On Generalizing the AGM Postulates

Giorgos FLOURIS, Dimitris PLEXOUSAKIS and Grigoris ANTONIOU
Institute of Computer Science, FO.R.T.H.
P.O. Box 1385, GR 71110, Heraklion, Greece
{fgeo, dp, antoniou}@ics.forth.gr

Abstract. The AGM theory is the dominating paradigm in the field of belief change but makes some non-elementary assumptions which disallow its application to certain logics. In this paper, we recast the theory by dropping most such assumptions; we determine necessary and sufficient conditions for a logic to support operators which are compatible with our generalized version of the theory and show that our approach is applicable to a broader class of logics than the one considered by AGM. Moreover, we present a new representation theorem for operators satisfying the AGM postulates and investigate why the AGM postulates are incompatible with the foundational model. Finally, we propose a weakening of the recovery postulate which has several intuitively appealing properties.

Keywords. AGM Postulates, Contraction, Belief Revision, Belief Change

Introduction

The problem of *belief change* deals with the modification of an agent's beliefs in the face of new, possibly contradictory, information [1]. Being able to dynamically change the stored data is very important in any knowledge representation system: mistakes may have occurred during the input; some new information may have become available; or the world represented by the Knowledge Base (KB) may have changed. In all such cases, the agent's beliefs should change to reflect this fact.

The most influential work on the field of belief change was developed by Alchourron, Gärdenfors and Makinson (AGM for short) in [2]. Instead of trying to find a specific algorithm (operator) for dealing with the problem of belief change, AGM chose to investigate the properties that such an operator should have in order to be intuitively appealing. The result was a set of postulates (the *AGM postulates*) that any belief change operator should satisfy. This work had a major influence on most subsequent works on belief change, being the dominating paradigm in the area ever since and leading to a series of works, by several authors, studying the postulates' effects, providing equivalent formulations, criticizing, commenting or praising the theory etc; a list of relevant works, which is far from being exhaustive, includes [1], [3], [4], [5], [6], [7], [8], [9].

The intuition upon which the AGM postulates were based is independent of the peculiarities of the underlying logic. This fact, along with the (almost) universal acceptance of the AGM postulates in the belief change community as the defining paradigm for belief change operators, motivate us to apply them in any type of logic. Unfortunately, this is not possible because the AGM theory itself was based on certain assumptions regarding the underlying logic, such as the admittance of standard boolean operators ($\neg$, $\wedge$, $\vee$, etc), classical tautological implication, compactness and others. These assumptions are satisfied by many interesting logics, such as Propositional Logic (PL), First-Order Logic, Higher-Order Logics, modal logics etc [10]; in the following, we will refer to such logics using the term *classical* logics.

On the other hand, there are several interesting logics, which do not satisfy these assumptions, such as various logics used in Logic Programming ([11]), equational logic ([12]), Description Logics ([13]) and others (see [14], [15]). Thus, the AGM postulates (as well as most of the belief change literature) are not applicable to such (non-classical) logics. As a result, we now have a very good understanding of the change process in classical logics, but little can be said for this process in non-classical ones, leading to problems in related fields (such as ontology evolution [16], [15]).

This paper is an initial attempt to overcome this problem by generalizing the most influential belief change theory, the AGM theory. We drop most AGM assumptions and develop a version of the theory that is applicable to many interesting logics. We determine the necessary and sufficient conditions required for an operator that satisfies the AGM postulates to exist in a given logic and develop a new representation theorem for such operators. The results of our research can be further applied to shed light on our inability to develop contraction operators for belief bases that satisfy the AGM postulates [4], [7]. Finally, we present a possible weakening of the theory with several appealing properties. Only informal sketches of proofs will be presented; detailed proofs and further results can be found in the full version of this paper ([14], [15]).

1. Tarski's Framework

When dealing with logics we often use Tarski's framework (see [4]), which postulates the existence of a set of expressions (the *language*, denoted by L) and an implication operator (function) that allows us to conclude facts from other facts (the *consequence operator*, denoted by Cn). Under this framework, a *logic* is actually a pair <L,Cn> and any subset $X \subseteq L$ is a *belief* of this logic. The implications Cn(X) of a belief X are determined by the consequence operator, which is constrained to satisfy three rationality axioms, namely *iteration* ($Cn(Cn(X))=Cn(X)$), *inclusion* ($X \subseteq Cn(X)$) and *monotony* ($X \subseteq Y$ implies $Cn(X) \subseteq Cn(Y)$). For X, $Y \subseteq L$, we will say that X *implies* Y (denoted by $X \vdash Y$) iff $Y \subseteq Cn(X)$.

This general framework engulfs all monotonic logics and will be our only assumption for the underlying logic throughout this paper, unless otherwise mentioned. Thus, our framework is applicable to all monotonic logics. On the other hand, the AGM theory additionally assumes that the language of the underlying logic is closed under the standard propositional operators ($\neg$, $\wedge$, $\vee$, etc) and that the Cn operator of the logic includes classical tautological implication, is compact and satisfies the rule of Introduction of Disjunctions in the Premises.

2. The AGM Theory and the Generalized Basic AGM Postulates for Contraction

In their attempt to formalize the field of belief change, AGM defined three different types of belief change (operators), namely *expansion*, *revision* and *contraction* [2]. Expansion is the trivial addition of a sentence to a KB, without taking any special provisions for maintaining consistency; revision is similar, with the important difference that the result should be a consistent set of beliefs; contraction is required when one wishes to consistently remove a sentence from their beliefs instead of adding

Table 1. Original and Generalized AGM Postulates

AGM Postulate	Original Version	Generalized Version
(K−1) Closure	K−x is a theory	K−X=Cn(K−X)
(K−2) Inclusion	K−x⊆K	K−X⊆Cn(K)
(K−3) Vacuity	If x∉Cn(K), then K−x=K	If X⊈Cn(K), then K−X=Cn(K)
(K−4) Success	If x∉Cn(∅), then x∉Cn(K−x)	If X⊈Cn(∅), then X⊈Cn(K−X)
(K−5) Preservation	If Cn({x})=Cn({y}), then K−x=K−y	If Cn(X)=Cn(Y), then K−X=K−Y
(K−6) Recovery	K⊆Cn((K−x)∪{x})	K⊆Cn((K−X)∪X)

one. AGM introduced a set of postulates for revision and contraction that formally describe the properties that such an operator should satisfy. This paper focuses on contraction, which is the most important operation for theoretical purposes [1], but some ideas on expansion and revision will be presented as well.

Under the AGM approach, a KB is a set of propositions K closed under logical consequence (K=Cn(K)), also called a *theory*. Any expression x∈L can be contracted from the KB. The operation of contraction can thus be formalized as a function mapping the pair (K, x) to a new KB K′ (denoted by K′=K−x). Similarly, expansion is denoted by K+x and revision is denoted by K∔x.

Of course, not all functions can be used for contraction. Firstly, the new KB (K−x) should be a theory itself (see (K−1) in table 1). As already stated, contraction is an operation that is used to *remove* knowledge from a KB; thus the result should not contain any new, previously unknown, information (K−2). Moreover, contraction is supposed to return a new KB such that the expression x is no longer believed; hence x should not be among the consequences of K−x (K−4). Finally, the result should be syntax-independent (K−5) and should remove as little information from the KB as possible (K−3), (K−6), in accordance with the *Principle of Minimal Change*. These intuitions were formalized in a set of six postulates, the *basic AGM postulates for contraction*, which can be found in table 1 (see also [2]). Reformulating these postulates to be applicable in our more general context is relatively easy; the generalized version of each postulate can be found in the rightmost column of table 1.

Notice that no assumptions (apart from the logic being expressible as an <L,Cn> pair) are made in the generalized set. In addition, the restriction on the contracted belief being a single expression was dropped; in our context, the contracted belief can be any set of propositions, which is read conjunctively. This is a crucial generalization, as, in our context (unlike the classical one), the conjunction of a set of propositions cannot be always expressed using a single one (see also [15]). The fact that X is read conjunctively underlies the formulation of the vacuity and success postulates and implies that we are dealing with *choice contraction* in the terminology of [17]. Moreover, the KB was allowed to be any set (not necessarily a theory). This is a purely technical relaxation done for aesthetic reasons related to the symmetry of the contraction operands. The full logical closure of K is considered when determining the contraction result, so this relaxation should not be confused with *foundational* approaches using *belief bases* ([18], [19]). It is trivial to verify that, in the presence of the AGM assumptions, each of the original AGM postulates is equivalent to its generalized counterpart. In the following, any reference to the AGM postulates will refer to their generalized version, unless stated otherwise.

3. Decomposability and AGM-compliance

One of the fundamental results related to the original AGM postulates was that all classical logics admit several contraction operators that satisfy them [2]. Unfortunately, it turns out that this is not true for non-classical logics (and the generalized postulates). Take for example $L=\{x,y\}$, $Cn(\emptyset)=\emptyset$, $Cn(\{x\})=\{x\}$, $Cn(\{y\})=Cn(\{x,y\})=\{x,y\}$. It can be easily shown that $<L,Cn>$ is a logic. Notice that, in this logic, all theories (except $Cn(\emptyset)$) contain x. So, if we attempt to contract $\{x\}$ from $\{x,y\}$ the result can be no other than $\emptyset$, or else the postulate of success would be violated. But then, the recovery postulate is not satisfied, as can be easily verified. Thus, no contraction operator can handle the contraction $\{x,y\}-\{x\}$ in a way that satisfies all six AGM postulates. As the following theorem shows, it is the recovery postulate that causes this problem:

Theorem 1. In every logic there exists a contraction operator satisfying (K−1)-(K−5).

Once the recovery postulate is added to our list of desirable postulates, theorem 1 fails, as shown by our counter-example. One of the major questions addressed in this paper is the identification of the distinctive property that does not allow the definition of a contraction operator that satisfies the AGM postulates in some logics. The answer to this question can be found by examining the situation a bit closer. Consider two beliefs $K, X\subseteq L$ and the contraction operation $Z=K-X$, which, supposedly, satisfies the AGM postulates. Let's first consider some trivial, limit cases:

- If $Cn(X)=Cn(\emptyset)$, then the recovery and closure postulates leave us little choice but to set $Z=Cn(K)$.
- If $Cn(X)\nsubseteq Cn(K)$, then the vacuity postulate forces us to set $Z=Cn(K)$.

These remarks guarantee that a proper Z can be found in both these cases, regardless of the logic at hand. The only case left unconsidered is the principal one, i.e., when $Cn(\emptyset)\subset Cn(X)\subseteq Cn(K)$. In this case, the postulates imply two main restrictions:

1. The result (Z) should be implied by K (inclusion postulate); however, Z cannot imply K, due to (K−4) and our hypothesis; thus: $Cn(Z)\subset Cn(K)$.
2. The union of X and Z should imply K (recovery postulate); however, our hypothesis and (K−2) imply that both X and Z are implied by K, so their union is implied by K as well. Thus: $Cn(X\cup Z)=Cn(K)$.

It turns out that the existence of a set Z satisfying these two restrictions for each (K,X) pair with the aforementioned properties is a necessary and sufficient condition for the existence of an operator satisfying the generalized (basic) AGM postulates for contraction in a given logic. To show this fact formally, we will need some definitions:

Definition 1. Consider a logic $<L,Cn>$ and two beliefs $K, X\subseteq L$. We define the set of *complement beliefs* of X with respect to K, denoted by $X^-(K)$ as follows:

- If $Cn(\emptyset)\subset Cn(X)\subseteq Cn(K)$, then $X^-(K)=\{Z\subseteq L \mid Cn(Z)\subset Cn(K)$ and $Cn(X\cup Z)=Cn(K)\}$
- In any other case, $X^-(K)=\{Z\subseteq L \mid Cn(Z)=Cn(K)\}$

Definition 2. Consider a logic $<L,Cn>$ and a set $K\subseteq L$:

- The logic $<L,Cn>$ is called *AGM-compliant with respect to the basic postulates for contraction* (or simply *AGM-compliant*) iff there exists an operator that satisfies the basic AGM postulates for contraction (K−1)-(K−6)
- The set K is called *decomposable* iff $X^-(K)\neq\emptyset$ for all $X\subseteq L$
- The logic $<L,Cn>$ is called *decomposable* iff all $K\subseteq L$ are decomposable

As already mentioned, the two notions of definition 2 actually coincide:

Theorem 2. A logic $<L,Cn>$ is AGM-compliant iff it is decomposable.

Theorem 2 provides a necessary and sufficient condition to determine whether any given logic admits a contraction operator that satisfies the AGM postulates. The following result shows that the set of complement beliefs ($X^{-}(K)$) of a pair of beliefs K, X contains exactly the acceptable (per AGM) results of K−X, modulo Cn:

Theorem 3. Consider a decomposable logic <L,Cn> and an operator '−'. Then '−' satisfies the basic AGM postulates for contraction iff for all K, $X \subseteq L$, there is some $Z \in X^{-}(K)$ such that $Cn(Z)=K-X$.

Theorem 3 provides yet another representation theorem for the AGM postulates for contraction, equivalent to partial meet functions [2], safe contraction [3], systems of spheres [6], epistemic entrenchment [8] etc. However, our representation theorem is unique in the sense that it is applicable to all monotonic logics, unlike the other results whose formulation assumes that the underlying logic satisfies the AGM assumptions.

As already mentioned, all classical logics admit several operators satisfying the original AGM postulates. This should also be true for the generalized postulates, so, we would expect all classical logics to be AGM-compliant. This is indeed true:

Theorem 4. All logics satisfying the AGM assumptions are AGM-compliant.

4. More Characterizations of AGM-compliance: Cuts and Max-cuts

Apart from theorem 2, an alternative (and equivalent) method for determining the AGM-compliance of a given logic can be devised, based on the notion of a *cut*. A cut is a family of subsets of Cn(K) with the property that every other subset of Cn(K) either implies or is implied by one of them. Such a family actually "cuts" the beliefs implied by K in two (thus the name). Cuts are related to AGM-compliance, so they have a special place in our theory. Formally, we define a cut as follows:

Definition 3. Consider a logic <L,Cn>, a set $K \subseteq L$ and a family $\wp$ of beliefs such that:

- For all $Y \in \wp$, $Cn(Y) \subset Cn(K)$
- For all $Z \subseteq L$ such that $Cn(Z) \subset Cn(K)$, there is a $Y \in \wp$ such that $Cn(Y) \subseteq Cn(Z)$ or $Cn(Z) \subseteq Cn(Y)$

Then $\wp$ is called a *cut* of K.

Now, consider a set $K \subseteq L$, a cut $\wp$ of K and a non-tautological set X that is implied by all the sets in the cut $\wp$. Take Z=K−X. Since $\wp$ is a cut and $Cn(Z) \subset Cn(K)$ by (K−2) and (K−4), Z will either imply or be implied by a set in $\wp$. If it implies a set in $\wp$, then it also implies X, so the operation Z=K−X does not satisfy success. If it is implied by a set in $\wp$ (say $Y \in \wp$), then both X and Z are implied by Y, so it is necessarily the case that $Cn(X \cup Z) \subseteq Cn(Y) \subset Cn(K)$, so recovery is not satisfied. Once we deal with some technicalities and limit cases, it turns out that this idea forms the basis for another equivalent characterization of AGM-compliant logics:

Theorem 5. Consider a logic <L,Cn> and a belief $K \subseteq L$. Then:

- K is decomposable iff $Cn(\cap_{Y \in \wp} Cn(Y)) = Cn(\emptyset)$ for every cut $\wp$ of K
- The logic <L,Cn> is AGM-compliant iff for every $K \subseteq L$ and every cut $\wp$ of K it holds that $Cn(\cap_{Y \in \wp} Cn(Y)) = Cn(\emptyset)$.

Cuts enjoy an interesting monotonic behavior; informally, the "larger" the beliefs that are contained in a cut the more likely it is to have a non-tautological intersection. This motivates us to consider the family of all maximal (proper) subsets of Cn(K), which is a special type of a cut, containing the strongest (largest) beliefs possible:

Definition 4. Consider a logic $<L,Cn>$, a set $K\subseteq L$ and a family $\wp$ of beliefs such that:

- For all $Y\in\wp$, $Cn(Y)\subset Cn(K)$
- For all $Z\subseteq L$ such that $Cn(Z)\subset Cn(K)$, there is a $Y\in\wp$ such that $Cn(Z)\subseteq Cn(Y)$
- For all $Y\in\wp$, $Y=Cn(Y)$
- For all Y, $Y'\in\wp$, $Cn(Y)\subseteq Cn(Y')$ implies $Y=Y'$

Then $\wp$ is called a *max-cut* of K.

It is easy to verify that a max-cut is a cut that contains exactly the maximal proper subsets of K. There is no guarantee that every belief will have a max-cut, but, if it does, then it is unique. The importance of max-cuts stems from the fact that, if a max-cut $\wp$ of K exists and $Cn(\cap_{Y\in\wp}Cn(Y))=Cn(\varnothing)$, then the same property will hold for all cuts of K as well (due to monotonicity). This fact allows us to determine the decomposability of a given belief with a single check, formally described in the following theorem:
Theorem 6. Consider a logic $<L,Cn>$ and a belief $K\subseteq L$ which has a max-cut $\wp$. Then K is decomposable iff $Cn(\cap_{Y\in\wp}Cn(Y))=Cn(\varnothing)$.

5. A Foundational AGM Theory

Knowledge in a KB can be represented using either *belief bases* or *belief sets* [18]. Belief sets are closed under logical consequence, while belief bases can be arbitrary sets of expressions. Belief sets contain explicitly all the knowledge deducible from the KB (so they are often large and infinite sets), while belief bases are (usually) small and finite, because they contain only the beliefs that were explicitly acquired through an observation, rule, measurement, experiment etc [9]; the rest of our beliefs are simply deducible from such facts via the inference mechanism (Cn) of the underlying logic.

The selected representation affects the belief change algorithms considered. In the belief set approach, all the information is stored explicitly, so there is no distinction between implicit and explicit knowledge. On the other hand, when changes are performed upon a belief base, we temporarily have to ignore the logical consequences of the base; in effect, there is a clear distinction between knowledge stored explicitly in the base (which can be changed directly) and knowledge stored implicitly (which cannot be changed, but is indirectly affected by the changes in the explicit knowledge). Thus, in the belief base approach, our options for the change are limited to the base itself, so the changes are (in general) more coarse-grained. On the other hand, belief sets are (usually) infinite structures, so, in practice, only a small subset of a belief set is explicitly stored; however, in contrast to the belief base approach, the implicit part of our knowledge is assumed to be of equal value to the explicitly stored one. Some thoughts on the connection between the two approaches appear in [4], [18].

A related philosophical consideration studies how our knowledge *should* be represented. There are two viewpoints here: *foundational* theories and *coherence* theories [19]. Under the foundational viewpoint, each piece of our knowledge serves as a justification for other beliefs; this viewpoint is closer to the belief base approach. On the other hand, under the coherence model, no justification is required for our beliefs; a belief is justified by how well it fits with the rest of our beliefs, in forming a coherent and mutually supporting set; thus, every piece of knowledge helps directly or indirectly to support any other. Coherence theories match with the use of belief sets as a proper knowledge representation format.

The AGM theory is based on the coherence model (using belief sets); however, some authors have argued that the foundational model is more appropriate for belief change [4], [9]. It can be shown that the foundational model is incompatible with the

Table 2. Generalized AGM Postulates (Coherence and Foundational Model)

AGM Postulate (Foundational Model)	Generalized Version (Coherence Model)	Generalized Version (Foundational Model)
(B–1) Base Closure	$K-X=Cn(K-X)$	$K-X\subseteq L$
(B–2) Base Inclusion	$K-X\subseteq Cn(K)$	$K-X\subseteq K$
(B–3) Base Vacuity	If $X\not\subseteq Cn(K)$, then $K-X=Cn(K)$	If $X\not\subseteq Cn(K)$, then $K-X=K$
(B–4) Base Success	If $X\not\subseteq Cn(\emptyset)$, then $X\not\subseteq Cn(K-X)$	If $X\not\subseteq Cn(\emptyset)$, then $X\not\subseteq Cn(K-X)$
(B–5) Base Preservation	If $Cn(X)=Cn(Y)$, then $K-X=K-Y$	If $Cn(X)=Cn(Y)$, then $K-X=K-Y$
(B–6) Base Recovery	$K\subseteq Cn((K-X)\cup X)$	$K\subseteq Cn((K-X)\cup X)$

AGM postulates in classical logics, because there are no operators satisfying the foundational version of the AGM postulates in such logics [4]. However, this result leaves open the question of whether there are any (non-classical) logics which are compatible with a generalized version of the foundational AGM postulates.

To address this question, we took a path similar to the one taken for the standard (coherence) case, by introducing a generalized foundational version of the AGM postulates (see table 2). This set of postulates can be derived either by generalizing the foundational postulates that appeared in [4], or by reformulating the postulates that appeared in table 1 so as to be applicable for the foundational case (the latter method is followed in table 2). Note that, in order to avoid confusion with the coherence model, a different notation was used in table 2 for the numbering of the foundational postulates. Notice that the main difference between the two sets of postulates is the fact that in belief set contraction the result should be a subset of the logical closure of the KB, while in belief base contraction the result should be a subset of the KB itself; moreover, the result of a contraction operation need not be a theory in the foundational case.

These seemingly small differences have some severe effects on the operators considered. Take for example the operation $\{x\wedge y\}-\{x\}$ in PL. Due to the base inclusion and base success postulates it should be the case that $\{x\wedge y\}-\{x\}=\emptyset$; but this violates the base recovery postulate, as can be easily verified. Thus, there can be no AGM-compliant base contraction operator that can handle this case. This simple example can be easily adapted to show that no classical logic can admit an operation satisfying the foundational version of the AGM postulates [4].

But what is the property that allows a logic to admit an AGM-compliant contraction operator for belief bases? The answer is simpler than one would expect. As already mentioned, the only additional requirement that is posed by our transition to the foundational model is the fact that the result should now be a subset of the belief base (instead of its logical closure, i.e., the belief set). Thus, the ideas that were used in the coherence case could be applied here as well, with the additional requirement that the result is a subset of the base. More formally:

Definition 5. Consider a logic $<L,Cn>$ and a set $K\subseteq L$:

- The logic $<L,Cn>$ is called *AGM-compliant with respect to the basic postulates of belief base contraction* (or simply *base-AGM-compliant*) iff there exists an operator that satisfies the basic foundational AGM postulates
- The set K is called *base-decomposable* iff $X^{-}(K)\cap P(K)\neq\emptyset$ for all $X\subseteq L$
- The logic $<L,Cn>$ is called *base-decomposable* iff all $K\subseteq L$ are base-decomposable

In definition 5, P(K) stands for the powerset of K, i.e., the family of all subsets of K (also denoted by 2^K). Using this definition, we can show the counterparts of theorems 1, 2 and 3 for the foundational case:

Theorem 7. In every logic there exists a contraction operator satisfying (B−1)-(B−5).

Theorem 8. A logic <L,Cn> is base-AGM-compliant iff it is base-decomposable.

Theorem 9. Consider a base-decomposable logic <L,Cn> and an operator '−'. Then '−' satisfies the basic foundational AGM postulates for contraction iff for all $K, X \subseteq L$, for which $Cn(X) \not\subseteq Cn(K)$, it holds that $K-X=K$ and for all $K, X \subseteq L$ for which $Cn(X) \subseteq Cn(K)$ it holds that $K-X \in X^{-}(K) \cap P(K)$.

As is obvious by theorems 7, 8 and 9, there is an exceptional symmetry between the foundational and coherence versions of both the postulates and the concepts and results related to the postulates. This symmetry is also extended to cuts:

Definition 6. Consider a logic <L,Cn>, a set $K \subseteq L$ and a family $\wp$ of beliefs such that:

- For all $Y \in \wp$, $Cn(Y) \subset Cn(K)$
- For all $Z \subseteq L$ such that $Cn(Z) \subset Cn(K)$ and $Z \subset K$, there is a $Y \in \wp$ such that $Cn(Y) \subseteq Cn(Z)$ or $Cn(Z) \subseteq Cn(Y)$

Then $\wp$ is called a *base cut* of K.

Theorem 10. Consider a logic <L,Cn> and a belief $K \subseteq L$. Then:

- K is base-decomposable iff $Cn(\cap_{Y \in \wp} Cn(Y)) = Cn(\varnothing)$ for every base cut $\wp$ of K
- The logic <L,Cn> is base-AGM-compliant iff for every $K \subseteq L$ and every base cut $\wp$ of K it holds that $Cn(\cap_{Y \in \wp} Cn(Y)) = Cn(\varnothing)$

At this point it would be anticipated (and desirable) that the counterpart of a max-cut for the foundational case could be defined (i.e., base-max-cuts). Unfortunately, the different nature of the foundational case makes the definition of base-max-cuts problematic; for this reason, base-max-cuts are not studied in this paper.

The theorems and definitions of this section show that base-AGM-compliance is a strictly stronger notion than AGM-compliance:

Theorem 11. If a logic <L,Cn> is base-AGM-compliant, then it is AGM-compliant.

6. Discussion

The original AGM theory focused on a certain class of logics and determined the properties that a belief change operator should satisfy in order to behave rationally. Our approach could be considered as taking the opposite route: starting from the notion of rationality, as defined by the AGM postulates, we tried to determine the logics in which a rational (per AGM) contraction operator can be defined. Our results determine the necessary and sufficient conditions for a logic to admit a rational contraction operator under both the coherence and the foundational model. The relation between the various classes of logics studied in this paper is shown in figure 1.

In section 3, we showed the intuitively expected result that all classical logics are AGM-compliant. A related question is whether there are any non-classical logics which are AGM-compliant (or base-AGM-compliant). The answer is positive; take $L=\{x_i \mid i \in I\}$ for any index set I with at least two indices and set $Cn(\varnothing)=\varnothing$, $Cn(\{x_i\})=\{x_i\}$ and

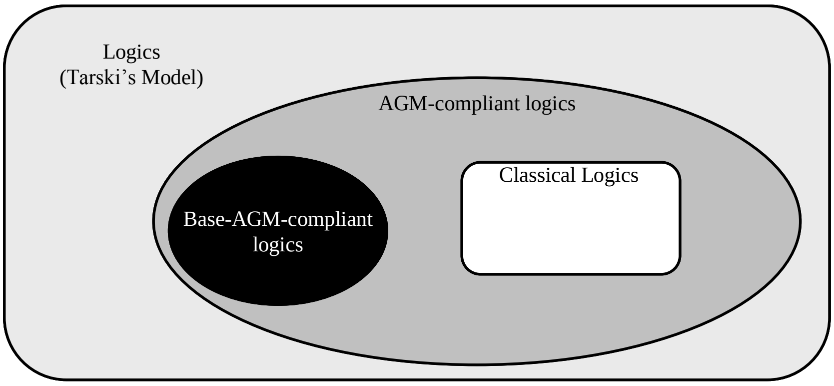

Figure 1. Topology of AGM-compliance

Cn(X)=L for all X such that |X|≥2. It can be easily shown that this construct is a base-AGM-compliant, non-classical logic (thus, it is AGM-compliant as well).

The example just presented is artificial; in [15], [20], it is shown that there are well-known, interesting, non-classical logics, used in real world applications (such as certain Description Logics), which are AGM-compliant. In such logics, the original AGM theory is not applicable, but our generalized theory is. Unfortunately, for the foundational case the situation is not so nice: the following theorem (theorem 12) shows that we should not expect to find any interesting base-AGM-compliant logic, while its corollary (theorem 13) recasts the result of [4] in our terminology:

Theorem 12. Consider a logic $<L,Cn>$. If there is a proposition $x \in L$ and a set $Y \subseteq L$ such that $Cn(\varnothing) \subset Cn(Y) \subset Cn(\{x\})$, then the logic is not base-AGM-compliant.

Theorem 13. No logic satisfying the AGM assumptions is base-AGM-compliant.

Our results could be further exploited by studying the connection of AGM-compliance with the concept of *roots* [14], [15] which is quite instructive on the intuition behind decomposability. Furthermore, an interesting connection of AGM-compliance with lattice theory [21] can be shown; AGM-compliance is a property that can be totally determined by the structure of the (unique, modulo equivalence) lattice that represents the logic under question [14], [15]. This is an important result for two reasons; first, lattice theory provides a nice visualization of the concepts and results presented in this paper; second, the various results of lattice theory that have been developed over the years can be used directly in our framework. The latter property may allow the development of deeper results regarding AGM-compliant logics.

As already mentioned, AGM dealt with two more operations, namely revision and expansion. Expansion can be trivially defined in the standard framework [2] using the equation: $K+x=Cn(K \cup \{x\})$. This equation can be easily recast in our framework, by setting $K+X=Cn(K \cup X)$. For revision, however, things are not so straightforward. For reasons thoroughly explained in [15], the reformulation of the AGM postulates for revision in our more general context requires the definition of a negation-like operator applicable to all monotonic logics. The exact properties of such an operator are a subject of current investigation [22]. A related question is whether contraction and revision in this context are interdefinable in a manner similar to the one expressed via the Levi and Harper identities in the classical setting [1].

Apart from the basic AGM postulates, AGM proposed two additional contraction postulates, the *supplementary* ones, for which a similar generalization can be done. Under certain conditions (satisfied by all finite logics), decomposability is a necessary and sufficient condition for the existence of a contraction operator satisfying both the basic and the supplementary postulates. However, it is an open problem whether this result extends to all logics; thus, the integration of the supplementary postulates in our framework is still incomplete. For more details on this issue see [15].

7. Towards a Universal AGM Theory

As already mentioned, the AGM postulates express common intuition regarding the operation of contraction and were accepted by most researchers. The only postulate that has been seriously debated is the postulate of recovery (K−6) and its foundational counterpart (B−6). Some works [4] state that (K−6) is counter-intuitive, while others [7], [8] state that it forces a contraction operator to remove too little information from the KB. However, it is generally acceptable that (K−6) (and (B−6)) cannot be dropped unless replaced by some other constraint that would somehow express the Principle of Minimal Change [1]. Theorems 1 and 7 show that the recovery (and base recovery) postulate is also the reason that certain logics are non-AGM-compliant (and non-base-AGM-compliant). These facts motivate us to look for a replacement of this postulate (say (K−6*) and (B−6*)) that would properly capture the Principle of Minimal Change in addition to being as close as possible to the original (K−6) and (B−6).

Let us first consider the coherence model. In order for such a postulate to be adequate for our purposes, it should satisfy the following important properties:

1. **Existence:** In every logic <L,Cn> there should exist a contraction operator satisfying (K−1)-(K−5) and (K−6*)
2. **AGM Rationality:** In AGM-compliant logics, the class of operators satisfying (K−1)-(K−5) and (K−6*) should coincide with the class of operators satisfying (K−1)-(K−6).

It turns out that the following postulate comes very close to satisfying both goals:

(K−6*) If $Cn((K-X)\cup X)\subset Cn(Y\cup X)$ for some $Y\subseteq Cn(K)$, then $Cn(\varnothing)\subset Cn(X)\subseteq Cn(Y)$

The original recovery postulate guarantees that the contraction operator will remove as little information from K as possible, by requiring that the removed expressions are all "related" to X; thus, the union of (K−X) with X is equivalent to K. Notice that, in the principal case of contraction (where $Cn(\varnothing)\subset Cn(X)\subseteq Cn(K)$), this result is maximal, due to the inclusion postulate. The idea behind (K−6*) is to keep the requirement that $Cn((K-X)\cup X)$ should be maximal, while dropping the requirement that this "maximal" is necessarily equivalent to K. Indeed, according to (K−6*), $Cn((K-X)\cup X)$ is maximal, because, if there is any Y that gives a "larger" union with X than (K−X), then this Y is necessarily a superset of X, so it cannot be considered as a result of (K−X), due to the success postulate.

Let's first verify that the second property (AGM rationality) is satisfied by (K−6*). Firstly, if postulate (K−6) is satisfied, then there is no Y satisfying the "if part" of (K−6*), so (K−6*) is satisfied as well. For the opposite, if the logic is AGM-compliant, then, by theorem 3, for every pair K, $X\subseteq L$ such that $Cn(\varnothing)\subset Cn(X)\subseteq Cn(K)$ there is some set $Z\subseteq L$ such that $Z\in X^{-}(K)$. For this particular Z it holds that $Cn(X)\nsubseteq Cn(Z)$ and $Cn(Z\cup X)=Cn(K)$, so $Cn(Z\cup X)$ is maximal among all $Z\subseteq Cn(K)$. Therefore, (K−6*) will not be satisfied unless Cn((K−X)∪X)=Cn(K), i.e., (K−6*) will be satisfied only for contraction operators for which (K−6) is also satisfied. Once we deal with some technicalities and limit cases we can show the following:

Theorem 14. Consider an AGM-compliant logic <L,Cn> and a contraction operator '−'. Then '−' satisfies (K−1)-(K−6) iff it satisfies (K−1)-(K−5) and (K−6*).

Unfortunately, the first of our desired properties, existence, cannot be guaranteed in general. However, finiteness, as well as decomposability, imply existence; this implication could have a "local" or a "global" character:

Theorem 15. Consider a logic <L,Cn>.

- There is a contraction operator (K−X) which satisfies (K−1)-(K−5) and (K−6*) whenever the belief set Cn(K) is decomposable or it contains a finite number of subsets, modulo logical equivalence
- If the logic is decomposable, or it contains a finite number of beliefs, modulo logical equivalence, then there is a contraction operator satisfying (K−1)-(K−5) and (K−6*)

As usual, similar results can be shown for the foundational case; the only change required is to restrict our search to all Y which are subsets of the belief base K, instead of its logical closure (Cn(K)). Moreover, notice that, in this case, the assumption of finiteness of the belief base is reasonable, so existence is (usually) achieved [22]:

(B−6*) If $Cn((K-X)\cup X)\subset Cn(Y\cup X)$ for some $Y\subseteq K$, then $Cn(\emptyset)\subset Cn(X)\subseteq Cn(Y)$

Theorem 16. Consider a base-AGM-compliant logic <L,Cn> and a contraction operator '−'. Then '−' satisfies (B−1)-(B−6) iff it satisfies (B−1)-(B−5) and (B−6*).

Theorem 17. Consider a logic <L,Cn>.

- There is a contraction operator (K−X) which satisfies (B−1)-(B−5) and (B−6*) whenever the belief base K is base-decomposable or it contains a finite number of subsets, modulo logical equivalence
- If the logic is base-decomposable, or it contains a finite number of beliefs, modulo logical equivalence, then there is a contraction operator satisfying (B−1)-(B−5) and (B−6*)

8. Conclusion and Future Work

We studied the application of the AGM theory in a general framework including all monotonic logics. We determined the limits of our generalization and showed that there are non-classical logics in which the AGM theory can be applied. We provided a new representation result for the AGM postulates, which has the advantage of being applicable to all monotonic logics, rather than just the classical ones, and studied the connection of the AGM theory with the foundational model. Finally, we proposed a weakening of the recovery postulate with several intuitively appealing properties. For more details refer to [14], [15]; for some applications of our work refer to [15], [20].

Future work includes the further refinement of (K−6*), (B−6*) so as to satisfy the existence property unconditionally, as well as the study of the operation of revision and the supplementary postulates. Finally, we plan to study the connection of other AGM-related results to AGM-compliance. Our ultimate objective is the development of a complete "AGM-like" theory suitable for all logics, including non-classical ones.

Acknowledgements

This paper is a heavily revised and extended version of [23]. The authors would like to thank Zhisheng Huang, Jeff Pan and Holger Wache for permission to include yet unpublished joint work (to appear in [22]).

References

[1] P. Gärdenfors. Belief Revision: An Introduction. In P. Gärdenfors (ed.) Belief Revision, pp. 1-20. Cambridge University Press, 1992.
[2] C. Alchourron, P. Gärdenfors, D. Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. Journal of Symbolic Logic 50:510-530, 1985.
[3] C. Alchourron, D. Makinson. On the Logic of Theory Change: Safe Contraction. Studia Logica 44:405-422, 1985.
[4] A. Fuhrmann. Theory Contraction through Base Contraction. Journal of Philosophical Logic 20:175-203, 1991.
[5] P. Gärdenfors, D. Makinson. Revisions of Knowledge Systems using Epistemic Entrenchment. In Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge, 1988.
[6] A. Grove. Two Modellings for Theory Change. Journal of Philosophical Logic, 17:157-170, 1988.
[7] S.O. Hansson. Knowledge-level Analysis of Belief Base Operations. Artificial Intelligence 82:215-235, 1996.
[8] S.O. Hansson. A Textbook of Belief Dynamics. Kluwer Academic Publishers, 1999.
[9] B. Nebel. A Knowledge Level Analysis of Belief Revision. In Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning (KR'89), pp. 301-311, 1989.
[10] H.B. Enderton. A Mathematical Introduction to Logic. Academic Press, New York, 1972.
[11] J. Lloyd. Foundations of Logic Programming. Springer, 1987.
[12] S.N. Burris. Logic For Mathematics And Computer Science. Prentice Hall, New Jersey, 1998.
[13] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. Patel-Schneider. The Description Logic Handbook: Theory, Implementation and Applications. Cambridge University Press, 2002.
[14] G. Flouris, D. Plexousakis, G. Antoniou. AGM Postulates in Arbitrary Logics: Initial Results and Applications, Technical Report, TR-336, ICS-FORTH, 2004.
[15] G. Flouris. On Belief Change and Ontology Evolution. Doctoral Dissertation, Department of Computer Science, University of Crete, February 2006.
[16] P. Haase, Y. Sure. D3.1.1.b State of the Art on Ontology Evolution. 2004. Available on the Web (last visited April, 2006): http://www.aifb.uni-karlsruhe.de/WBS/ysu/publications/SEKT-D3.1.1.b.pdf
[17] A. Fuhrmann, S.O. Hansson. A Survey of Multiple Contractions. Journal of Logic, Language and Information, 3:39-76, 1994.
[18] S.O. Hansson. Revision of Belief Sets and Belief Bases. In Handbook of Defeasible Reasoning and Uncertainty Management Systems, Vol. 3, pp. 17-75, 1998.
[19] P. Gärdenfors. The Dynamics of Belief Systems: Foundations Versus Coherence Theories. Revue Internationale de Philosophie, 44:24-46, 1992.
[20] G. Flouris, D. Plexousakis, G. Antoniou. On Applying the AGM Theory to DLs and OWL. In Proceedings of the 4th International Semantic Web Conference (ISWC-05), pp. 216-231, 2005.
[21] G. Grätzer. Lattice Theory: First Concepts and Distributive Lattices. W. H. Freeman & co, 1971.
[22] G. Flouris, Z. Huang, J. Pan, D. Plexousakis, H. Wache. Inconsistencies, Negations and Changes in Ontologies. To appear in Proceedings of the 21st National Conference on Artificial Intelligence, 2006.
[23] G. Flouris, D. Plexousakis, G. Antoniou. Generalizing the AGM Postulates: Preliminary Results and Applications. In Proceedings of the 10th International Workshop on Non-Monotonic Reasoning, 2004.