowicz, editors, *Logic for a change*, number 9 in Uppsala Prints and Preprints in Philosophy, pages 136–154. Dep. of Philosophy, Uppsala University, 1995.

[Segerberg, 1997] Krister Segerberg. Irrevocable belief revision in dynamic doxastic logic. In *Uppsala Prints and Preprints in Philosophy*, pages –. Dep. of Philosophy, Uppsala University, 1997.

[Spohn, 1987] Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change and Statistics*, volume 2, pages 105–134. D. Reidel, Dordrecht, 1987.

[Stalnaker, 1968] Robert Stalnaker. A theory of conditionals. In N. Rescher, editor, *Studies in logical theory*, number 2 in American philosophical quarterly monograph series. Blackwell, Oxford, 1968. Also appears in*Ifs*, (ed., by W. Harper,R. C. Stalnaker and G. Pearce), Reidel, Dordrecht, 1981.

[Williams, 1994] Mary-Anne Williams. Transmutations of knowledge systems. In Jon Doyle, Erik Sandewall, and Pietro Torasso, editors, *Proceedings of the fourth International Conference on Principles of Knowledge Representation and Reasoning*, pages 619–629. Morgan Kaufmann, Bonn, Germany, May 1994.