Appeared in *Proceedings of BR'98. Trento. Italy.*

Iterable AGM functions

Carlos Areces
University of Amsterdam
ILLC
The Netherlands
carlos@wins.uva.nl

Verónica Becher
Universidad de Buenos Aires
Departamento de Computación
Argentina
vbecher@dc.uba.ar

July, 1998

Abstract

The AGM model 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: an iterable function. We formulate an iterable version of the AGM model, showing that the AGM formalism is in fact compatible with iteration. Following Alchourrón and Makinson's early reference to an iterable form of a safe contraction function, we give an iterable construction 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 is its lack of definition with respect to iterated change.

Informally, let K be a collection of data, 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 this 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*. We say that a function is iterable if its re-application is well defined. The basic idea of this article dates back to Alchourrón and Makinson's work on safe contractions [3]. 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 on a definition of an iterable function for each of the five AGM presentations: postulates [2], systems of spheres [12], meet functions [2], epistemic entrenchments [10] and safe hierarchies [3].

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. We use Greek letters α, β, γ for sentences of $\mathcal{L}$, capital letters K, K', H for theories of $\mathcal{L}$ (subsets of $\mathcal{L}$ closed under Cn), and denote with $I\!\!K$ 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 relative to the original theory. To deal with successive change it is necessary to possess *beforehand* the complete set of change functions, one for each possible theory. 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, if there are no properties linking the different change functions the result obtained can be unexpected and the corresponding behavior erratic.

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). This can be done in a qualitative way (remaining close to the AGM spirit) as in [5, 6, 14, 18, 16], or by enriching the model with numbers [7, 21, 22]. These approaches are very rich, permitting a whole new landscape of variations, but they are usually complex.

We will consider the limit case, an *iterable* change function: 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 closer to the AGM model. 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:

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

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 deterministic schemes 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* [15], 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 [4, 5] 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 iterable functions that depend on no order at all, as are the original AGM expansion and the full meet contraction function [2]. The AGM representation theorem for the function $+$ states that $K + \alpha = \mathrm{Cn}(K \cup \{\alpha\})$. Expansion inherits its iterability from the consequence relation, which is applicable to any set. 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 *global* function $+ : \mathbb{K} \times \mathcal{L} \to \mathbb{K}$ while, as it stands in the original AGM model, contraction is a *local* function relative to a theory K, $-^{K} : \mathcal{L} \to \mathbb{K}$ and it is based on a selection function s^{K} relative to the theory K.

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

where the set $K \perp \alpha$ contains the maximal subsets of K that do not imply α and the function $s^{K} : \mathcal{L} \to \mathcal{P}(K) \setminus \{\emptyset\}$ 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 function that solely depends on the explicit argument K and α, i.e. if $-$ is a full meet, again we have $- : \mathbb{K} \times \mathcal{L} \to \mathbb{K}$. Moreover, the full meet contraction function is equivalent to $K - \alpha = K \cap \mathrm{Cn}(\neg\alpha)$. Since they depend on no underlying structure, relative order or selection function, full meet contractions are applicable to any theory: they are iterable functions. The following two properties reveal the simplicity of full meet contractions functions.

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 contractions 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 [3] Alchourrón and Makinson define a new kind of contraction function based on a hierarchical order of the elements of a theory K according to what they call their "degree of safety". The safe contraction of a theory K with formula α is defined as the set of safe elements of the theory that do not imply α. They show that every safe contraction over a theory K is a partial meet contraction function over K. They also prove the converse result but 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 [19].

They also study some properties of the contraction function with respect to the intersection or union of theories and also properties of "multiple contractions". They say ([3], 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 $Z - \alpha$, when A is a theory and $A \subseteq Z$. 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 safety hierarchy.

Reusing the same fixed order makes sense, for example, when involved in hypothetical reasoning or explanation search: a fix set of facts which are known beforehand (rules of "how the world works") constitute the background knowledge from which a sequence of consistent inference steps are performed to obtain an explanation or justification[1].

The following subsections are devoted to the definition of iterable AGM functions in each of the classical presentations. 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 restricted to the search for a contraction (or revision) function (see Section 4 for further details).

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

3.1 Extended Safe Contraction

A relation $<_{sf}$ over a set A is a *hierarchy* if it is acyclic: for any set of elements $\alpha_1, \ldots, \alpha_n \in A, n \geq 1$, not $\alpha_1 <_{sf} \alpha_2 <_{sf} \ldots \alpha_n <_{sf} \alpha_1$. 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 up Cn, and let α be a sentence. The safe contraction function $-_{sf}$ is defined as

$$\begin{aligned} K -_{sf} \alpha \;=\; & \{\beta \in K \mid \text{ for every } K' \subseteq K, \text{ such that } \alpha \in \mathrm{Cn}(K') \text{ and } K' \text{ is} \\ & \subseteq\text{-minimal, } \beta \notin K' \text{ or there is } \gamma \in K' \text{ such that } \gamma <_{sf} \beta\}. \end{aligned}$$

The elements of $K -_{sf} \alpha$ are called the safe elements of K with respect to α since they can not be "blamed" of 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.

Following Alchourrón and Makinson's idea of restricting a hierarchical order, we can define the iterable safe contraction function $-_{sf}$ based on a hierarchy over all the sentences of $\mathcal{L}$.

Definition 3.1 (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 $<^K_{sf} = <_{sf} \cap (K \times K)$.*

Proposition 3.2 *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.*

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

Definition 3.3 (Iterable Safe Contraction) *Let $<_{sf}$ be a virtually connected hierarchy that continues up* Cn *in $\mathcal{L}$. The iterable AGM contraction $-_{sf} : \mathbb{K} \times \mathcal{L} \rightarrow \mathbb{K}$ is defined as*

$$\begin{aligned} K -_{sf} \alpha \;=\; & \{\beta \in K \mid \textit{ for every } K' \subseteq K, \textit{ such that } \alpha \in \mathrm{Cn}(K') \textit{ and } K' \textit{ is} \\ & \subseteq \textit{-minimal, } \beta \notin K' \textit{ or there is } \gamma \in K' \textit{ such that } \gamma <^K_{sf} \beta\}, \end{aligned}$$

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.

In the definitions above we started from a hierarchy $<_{sf}$ for $\mathcal{L}$ and defined its restriction $<_{sf}^{K}$. A valid question is whether the converse can also be achieved. Given a hierarchy for K defining a local (non iterable) contraction function $-^{K}$ can a hierarchy for $\mathcal{L}$ be defined such that the iterable functions agree with $-^{K}$ when applied to K?

Proposition 3.4 *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 $<_{sf}^{K}$ the order associated to $-^{K}$, define $<_{sf}$ as follows: $\alpha_1 <_{sf} \alpha_2$ iff $(\alpha_1, \alpha_2 \in K$ and $\alpha_1 <_{sf}^{K} \alpha_2)$ or $(\alpha_1 \in K$, $\alpha_2 \notin K$ and there is $\alpha_3 \in K$ such that $\alpha_1 <_{sf}^{K} \alpha_3)$. From the definition $<_{sf}^{K} = <_{sf} \cap (K \times K)$ and it is not hard to check that $<_{sf}$ is a virtually connected hierarchy that continues up over Cn in $\mathcal{L}$. QED

In [13] the safe contraction approach is generalized by Hansson in the kernel contraction approach. Instead of implementing a relational way of defining "safe elements", selection functions (called incision functions) are introduced which make a cut in each minimal proof of α in K as a way to define $K - \alpha$. Our results for safe contraction can be extended to kernel contraction easily.

3.2 Extended Partial Meet Contraction Functions

Selection functions provide another way to construct a contraction function. 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.

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

A selection function s is a function which, when applied to a non-empty set, returns a nonempty subset of it. Let K be a theory, we note as $s^{K} : \mathcal{L} \to \mathcal{P}(\mathcal{P}(K))$, a selection function for the sets $K \perp \alpha$, 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, which means that for each $\alpha \in \mathcal{L}$ the selection function must return the minimal 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 given 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 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 set of $\mathcal{L}$ that fails to imply α. 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 3.5 (Derived Selection Functions) *Let $s^{\mathcal{L}}$ be a selection functions for $\mathcal{L}$. Then, for any theory K the selection functions 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 3.6 *If $s^{\mathcal{L}}$ is a transitively relational selection function, then for any theory K, s^K is a transitively relational selection function.*

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

Definition 3.7 (Iterable Partial Meet Contraction) *Let $s^{\mathcal{L}}$ be a 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 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 3.8 *Let $-^{K}$ be an AGM 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 a selection function $s^{\mathcal{L}}$. Clearly, for each $H \in K \perp \alpha$ there is $H' \in \mathcal{L} \perp \alpha$ such that $H \subseteq H'$. Then we can define $s^{\mathcal{L}}(\alpha)$ by just putting $H' \in s^{\mathcal{L}}(\alpha)$ iff ($H \subseteq H'$ and $H \in s^{K}(\alpha)$). 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 minimal elements of R. It follows that $s^{\mathcal{L}}$ is transitively relational. QED

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 [12].

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

[2]Personal communication.

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, 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 3.9 (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 3.10 *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 3.11 (Iterable Sphere Contraction) *Let S be a system of spheres for $\mathcal{L}$. The iterable AGM contraction $-_{ss} : I\!\!K \times \mathcal{L} \rightarrow I\!\!K$ 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 as it is a global function. 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 by showing that any system of spheres can be transformed into another with an empty center.

Proposition 3.12 *Let S^K be a system of spheres for K. Then S^K can be extended to a system of spheres for $\mathcal{L}$.*

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

3.4 Extended Epistemic Entrenchment

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 [10] specifies 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. For any α, β in K, $\alpha \leq_{ee} (\alpha \wedge \beta)$ 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 3.13 (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:*

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

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

Proposition 3.14 *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 3.15 (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 $K -_{ee} \alpha = \{\beta \in K \mid \alpha \in \mathrm{Cn}(\emptyset)$ or $\alpha <_{ee}^{K} (\alpha \vee \beta)\}$, where $<_{ee}^{K} = \leq_{ee}^{K} - (\leq_{ee}^{K})^{-1}$, for $\leq_{ee}^{K}$ the derived epistemic entrenchment relation for K.*

It remains to show that every epistemic entrenchment contraction function can be extended to an iterable contraction function.

Proposition 3.16 *Let $\leq_{ee}^{K}$ be an epistemic entrenchment relation for K. Then there is an epistemic entrenchment relation $\leq_{ee}$ for $\mathcal{L}$ such that $\leq_{ee}|K = \leq_{ee}^{K}$ (where $R|X$ is the restriction of R to the elements in X).*

Proof. 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 $EE1$, $EE2$ and $EE5$ does not refer to the specific theory so they hold also trivially for $\mathcal{L}$. Condition $EE4$ does not apply as $\mathcal{L}$ is inconsistent. Condition $EE3$ is satisfied as follows. Take $\alpha, \beta \in \mathcal{L}$. If $\alpha, \beta \in K$ then the condition follows. If either $\alpha \notin K$ or $\beta \notin K$ then as $EE4$ is not trivial for K, either $\alpha \leq_{ee}^{K} (\alpha \wedge \beta)$ or $\beta \leq_{ee}^{K} (\alpha \wedge \beta)$. QED

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 method to restore consistency; an encoding of consistent-removal criteria. As every theory is a subset of the inconsistent theory this consistent-removal criteria can be projected 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 $K*1$- $K*9$ 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 the 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, or as reported by Friedman and Halpern in [9], the qualitative analogue of the Markov Assumption.

The revision postulate $K*9$ is sound with the interpretation of revision as a kind of hypothetical reasoning. 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 and we accommodate the last supposition in accordance with our (fixed and pre-established) criteria.

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.

A version of postulate $K-9$ appears in [2] as a property of safe contractions when multiple sets are considered (Lemma 7.4) and also in [20]. But in neither case they take this property as defining a policy of iteration. Remarkably, Freund and Lehmann [8] propose precisely the same postulate $K*9$ and show the correspondence between an AGM revision operation satisfying it and a rational consistency preserving consequence relation. They 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. While they relate it with a nonmonotonic consequence relation we consider the the iterable functions arising in the different AGM presentations.

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 in systems of spheres.

Theorem 3.17 (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 $f_S^K(\neg\alpha) = [K]$, 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. QED

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

Theorem 3.18 (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 $\alpha <_{ee}^K (\alpha \vee \beta)\}$. 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 $K-9$. 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 $\mathcal{L}$. We now prove that $-_{pm}$ and $-_{ss}$ are in fact equivalent.

Theorem 3.19 (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 one to one correspondence between systems of spheres centered in $\emptyset$ and transitive relations (that satisfy the Limit Assumption) over maximal consistent sets of $\mathcal{L}$, since the set inclusion order in a system of spheres S induces a total preorder of worlds, and conversely. QED

Finally, by using results of Rott in [19] 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 3.20 (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, $K -_{ee} \alpha = \{\beta \in K \mid$ for every $K' \subseteq K$, s.t. $\alpha \in \mathrm{Cn}(K')$ and K' $\subseteq$-minimal, $\beta \notin K'$ 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 [19], 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 [19] 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 $-_{sf}$, $-_{ee}$ are AGM functions, $K -_{sf} \alpha = K = \{\beta \in K \mid \alpha \in \mathrm{Cn}(\emptyset) \text{ or } \alpha <^{K}_{ee} (\alpha \vee \beta)\}$. If $\alpha \notin \mathrm{Cn}(\emptyset)$ and $\alpha \in K$, as the functions satisfy $K{-}9$, $K -_{sf} \alpha = (\mathcal{L} -_{sf} \alpha) \cap K$ $= (\mathcal{L} -_{ee} \alpha) \cap K = \{\beta \in K \mid \alpha \in \mathrm{Cn}(\emptyset) \text{ or } \alpha <^{K}_{ee} (\alpha \vee \beta)\}$. QED

4 Properties of Iterable AGM Functions

As we already argued, AGM full meet operations are the most elementary iterable constructions that cumply with the AGM postulates, but they are commutatitive. However, despite their modest definition our iterable AGM functions do not validate commutativity, since 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 [14] 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 [1] as an intuitive and appealing property for change functions. Functions satisfying

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

were called *permutable* and the question of under which condition an AGM function is permutable was left open in that paper. In [2] 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 functions equivalently via **Levi** or **R-Levi**.

Proposition 4.1 *Iterable AGM contraction functions are permutable.*

A direct proof of the above is immediate but the result derives from Lemma 7.1 in [2].

As iterable AGM contractions induce safe contractions functions for each theory K, the results proved in [3] carry over.

Proposition 4.2
i.) If $\alpha \in K_1 \cap K_2$ then

$$(K_1 \cap K_2) - \alpha = (K_1 - \alpha) \cap (K_2 - \alpha)$$

$$(K_1 \cup K_2) - \alpha = (K_1 - \alpha) \cup (K_2 - \alpha).$$

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.

In [7] Darwiche and Pearl present a number of properties for iterated change. In our notation, they are:

- **C1.** If $\alpha \in \mathrm{Cn}(\beta)$ then $(K * \alpha) * \beta = K * \beta$.
- **C2.** If $\neg\alpha \in \mathrm{Cn}(\beta)$ 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 [6]) are considered too demanding. Furthermore in [16] property $C2$ was proved inconsistent with AGM postulates $K{*}7$ and $K{*}8$. It is not difficult to prove the following.

Proposition 4.3
i) All iterable AGM functions, satisfy $C1$, $C3$ and $C4$.
ii) There exist iterable AGM functions violating $C2$, $C5$ and $C6$.

Condition $C2$ deals with with revision by inconsistent inputs, and its violation is to be expected in our approach. As we already said, when inconsistent hypothesis are confronted, preservation of the hypothetic reasoning carried out up to the conflictive point is not necessarily preserved.

Most noticeably, our iterable AGM functions validate six of the seven postulates of Lehman's widening ranked models [16]. Lehmann argues that these structures are suitable for iterative change and proposes seven postulates that fully characterize revision functions based on these widening 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 forces dependency between two revision steps and consequently enforces (at least to some extent) the property of "historical memory", which iterable AGM revisions lack.

Proposition 4.4
i) All iterable AGM functions, satisfy $I1$-$I6$.
ii) There exist iterable AGM functions violating $I7$.

Finally, in [17] 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. $(K_1 \cap K_2) * \alpha = (K_1 * \alpha) \cap (K_2 * \alpha)$ or $K_1 * \alpha$ or $K_2 * \alpha$.
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 4.2, in general it is not validated by our iterable AGM functions. We can prove the following proposition:

Proposition 4.5
i) All iterable AGM functions, satisfy $D2$ and $D3$.
ii) There exist iterable AGM functions violating $D1$.

5 Further Work

Given the link existing between conditionals and theory change, as pursued for example by Grahne [11] and Boutilier [6], it seems interesting to investigate a conditional logic for our iterable framework. In such a logic our contraction and revision functions would become connectives in the object language and only finite axiomatisable theories would be considered. The iterability of our functions would be reflected as logical formulae with nested occurrence of the change operators. Presumably this logic would provide further light over new properties of the iterable model and establish a closer link between theory change and the field of conditional logics.

Acknowledgments: This work has benefited from invaluable discussions with David Makinson and Hans Rott. We thank Isaac Levi for his smart interpretation of our revision operator. We also acknowledge Eduardo Fermé for his helpful comments on previous versions of this work. Carlos Areces is under grant Nr. ARG0100049 of the British Council.

References

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

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

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

[4] V. Becher. Unified semantics for revision and update, or the theory of lazy update. In *Annals of the 24th JAIIO*, Argentina, 1995.

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

[6] C. Boutilier. Iterated revision and minimal revision of conditional beliefs. *J. of Philosophical Logic*, 1995.

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

[8] M. Freund and D. Lehmann. Belief Revision and Rational Inference. Technical Report of The Leibniz Center 94-16, Hebrew University, Jerusalem, Israel. 1994.

[9] 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.

[10] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books/MIT Press, Cambridge (Mass.), 1988.

[11] G. Grahne. Updates and counterfactuals. In *P. of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 269–276, Cambridge, 1991.

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

[13] S. Hansson. Kernel contraction. *J. of Symbolic Logic*, 59:845–859, 1994.

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

[15] H. Katsuno and A. Mendelzon. On the difference between updating a knowledge database and revising it. In *P. of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 387–394, Cambridge, 1991.

[16] D. Lehmann. Belief revision, revised. In *P. of the 14th IJCAI*, pages 1534-1541, Montreal, Canada, August 1995, Morgan Kaufmann.

[17] D. Lehmann, K. Schlechta, and M. Magidor. Distance semantics for belief revision. Technical Report of The Leibniz Center 98-10, Hebrew University, Jerusalem, Israel. 1998.

[18] A. Nayak. Iterated belief change based in epistemic entrenchment *Erkenntnis*, 41:353–390, 1994.

[19] 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.

[20] H. Rott. Coherence and Conservatism in the Dynamics of Belief. Part I, Part II. Submitted to *Erkenntnis*, 1998.

[21] W. Spohn. Ordinal conditional functions. A dynamic theory of epistemic states. In W. Harper and B. Skyrms editors, *Causation in Decisions, Belief, Change and Statistics*, Dordrecht; Kluwer Academic Press, pages 105-134, 1988.

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