CARLOS ARECES & VERÓNICA BECHER

ITERABLE AGM FUNCTIONS

CARLOS ARECES & VERÓNICA BECHER

ITERABLE AGM FUNCTIONS

ABSTRACT: The AGM model [Alchourrón *et al.*, 1985] has been criticized for not addressing the problem of iterated change. This is true, but in some cases the invalid claim that "AGM does not allow iteration" has been made. In this paper we examine the most elementary scheme of iteration: a binary operation * defined for any theory K and any formula α. We naturally extend the AGM model showing that it is in fact compatible with iteration. Following Alchourrón and Makinson's early reference to an iterated form of the safe contraction function, we provide extended constructions for each of the five AGM presentations (meet functions, systems of spheres, postulates, epistemic entrenchments and safe hierarchies) and prove their equivalence.

1 INTRODUCTION

The issue of iterated theory change is indeed interesting. Legal codes are under constant modification, new discoveries shape scientific theories, and robots ought to update their representation of the world each time a sensor gains new data. A pertinent criticism to the AGM formalism of theory change [Alchourrón *et al.*, 1985] is its lack of definition with respect to iterated change. Let's start by introducing the basic elements in the AGM framework

Informally, let K be a collection of data (which we call a theory), and α a piece of information, then the AGM model characterizes three different kinds of theory change: expansion ($K+\alpha$) when information is *simply added* to the theory, revision ($K*\alpha$) when information is *consistently added* to the theory, and contraction ($K-\alpha$) when information is *eliminated* from the theory. The iteration of each of these operations 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.

The aim of the present work is to define the simplest scheme of iteration for AGM: *iterable functions*, binary functions whose behaviour is specified for any theory K and formula α. The basic idea of this article dates back to Alchourrón and Makinson's work on safe contractions [Alchourrón and Makinson, 1985]. In that article they briefly discuss properties of the contraction function with respect to the intersection and union of theories, and they present some properties of "multiple contractions." Following this idea we elaborate upon the definition of an iterable function for each of the five AGM presentations: postulates [Alchourrón *et al.*, 1985], systems of spheres [Grove, 1988], meet functions [Alchourrón and Makinson, 1982], epistemic entrenchments [Gärdenfors and Makinson, 1988] and safe hierarchies [Alchourrón and Makinson, 1985].

In this work we assume knowledge of the AGM model as well as basic knowledge of classical logic. About notation, we fix a propositional language $\mathcal{L}$ and a consequence operation Cn that includes the classical one, is compact and satisfies disjunction in the antecedent. We use Greek letters α, β, γ for sentences of $\mathcal{L}$, capital Roman letters K, K', H for theories of $\mathcal{L}$ (subsets of $\mathcal{L}$ closed under Cn), and denote with $\mathbb{K}$ the class of all theories of $\mathcal{L}$. For the purposes of this work we consider the terms possible world, maximal consistent subset of $\mathcal{L}$ and valuation on $\mathcal{L}$ interchangeable. The letters U, V, W range over sets of possible worlds. M is the set of all possible worlds. $[\alpha]$ denotes the set of all possible worlds containing α. If K is a theory, $[K]$ denotes the set of possible worlds containing K. Given U a set of possible worlds, $\mathrm{Th}(U)$ returns the associated theory.

2 AGM AND THE NOTION OF ITERATION

The AGM model has as points of departure a theory to be modified, a formula to be considered the new information and change functions. The framework seems then to be naturally that of binary functions $\bullet$ which when applied to a theory K and a formula α return a new theory K'. In infix notation, $K \bullet \alpha = K'$. However, when we study the AGM model, we immediately realize that change functions are *relative to the original theory*. This is an important issue: although change functions are presented as binary operators they are characterized only for a given fixed theory K. I.e. the first argument (the theory to be modified) has been fixed obtaining in such a way a family of functions $\bullet^K$. The binary approach has been abandoned and AGM only aims to study one of the indexed unary functions without trying to correlate it with the rest of the family.

Turning to iterated change now, once that 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 change functions, one for each possible theory, we can freely perform successive changes. Consider any pair of formulas α, β and a particular theory K. In order to calculate the successive changes of K, first by α and then by β, we need the change function $\bullet_1$ for K, and also the change function $\bullet_2$ relative to $(K \bullet_1 \alpha)$. The result of the successive change is the theory $(K \bullet_1 \alpha) \bullet_2 \beta$. But beware, if there are no 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.

These facts have already been recognized and discussed in the theory change community and many proposals have in common the following characteristic: the AGM model is expanded 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 build 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; Hansson, To come; Nayak, 1994; Lehmann, 1995], or by enriching the model with numbers [Darwiche and Pearl, 1997; Spohn, 1987; Williams, 1995]. Notice that in these approaches we are not really going back to a binary function $\bullet$ and returning the theory K to its original role of argument. The "construction" method is more flexible than considering a binary function. Given a theory K and a formula α, change functions associated to $K \bullet \alpha$ are not uniquely determined and depend really on a third argument: the change function for K. These approaches are very rich —they can circumvent the problem of functional behaviour of the change function with respect to the theory under change, discussed for example in [Lehmann, 1995]. But they are usually complex.

We will consider in this paper the case of binary change functions. More specifically, AGM binary (non indexical) functions which are total in their two arguments. Any such function can trivially be iterated, $(K \bullet \alpha) \bullet \beta$ is well defined. We will call these functions *iterable AGM functions*. Notice that iterable functions can be considered a particular case of the construction method described above: after applying it to a theory, the same function is "returned" and can be applied again to the resulting theory. Some advantages of this approach are evident: it is mathematically elegant, simple and remains close to the AGM model (each time the theory argument is fixed a standard AGM indexical function is obtained). But, while formally attractive, an iterable function makes a strong simplifying assumption since it lacks historic memory; each theory is modified in a predetermined way independently of how we have obtained such a theory. An iterable function $\bullet$ is deterministic with respect to the theory to be modified. I.e. it satisfies:

$$\text{If } K = (\ldots(H \bullet \alpha_1)\ldots \bullet \alpha_n) \text{ then } K \bullet \alpha = (\ldots(H \bullet \alpha_1)\ldots \bullet \alpha_n) \bullet \alpha,$$

which is posed as undesirable in [Lehmann, 1995] (where it is referred to as a non-postulate.) But if K is really considered an *argument* of the function $\bullet$, this is to be expected. If f is a function, it is required that $f(a) = f(b)$ whenever $a = b$.

In spite of the modesty of functional change operators they vary in the sophistication of their fixed structures and associated behavior. A quite elaborate one, outside the AGM framework, is Katsuno and Mendelzon's account of *update* [Katsuno and Mendelzon, 1991], which is 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. With the same spirit of update, Becher's lazy update [Becher, 1995; Becher, 1998] formalizes an iterable version of the AGM model defined in terms of a fixed set of wellfounded orders over possible worlds.

In the simplest case we have binary change functions that depend on no order at all, as the original AGM expansion and the full meet contraction function [Alchourrón and Makinson, 1982]. The AGM representation theorem for the function $+$ states that $K+\alpha = \mathrm{Cn}(K \cup \{\alpha\})$. Expansion inherits its capacity of iteration from the consequence relation, which is applicable to any set of formulas. For any theory K and formulas α, β, $(K+\alpha)+\beta = \mathrm{Cn}(\mathrm{Cn}(K \cup \{\alpha\}) \cup \{\beta\}) = \mathrm{Cn}(K \cup \{\alpha\} \cup \{\beta\})$. As indicated by its signature, expansion is a total binary function $+ : I\!\!K \times \mathcal{L} \rightarrow I\!\!K$. While, as it stands in the original partial meet construction presented in [Alchourrón *et al.*, 1985], contraction is a unary function relative to a theory K, $-^{K} : \mathcal{L} \rightarrow I\!\!K$, based on a selection function s^{K} depending on K.

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

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

The limiting case in which the function s^{K} returns the whole set $K \perp \alpha$ originates the full meet contraction function. The selection function s^{K} relative to K disappears, yielding a contraction that solely depends on the explicit argument K and α, i.e. if $-$ is a full meet, again we have a binary total function $- : I\!\!K \times \mathcal{L} \rightarrow I\!\!K$. Moreover, for full meet contraction functions $K - \alpha = K \cap \mathrm{Cn}(\neg\alpha)$. As expansion functions, they depend on no underlying structure, relative order or selection function, and are applicable to any theory.

But there exists important differences between expansion and full meet contraction. While expansion is appropriately modeled via the consequence relation, full meet contraction is ultimately unsatisfactory. The following two properties reveal its over-simplicity.

Elimination: $(K - \alpha) - \beta = K - (\alpha \wedge \beta)$, from which
Commutativity: $(K - \alpha) - \beta = (K - \beta) - \alpha$ follows.

Alchourrón, Gärdenfors and Makinson argue that full meet contraction functions suffer from too much loss of information and take them as a demarcation of the limiting case. The aim of the next section is to provide iterable AGM functions which are more interesting than full meet functions.

3 ITERABLE AGM FUNCTIONS

In [Alchourrón and Makinson, 1985] Alchourrón and Makinson define a new kind of contraction function called safe contraction, based on a hierarchical order of the elements of a theory K according to what they call their "degree of safety." They 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 relation Cn partitions the elements of K into a finite number of equivalence classes). The general case for infinite theories was solved by Rott in [Rott, 1992a].

Interestingly, Alchourrón and Makinson study in their paper some properties of safe contraction with respect to the intersection or 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 A (the set of propositions) is allowed to vary. [...] But in the case of safe contraction the way of dealing with variations of A is quite straightforward. As we are working with a relation $<$ over A the natural relation to consider over a subset A' of A is simply the restriction $< \cap (A' \times A')$ of $<$ to A'."

They obtain a general result relating $A' - \alpha$ to $A - \alpha$, when $A' \subseteq A$. As a special case they apply it to $(A - \beta) - \gamma$, since $A - \beta \subseteq A$ 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 $\mathcal{L}$ 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 an iterable contraction function based on a unique fixed order of all the formulas, the safe hierarchy.

Reusing the same fixed order makes sense, for example, 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[1]. We will come back to this idea in Section 3.5.

The following subsections 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 inter-definable 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 for further details).

3.1 Extended Safe Contraction Functions

A relation $<_{sf}$ over a set A is a *hierarchy* if it is acyclic: for any set of elements $\alpha_1, \ldots, \alpha_n \in A, n \geq 1$, it is not the case that $\alpha_1 <_{sf} \alpha_2 <_{sf} \ldots \alpha_n <_{sf} \alpha_1$. A relation $<_{sf}$ over A *continues up* Cn if for every $\alpha_1, \alpha_2, \alpha_3 \in A$, if $\alpha_1 <_{sf} \alpha_2$ and $\alpha_3 \in \mathrm{Cn}(\alpha_2)$ then $\alpha_1 <_{sf} \alpha_3$. A relation $<_{sf}$ over A is *virtually connected* if for every $\alpha_1, \alpha_2, \alpha_3 \in A$ if $\alpha_1 <_{sf} \alpha_2$ then either $\alpha_1 <_{sf} \alpha_3$ or $\alpha_3 <_{sf} \alpha_2$. Let $<_{sf}$ be a virtually connected hierarchy over a theory K that continues

[1] This interpretation was pointed out to us by Isaac Levi.

up Cn, and let α be a sentence. The safe contraction function $-$ is defined as

$$K - \alpha = K \cap \mathrm{Cn}(\{\beta \mid \forall K' \subseteq K, \text{ s.t. } \alpha \in \mathrm{Cn}(K') \text{ and } K' \text{ is } \subseteq\text{-minimal with this property, } \beta \notin K' \text{ or there is } \gamma \in K' \text{ s.t. } \gamma <_{sf} \beta\}).$$

The elements of $K -_{sf} \alpha$ are called the safe elements of K with respect to α since they can not be "blamed" for implying α. An element is safe for α if it does not belong to any of the $\subseteq$-minimal subsets of K that imply α, 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 $-_{sf}$ based on a hierarchy over all the sentences of $\mathcal{L}$.

DEFINITION 1. **(Derived Order)** Let $<_{sf}$ be a hierarchy over the language $\mathcal{L}$. 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).

PROPOSITION 2. *Let $<_{sf}$ be a virtually connected hierarchy that continues up* Cn *in $\mathcal{L}$, 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. ■

Once this result is obtained, to define an iterable safe contraction is straightforward.

DEFINITION 3. **(Iterable Safe Contraction)** Let $<_{sf}$ be a virtually connected hierarchy that continues up Cn in $\mathcal{L}$. The iterable AGM contraction $-_{sf} : I\!\!K \times \mathcal{L} \to I\!\!K$ is defined as

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

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 1 and 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 $\mathcal{L}$ and defined its restriction $<^K_{sf}$. A valid question is whether the converse can also be achieved. Given a hierarchy for K can a hierarchy for $\mathcal{L}$ be defined such that the iterable function agrees with $-^K$ when applied to K?

PROPOSITION 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 α, $K -_{sf} \alpha = -^K(\alpha)$.*

Proof. Given $<^K_{sf}$ the order associated to $-^K$, define $<_{sf}$ as follows: $\alpha <_{sf} \beta$ iff either $(\alpha \notin K)$ or $(\alpha, \beta \in K$ and $\alpha <_{sf} \beta)$. 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 $<^K_{sf} = <_{sf}|K$, and it is not hard to check that $<_{sf}$ is a virtually connected hierarchy that continues up over Cn in $\mathcal{L}$. ■

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

3.2 Extended Partial Meet Contraction Functions

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

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

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

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

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 $K-8$, s^K must be *transitively relational*, i.e. for each $\alpha \in \mathcal{L}$ the selection function returns the smallest elements according to some transitive relation defined over $K \perp \alpha$.

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

DEFINITION 5. **(Derived Selection Functions)** Let $s^{\mathcal{L}}$ be a selection function for $\mathcal{L}$. Then, for any theory K the selection function s^K is

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

It is immediate to see that each derived s^K is indeed a selection function. What is more interesting is to check whether each s^K is transitively relational whenever $s^{\mathcal{L}}$ is.

PROPOSITION 6. *If $s^{\mathcal{L}}$ 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 $s^{\mathcal{L}}$ is transitively relational there is a transitive relation R defined over $\mathcal{L} \perp \alpha$ whose smallest elements are selected by $s^{\mathcal{L}}(\alpha)$. This relation R can be projected over each $K \perp \alpha$ to show that $s^K(\alpha)$ selects the smallest elements of a transitive relation. ■

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

DEFINITION 7. **(Iterable Partial Meet Contraction)** Let $s^{\mathcal{L}}$ be a transitively relational selection function over $\mathcal{L}$. The iterable AGM contraction $-_{pm} : I\!\!K \times \mathcal{L} \to I\!\!K$ is defined as $K -_{pm} \alpha = \bigcap s^K(\alpha)$, 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.

PROPOSITION 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* α*,* $K -_{pm} \alpha = -^{K}(\alpha)$*.*

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

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

3.3 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[2]. We first turn to Grove's original framework [Grove, 1988].

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

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

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

Using the function c_S, the function $f_S : \mathcal{L} \to \mathcal{P}(M)$ is defined as $f_S(\alpha) = [\alpha] \cap c_S(\alpha)$. Given a sentence α, $f_S(\alpha)$ returns the closest elements (with respect to theory K) where α holds. Grove shows that the function defined as $-^{K}(\alpha) = \mathrm{Th}([K] \cup f_S(\neg\alpha))$ is an AGM contraction function. And conversely,

[2] Personal communication.

for any AGM contraction function relative to a theory K there is a system of spheres S centered in $[K]$ 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 $\mathcal{L}$ has the particular property that its innermost sphere is the empty set, since $[\mathcal{L}] = \emptyset$. Given a system of spheres S centered in $\emptyset$ we define for any theory K a derived system S^K centered in $[K]$ simply by "filling in" the innermost sphere of S with $[K]$.

DEFINITION 9. **(Derived System of Spheres)** Let S be a system of spheres for $\mathcal{L}$. Then for any theory K the derived system of spheres S^K is defined as $S^K = \{[K] \cup S_i \mid S_i \in S\}$.

PROPOSITION 10. *Let S be a system of spheres for $\mathcal{L}$. Then for any theory K, S^K is a system of spheres centered in 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 $-_{ss} : \mathbb{K} \times \mathcal{L} \to \mathbb{K}$, applicable to every theory K and every formula α.

DEFINITION 11. **(Iterable Sphere Contraction)** Let S be a system of spheres for $\mathcal{L}$. The iterable AGM contraction $-_{ss} : \mathbb{K} \times \mathcal{L} \to \mathbb{K}$ is defined as $K -_{ss} \alpha = \mathrm{Th}([K] \cup f_S^K(\neg\alpha))$, 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.

PROPOSITION 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 α, $K -_{ss} \alpha = -^K(\alpha)$.*

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 $\mathcal{L}$. Define S centered in $\emptyset$ as $S = S^K \cup \{\emptyset\}$. Clearly, S validates $S1$ to $S4$ for $\mathcal{L}$. ■

3.4 Extended Epistemic Entrenchments

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

EE1. If $\alpha \leq_{ee} \beta$ and $\beta \leq_{ee} \delta$ then $\alpha \leq_{ee} \delta$.
EE2. If $\beta \in \mathrm{Cn}(\alpha)$ then $\alpha \leq_{ee} \beta$.
EE3. $\alpha \leq_{ee} (\alpha \wedge \beta)$ or $\beta \leq_{ee} (\alpha \wedge \beta)$.
EE4. If theory K is consistent then $\alpha \notin K$ iff $\alpha \leq_{ee} \beta$ for every β.

EE5. If $\beta \leq_{ee} \alpha$ for every β then $\alpha \in \mathrm{Cn}(\emptyset)$.

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

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

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 $\mathcal{L}$. If we consider as a point of departure an epistemic entrenchment over the contradictory theory $\mathcal{L}$, 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 13. **(Derived Epistemic Entrenchment)** Let $\leq_{ee}$ be an epistemic entrenchment relation for $\mathcal{L}$. Then for any theory K the derived epistemic entrenchment relation $\leq_{ee}^{K}$ is defined as:

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

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

PROPOSITION 14. *Let $\leq_{ee}$ be an epistemic entrenchment relation for $\mathcal{L}$, then for any theory K, $\leq_{ee}^{K}$ is an epistemic entrenchment relation for K.*

DEFINITION 15. **(Iterable Epistemic Entrenchment Contraction)** Let $\leq_{ee}$ be an epistemic entrenchment relation for $\mathcal{L}$. The iterable AGM contraction $-_{ee} : \mathbb{K} \times \mathcal{L} \to \mathbb{K}$ is defined as $K -_{ee} \alpha = \{\beta \in K \mid \alpha \in \mathrm{Cn}(\emptyset) \text{ or } \alpha <_{ee}^{K} (\alpha \vee \beta)\}$, 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.

PROPOSITION 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 $-_{ee}$ based on epistemic entrenchments such that for every α, $K -_{ee} \alpha = -^{K}(\alpha)$.*

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

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

3.5 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 $\mathcal{L}$, 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:

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

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

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

In Section 4 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 $K{*}1$ to $K{*}9$ is in fact iterable: for any $\alpha, \beta \in \mathcal{L}$, $K * \alpha * \beta$ is well defined: If $\neg\beta \in K * \alpha$ then $K * \alpha * \beta = (\mathcal{L} * \beta)$; else $K * \alpha * \beta = (K * \alpha) + \beta$. 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 Section 2, or as reported in [Friedman and Halpern, 1996], the qualitative analogue of the Markov Assumption.

The revision postulate $K{*}9$ is sound with the interpretation of revision as a kind of tentative reasoning. The revision function for $\mathcal{L}$ 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 in [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 [Gärdenfors and Makinson, 1991] 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.

3.6 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 17. **(Postulates/Systems of Spheres)** *Given an iterable AGM contraction $-$ satisfying $K{-}1$ to $K{-}9$, there exists a system of spheres S for $\mathcal{L}$ such that for every K and every α, $K-\alpha = \mathrm{Th}([K] \cup f_S^K(\neg\alpha))$. Conversely, every $-_{ss}$ based on a system of spheres S for $\mathcal{L}$ 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 $\mathcal{L}$ that satisfies $K{-}1$ to $K{-}8$ there is a system of spheres S for $\mathcal{L}$ such that $\mathcal{L}-\alpha = \mathrm{Th}(f_s(\neg\alpha))$. By definition $S^K = \{[K] \cup S_i \mid S_i \in S\}$. There are two cases. For any $\alpha \notin K$, clearly $f_S^K(\neg\alpha) = [K] \cap [\neg\alpha]$, then $\mathrm{Th}([K] \cup f_S^K(\neg\alpha)) = K$ and by postulate $K{-}3$, $K = K-\alpha$, so we are done. For $\alpha \in K$, $f_S^K(\neg\alpha) = f_S^{\mathcal{L}}(\neg\alpha)$, then $\mathrm{Th}([K] \cup f_S^K(\neg\alpha)) = \mathrm{Th}([K] \cup f_S^{\mathcal{L}}(\neg\alpha)) = K \cap \mathrm{Th}(f_S^K(\neg\alpha)) = K \cap (\mathcal{L}-\alpha)$, and we are done. ■

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

THEOREM 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 $\mathcal{L}$ such that for every K and every*

α, $K - \alpha = \{\beta \in K \mid \alpha \in \mathrm{Cn}(\emptyset)$ or $\alpha <^{K}_{ee} (\alpha \vee \beta)\}$. 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 $\mathcal{L}$.

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

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

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

THEOREM 19. **(Meet Functions/Systems of Spheres)** *For each iterable partial meet contraction $-_{pm}$ there exists a system of spheres S for $\mathcal{L}$ such that for every theory K and every α, $K -_{pm} \alpha = \mathrm{Th}([K] \cup c^{K}_{S}(\neg\alpha))$. Conversely, for each iterable contraction $-_{ss}$ defined by a system of spheres there exists a selection functions $s^{\mathcal{L}}$ such that for every theory K and every α, $K -_{ss} \alpha = \bigcap s^{K}(\alpha)$.*

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

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 to iteration presented in the article are indeed five faces of the same phenomenon.

THEOREM 20. **(Epistemic Entrenchments/Safe Hierarchies)** *For each iterable epistemic entrenchment contraction $-_{ee}$ there exists a virtually connected hierarchy $<_{sf}$ that continues up* Cn *in $\mathcal{L}$, such that for every theory K and every α, $K -_{ee} \alpha = K \cap \mathrm{Cn}(\{\beta \mid \forall K' \subseteq K$, s.t. $\alpha \in \mathrm{Cn}(K')$ and K' is a $\subseteq$-minimal with this property, $\beta \notin K'$ or there is $\gamma \in$*

K' s.t. $\gamma <^K_{sf} \beta\}$). Conversely, for each safe iterable contraction $-_{sf}$ there exists an epistemic relation $\leq_{ee}$ for $\mathcal{L}$, such that for every theory K and every α $K -_{sf} \alpha = \{\beta \in K \mid \alpha \in \mathrm{Cn}(\emptyset)$ or $\alpha <^K_{ee} (\alpha \vee \beta)\}$.

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 $\mathcal{L}$ 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 $\mathcal{L}$. Take $\leq_{ee}$ as the basis for our epistemic entrenchment iterable contraction function $-_{ee}$. If $\alpha \in \mathrm{Cn}(\emptyset)$ or $\alpha \notin K$, then as both $-_{sf}$ and $-_{ee}$ are AGM functions, $K -_{sf} \alpha = K = \{\beta \in K \mid \alpha \in \mathrm{Cn}(\emptyset)$ or $\alpha <^K_{ee} (\alpha \vee \beta)\}$. If $\alpha \notin \mathrm{Cn}(\emptyset)$ and $\alpha \in K$, as the functions satisfy $K-9$, $K -_{sf} \alpha = (\mathcal{L} -_{sf} \alpha) \cap K = (\mathcal{L} -_{ee} \alpha) \cap K = \{\beta \in K \mid \alpha \in \mathrm{Cn}(\emptyset)$ or $\alpha <^K_{ee} (\alpha \vee \beta)\}$. ■

4 PROPERTIES OF ITERABLE AGM FUNCTIONS

As we already argued, AGM full meet operations are the most elementary iterable constructions that comply with the AGM postulates, but they are commutative. However, despite their modest definition our iterable AGM functions do not validate commutativity: in general, $(K-\alpha)-\beta$ is different from $(K-\beta)-\alpha$.

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 let us define iterable revision functions:

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

This is important since it allows for sequences of different kinds of changes, like for example $(\ldots((K+\alpha)-\beta)*\delta\ldots*\gamma)$. In [Hansson, To come] Hansson 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. $K * \alpha = (K + \alpha) - \neg\alpha.$

Originally, this property was studied in [Alchourrón and Makinson, 1982] as an intuitive property for change functions. Functions satisfying

$$(K - \neg\alpha) + \alpha = (K + \alpha) - \neg\alpha$$

were called *permutable* and the question under which condition an AGM function is permutable was left open in that paper. In [Alchourrón and Makinson, 1985] they show that safe contractions are permutable under the appropriate conditions (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**.

PROPOSITION 21. *Iterable AGM contraction functions are permutable.*

A direct proof of the above is immediate but the result 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:

PROPOSITION 22.
i.) If $\alpha \in K_1 \cap K_2$ *then* $(K_1 \cap K_2) - \alpha = (K_1 - \alpha) \cap (K_2 - \alpha)$ *and* $(K_1 \cup K_2) - \alpha = (K_1 - \alpha) \cup (K_2 - \alpha)$.
ii.) If $\alpha \in K - \beta$ *and* $\beta \in K - \alpha$ *then* $(K - \alpha) - \beta = (K - \alpha) \cap (K - \beta) = (K - \beta) - \alpha$.

And similarly for iterable revisions. These properties hold not just for two theories but also for indefinitely many.

Darwiche and Pearl present in [Darwiche and Pearl, 1997] a number of properties for iterated change. In our notation, they are:

C1. If $\alpha \in \text{Cn}(\beta)$ then $(K * \alpha) * \beta = K * \beta$.
C2. If $\neg\alpha \in \text{Cn}(\beta)$ then $(K * \alpha) * \beta = K * \beta$.
C3. If $\alpha \in K * \beta$ then $\alpha \in (K * \alpha) * \beta$.
C4. If $\neg\alpha \notin K * \beta$ then $\neg\alpha \notin (K * \alpha) * \beta$.
C5. If $\neg\beta \in K * \alpha$ and $\alpha \notin K * \beta$ then $\alpha \notin (K * \alpha) * \beta$.
C6. If $\neg\beta \in K * \alpha$ and $\neg\alpha \in K * \beta$ then $\neg\alpha \in (K * \alpha) * \beta$.

Properties $C1$ to $C4$ are considered in that paper as plausible for iterated change, while $C5$ and $C6$ (satisfied by iterated functions like those in [Boutilier, 1996]) are considered too demanding. In [Lehmann, 1995] property $C2$ was proved inconsistent with AGM postulates $K{*}7$ and $K{*}8$. It is not difficult to prove the following.

PROPOSITION 23.
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 six of the seven postulates of Lehmann's widening ranked models [Lehmann, 1995]. He argues that these structures are suitable for iterative change and proposes seven postulates that fully characterize revision functions based on these structures. His postulates in our notation are:

I1. $K * \alpha$ is a consistent theory.
I2. $\alpha \in K * \alpha$.
I3. If $\beta \in K * \alpha$, then $\alpha \supset \beta \in K$.
I4. If $\alpha \in K$ then $K * \beta_1 * \ldots * \beta_n \equiv K * \alpha * \beta_1 * \ldots * \beta_n$ for $n \geq 1$.
I5. If $\alpha \in Cn(\beta)$, then $K * \alpha * \beta * \beta_1 * \ldots * \beta_n \equiv K * \beta * \beta_1 * \ldots * \beta_n$.

I6. If $\neg\beta \notin K*\alpha$ then $K*\alpha*\beta*\beta_1*\ldots*\beta_n \equiv K*\alpha*(\alpha\wedge\beta)*\beta_1*\ldots*\beta_n$.
I7. $K*\neg\beta*\beta \subseteq Cn(K\cup\beta)$.

The last property, when β is in K, forces dependency between two proper revision steps (in the sense that it constrains a given revision operation on the base of a previous one). Consequently, $I7$ introduces historic memory (at least to some extent), which iterable AGM revisions lack.

PROPOSITION 24.
i) All iterable AGM functions, satisfy $I1$*-*$I6$*.*
ii) There exist iterable AGM functions violating $I7$*.*

Finally, in [Lehmann *et al.*, 1996] the following three properties are suggested as natural properties for iterated change, which can be shown to hold when change functions are defined on "pseudo-distances." For any pair of theories K_1, K_2 and sentences α, β, γ, δ

D1. $(K_1\cap K_2)*\alpha \in \{(K_1*\alpha)\cap(K_2*\alpha), K_1*\alpha, K_2*\alpha\}$.
D2. If $\delta\in(K*\alpha)*\gamma$ and $\delta\in(K*\beta)*\gamma$ then $\delta\in(K*(\alpha\vee\beta))*\gamma$.
D3. If $\delta\in(K*(\alpha\vee\beta))*\gamma$ then $\delta\in(K*\alpha)*\gamma$ or $\delta\in(K*\beta)*\gamma$.

Although property $D1$ is a variation of one of the properties appearing in Proposition 22, in general it is not validated by our iterable AGM functions. We can prove the following proposition:

PROPOSITION 25.
i) All iterable AGM functions satisfy $D2$ *and* $D3$*.*
ii) There exist iterable AGM functions violating $D1$*.*

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.

Acknowledgments: This work has benefited from invaluable discussions with David Makinson and Hans Rott. We are indebted to one of the anonymous referees for his very careful review of the paper. We also acknowledge Eduardo Fermé for his helpful comments on previous versions of this work.

REFERENCES

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

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

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

[Becher, 1995] V. Becher. Unified semantics for revision and update, or the theory of lazy update. In *Annals of the 24 Jornadas Argentinas de Informática e Investigación Operativa (JAIIO)*, 1995.

[Becher, 1998] V. Becher. Iterable AGM revision based on fixed orderings. Technical report, Departamento de Computación. FCEyN. UBA, January 1998.

[Boutilier, 1996] C. Boutilier. Iterated revision and minimal revision of conditional beliefs. *Journal of Philosophical Logic*, 25:263–305, 1996.

[Darwiche and Pearl, 1997] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1-2):1–29, 1997.

[Freund and Lehmann, 1994] M. Freund and D. Lehmann. Belief revision and rational inference. Technical Report 94-16, The Leibniz Center. Hebrew University, Jerusalem, Israel, 1994.

[Friedman and Halpern, 1996] N. Friedman and J. Halpern. A qualitative Markov assumption and its implications for belief change. In *Uncertainty in artificial intelligence (Portland, OR, 1996)*, pages 263–273. Morgan Kaufmann, San Francisco, CA, 1996.

[Gärdenfors and Makinson, 1988] P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge (Pacific Grove, CA, 1988)*, pages 83–95, Los Altos, CA, 1988. Morgan Kaufmann.

[Gärdenfors and Makinson, 1991] P. Gärdenfors and D. Makinson. Relations between the logic of theory change and nonmonotonic logic. In *The logic of theory change (Konstanz, 1989)*, pages 185–205. Springer, Berlin, 1991.

[Grove, 1988] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

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

[Hansson, To come] S. Hansson. *A Textbook of Belief Dynamics*. Kluwer Academic Publisher, To come. Manuscript.

[Katsuno and Mendelzon, 1991] H. Katsuno and A. Mendelzon. On the difference between updating a knowledge database and revising it. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, pages 387–394, Cambridge, 1991.

[Lehmann *et al.*, 1996] D Lehmann, M. Magidor, and K. Schlechta. Distance semantics for belief revision. In Y. Shoham, editor, *Proceedings of the Sixth Conference of Theoretical Aspects of Rationality and Knowledge*, pages 137–146, San Francisco, 1996. Morgan Kaufmann.

[Lehmann, 1995] D. Lehmann. Belief revision, revised. In *IJCAI-95, Vol. 1, 2 (Montreal, PQ, 1995)*, pages 1534–1540. Morgan Kaufmann, San Francisco, CA, 1995.

[Makinson, 1993] D. Makinson. Five faces of minimality. *Studia Logica. An International Journal for Symbolic Logic*, 52(3):339–379, 1993.

[Nayak, 1994] A. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis. An International Journal of Analytic Philosophy*, 41(3):353–390, 1994.

[Rott, 1992a] H. Rott. On the logic of theory change: More maps between different kinds of contraction functions. In P. Gärdenfors, editor, *Belief Revision*, volume 29, pages 122–141. Cambridge Univ. Press, Cambridge, 1992.

[Rott, 1992b] H. Rott. Preferential belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information*, 1:45–78, 1992.

[Spohn, 1987] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W. L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change and Statistics*, volume 2, pages 105–134. D. Reidel, Dordrecht, 1987.

[Williams, 1995] M. Williams. Iterated theory base change: A computational model. In *IJCAI-95 Proceeding of the 14th International Joint Conference on Artificial Intelligence*, pages 1541–1550. Morgan Kauffman, 1995.