An Inconsistency Tolerant Model for Belief Representation and Belief Revision

Samir Chopra
Department of Philosophy
City University of New York (GSUC)
New York, NY, U.S.A
sc57@is4.nyu.edu

Rohit Parikh
Computer Science, Mathematics and Philosophy
City University of New York (GSUC)
New York, NY, U.S.A
ripbc@cunyvm.cuny.edu

July 21, 1999

Abstract

We propose a new model for representing and revising belief structures, which relies on a notion of partial language splitting and tolerates some amount of inconsistency while retaining classical logic. The model preserves an agent's ability to answer queries in a coherent way using Belnap's four-valued logic. Axioms analogous to the AGM axioms hold for this new model. The distinction between implicit and explicit beliefs is represented and psychologically plausible, computationally tractable procedures for query answering and belief base revision are obtained.

1 Introduction

This paper is motivated by several concerns in the area of belief representation and belief revision.

Minimal Change: When a belief structure is revised in the face of new information, we would like the new belief structure to be as similar to the old one as possible. This requirement goes back to Quine and has been repeatedly emphasized by Gärdenfors as the *preservation criterion* (cf. [Gärdenfors, 1988].

Computational Tractability: Since we are seeking to represent the beliefs of real people, and since the satisfiability problem is **NP**-complete, even at the propositional level we need to be aware of the computational limitations of real people or even real computers and try, if possible, to take them into account in our belief representation models.

Explicit vs Implicit Beliefs: There are beliefs which are explicit - they may be actually asserted or at least agreed to. Other beliefs are implicit in the sense that an agent can be made to agree to them after a discussion. If the set of explicit beliefs is consistent, then we can define the set of implicit beliefs as just their logical consequences. However, this becomes implausible if – unknown to the agent – the explicit beliefs are inconsistent.

Inconsistency Tolerance: The beliefs of real individuals are usually inconsistent from a global point of view. This is also likely to be the case with large data bases. So we need to model how an agent deals with possible inconsistencies in its belief base.

In this paper we propose a model which comes to grips with all four issues. Our work is related to the *splitting languages* framework of [Parikh, 1996]. That paper showed how an agent's beliefs can be (uniquely) subdivided into sub-areas and how this division can be used in belief revision. However, it did not address the issue of explicit vs implicit beliefs, not did it tackle the question of inconsistent explicit beliefs. In the current paper we do both.

We take the celebrated AGM axioms [Alchourron *et. al*, 1985] as our starting point, but modify them somewhat and represent an agent's beliefs, not by a *theory* as is usual, but by what we call a B-structure, a notion which generalizes the notion of a theory. We begin by giving some notation and re-stating the AGM axioms.

Notation: In the following, L is a finite propositional language. (However, most results continue to hold for a countably infinite first order language without equality.) The constants *true, false* are in L. The letter L stands both for a set of propositional symbols and for the formulae generated by that set. It will be clear from the context which is meant. $A \Leftrightarrow B$ means that $A \leftrightarrow B$ is a tautology, i.e. true under all truth assignments. Similarly, $A \Rightarrow B$ means that $A \rightarrow B$ is a tautology. If X is a set of formulae then $Cn(X)$ is the logical closure of X. Thus X is a theory iff $X = Cn(X)$. Letters T, T' denote theories. $T * A$ is the revision of T by A, and finally, $T \dot{+} A$ is $Cn(T \cup \{A\})$, i.e. the result of a brute addition of A to T (followed by logical closure) without considering the need for consistency. Slightly departing from the AGM usage, we always assume that both T and A are *individually* consistent, only their union might not be.

AGM have proposed the following widely accepted axioms for the revision operator $*$. We omit their axioms 7,8 having to do with conjunction. See [Parikh 1996] for a discussion of these two axioms and why they are less plausible than the other six.

1. $T * A$ is a theory.
2. $A \in T * A$
3. If $A \Leftrightarrow B$, then $T * A = T * B$.
4. $T * A \subseteq T \dot{+} A$
5. If A is consistent with T, i.e. it is not the case that $\neg A \in T$, then $T * A = T \dot{+} A$.
6. $T * A$ is consistent if A is.

However, the AGM axioms are consistent with the *trivial update*, which is defined by:

If A is consistent with T, then $T * A = T \dot{+} A$, otherwise $T * A = Cn(A)$.

Thus in case A is inconsistent with T, under this update, all information in T is simply discarded, clearly something which is not in the spirit of the minimal change requirement. The work in [Parikh, 1996], discussed below, was motivated in part by a desire to block the trivial update.

We now give some details of the **LS** model on which our new notion of a B-structure is based.

2 The LS Model

[Parikh, 1996] introduced a language splitting model which addressed *two* of the four issues we mentioned at the beginning, namely minimal change and computational tractability. We review his results below before going on to describe our new B-structures model which addresses inconsistency tolerance as well.

The intuition behind the **LS** model is that our beliefs are subdivided into *disjoint* areas which do not affect each other. This separation allows us to revise our beliefs *locally* and to minimize the amount of computation we have to do when we receive a new piece of information, both in checking whether it is consistent with the old set of beliefs and also in revising our beliefs in view of the new information.

Definition LS1: 1) Suppose T is a theory in the language L and let $L_1, L_2, .., L_n$ be a partition of L. Then $L_1, ..., L_n$ *split* the theory T if there exist formulae $A_i \in L_i$ such that $T = Cn(A_1, ..., A_n)$. We may also say that $\{L_1, ..., L_n\}$ is a T-splitting.

2) If $L_1 \subset L$ then we say that T is *confined* to L_1 if $T = Cn(T \cap L_1)$.

In part 1 of the definition, we can think of T as being generated by the various T_i in languages L_i. Then the condition implies that T contains no "cross-talk" between L_i and L_j for distinct i, j. Part 2 of the definition says that T knows nothing about the part $L - L_1$ of L.

Lemma LS1: [Parikh, 1996] Given a theory T in the language L, there is a unique *finest* T-splitting of L, i.e. one which refines *every other* T-splitting.

Lemma 1 says that there is a unique way to think of T as being composed of disjoint information about certain subject matters.

Lemma LS2: Given a formula A, there is a *smallest* language L' in which A can be expressed, i.e., there is $L' \subseteq L$ and a formula $B \in L'$ with $A \Leftrightarrow B$, and for all L'' and B'' such that $B'' \in L''$ and $A \Leftrightarrow B''$, $L' \subseteq L''$.

Although A is equivalent to many different formulas in different languages, lemma 2 tells us that nonetheless, the question, "What is A actually *about*?" can be uniquely answered by providing a smallest language in which (a formula equivalent to) A can be stated. For example if A is the formula $P \wedge (Q \vee \neg Q)$, then $A \Leftrightarrow P$ and hence $L(A) = \{P\}$. The language L' will be referred to as $L(A)$. If $A \Leftrightarrow B$, then $L(A) = L(B)$.

The general rationale for **axiom P**, below, is that if we have information about two or more subject matters which, as far as we know, are unrelated (are split) then when we receive information about *one* of them, we should only update our information in that subject and leave the rest of our beliefs unchanged.

Axiom P: If $T = Cn(A, B)$ where A, B are in L_1, L_2 respectively and C is in L_1, then $T * C = Cn(A) *' C \dot{+} B$, where $*'$ is a (local) update operator for the sub-language L_1.

It follows from axiom P that if T is split between L_1 and L_2, A, B are in L_1 and L_2 respectively, then $T * A * B = T * B * A$.

Axiom P is thought of as *additional* to the AGM axioms. We now get

Theorem: [Parikh, 1996] There is an update procedure which satisfies the six AGM axioms and axiom P. Moreover, the trivial update procedure does not satisfy axiom P.

2.1 Comments on the LS Model

The **LS** model insists that the division of our beliefs uses *disjoint* sub-languages. But then global inconsistencies cannot be represented. Suppose an agent believes theories T_i in languages L_i and the L_i are mutually disjoint. Then if the T_i are individually consistent, they are also jointly consistent. Thus the **LS** model cannot explain how an agent can be locally logically omniscient – i.e. derive logical consequences *within* each L_i but still fail to be globally consistent.

We would like to model a more realistic agent, i.e. one who is locally logically omniscient, who does not believe any outright inconsistencies, but whose global belief structure might well be inconsistent without his being aware of it. This means that the languages L_i must be allowed to overlap on occasion. The B-structure model which we introduce below does just that and shows how the logic is handled using Belnap's four truth values, $\{\perp, true, false, \top\}$ [Belnap, 1977].

We feel that this relaxation is psychologically realistic. E.g. we may imagine that Clinton's problems with Lewinsky might be connected with the bombing of Iraq, which may affect the price of oil and which in turn may affect our airfare to a conference. But it is unrealistic to think that we might reason with all these beliefs together at one time. Rather they will be separated, though only partially; the connections may be noticed on occasion, if there is an article in the newspaper connecting a price hike in airfares with the bombing of Iraq.

Note that the sort of model we are considering will also be appropriate for group reasoning. Each agent will have its own language and be individually consistent in it. But the languages of different agents may overlap and a particular question A may be resolved by consulting those agents whose languages overlap with the language of A. Options A and B discussed below will show how this is handled at a technical level. Even a single agent may be thought of as consisting of a collective entity, see [Minsky, 1986].

3 The B-structure model

Definition 1: a belief structure $\mathcal{B}$ on L is a set $\{(L_1, T_1), ..., (L_n, T_n)\}$ such that
$L = \bigcup L_i : i \leq n$, and each T_i is a theory in L_i. (We can imagine that $T_i = Cn(\Gamma_i)$ where the Γ_i are the explicit beliefs of some agent in language L_i. If n =1, then $\mathcal{B}$ will be just a theory.)

As in the case of belief bases, the object of revision is a set of beliefs that is not necessarily logically closed.

While we are not requiring the L_i to be disjoint, we do want to retain *some* amount of disjointness. The following notion does that.

Definition 2: Given a finite propositional language L and a set of languages $\{L_1 \ldots L_n\}$ such that $L = L_1 \cup \ldots \cup L_n$, $\{L_1 \ldots L_n\}$ is a *k-partition* of L if any propositional symbol q occurs in at most k of the languages L_i.

The presence of such an overlap of symbols may be thought of as "cross-talk" amongst the T_i; beliefs in some of the T_i may be relevant or related to beliefs in others. A k-partition with a smaller k will be *more* disjoint than one with a larger k and indicates a finer partitioning of her beliefs by the agent. In particular, a normal partition is just a 1-partition. A k-partition is *a fortiori* a $k+1$-partition and hence, for instance, a maximally fine 2-partition will usually be finer than the finest 1-partition whose existence is guaranteed by lemma **LS1**. Tolerating some overlap makes it easier to organize one's beliefs in smaller chunks. For example, if T is axiomatized by the formula $(P \wedge Q) \vee (\neg Q \wedge R)$ then $L(T) = \{P, Q, R\}$ has only a trivial 1-partition. However, there is a 2-partition into sets $\{P, Q\}$ and $\{Q, R\}$ with the two theories being generated by $P \wedge Q$ and $\neg Q \wedge R$ respectively.

Definition 3: $\mathcal{B}$ as in definition 1 is m-consistent if any m of the T_i are jointly consistent.

Again, we do not require that the whole collection be consistent, but the *amount* of inconsistency is limited by this requirement. If $\mathcal{B}$ is 1-consistent, that merely says that all the T_i are individually consistent. At the other extreme, if $\mathcal{B} = \{(L_1, T_1), ..., (L_n, T_n)\}$ is n-consistent, then it will be consistent, period.

A very natural example of such a collection arises in the context of the *bald person paradox.* Suppose we say that a person with 10,000 hairs is not bald but that a person with 0 hairs *is* bald. We also say that a non-bald person cannot become bald through the removal of just one hair. This gives us the axioms $B(0)$, $B(n) \rightarrow B(n+1)$, for $n = 0$ to $n = 9,999$, and finally, $\neg B(10,000)$. These 10,002 axioms are inconsistent, but any 10,001 of them are consistent. This gives us a 10,001-consistent, inconsistent collection of axioms, all of which are accepted in practice. For a more practical example, consider an airline which has 100 seats on a flight, and accepts 120 reservations. If customer number i is C_i, then the airline is asserting, for $1 \leq i \leq 120$, that C_i has a seat and also that there are only 100 seats. This is an inconsistent set of assertions, whose inconsistency will not matter if 100 or fewer of the customers remind them of the respective assertions.

Our definitions clearly indicate a departure from conventional representations of belief sets as consistent, logically closed sets. The advantages of such a departure should be clear in the following discussion.

To represent how an agent *uses* his B-structure and revises it, we make use of the four valued semantics [Belnap, 1977], [Fitting,1989], where the four truth values are $\perp$, *true, false* $\top$, which stand for, respectively, *no information, true, false*, and *over-defined* (or inconsistent). The first three truth values are intuitively obvious. The last one, *over-defined* can arise, for instance when we are given contradictory information by two people both of whom are trustworthy. In this four valued lattice there are two orderings, the *knowledge* ordering under which $\perp <_k$ *true, false* $<_k \top$ and the *truth* ordering under which *false* $<_t \perp, \top <_k$ *true. true* is truer than $\top$ which is truer than *false.* But $\perp, \top$ are incomparable under the truth ordering. Similarly, *true,* *false* are incomparable under the knowledge ordering. The four truth values with their orderings are displayed in the double Hasse diagram below.

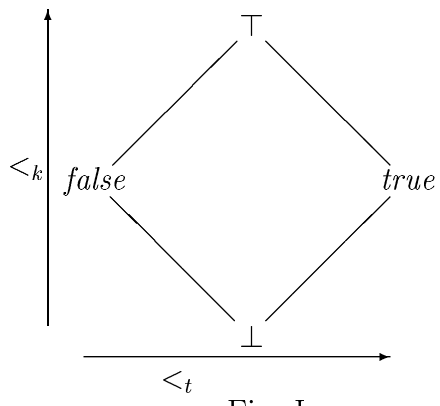

Fig. I

Now we consider how a B-structure responds to queries and how it is revised.

4 Answering queries

Given a formula A and B-structure $\mathcal{B}$ the query "A?" is answered by giving the value $v_{\mathcal{B}}(A)$ – abbreviated as $v(A)$ – which is defined as follows: Let $\Gamma_A = \bigcup T_i : L_i \cap L(A) \neq \emptyset$, where $L(A)$ is the *smallest language* in which A can be expressed (cf. Lemma LS2).
If Γ_A is consistent, then
  if $\Gamma_A \vdash A$, then $v(A) = true$
  if $\Gamma_A \vdash \neg A$, then $v(A) = false$, and
  $v(A) = \bot$ otherwise.
If Γ_A is inconsistent, then $v(A) = \top$.

Intuitively we see *which* of the theories T_i *could be* relevant to A and put them together to get Γ_A. Γ_A is used to answer questions about A and the rest of the theories T_j are *not* brought into play.

Controlling Inconsistency: Our approach, while seeking to handle this problem, is distinct from the usual paraconsistent approach in that we do allow full use of classical logic, albeit locally.

Limited consistency: Suppose that A has at most l distinct symbols, the

L_i are a k-partition, and $m \geq k \times l$. Then Γ_A will be a union of at most m of the theories T_i. If the collection $\{T_i : i \leq n\}$ is m-consistent, then Γ_A will be consistent and *exactly one* of the first three values will be given. In other words, an agent whose B-structure is *fairly consistent*, and who is responding to a *short* query, will *always* give a consistent response.

We say that the person with belief base $\mathcal{B}$ *implicitly believes* A if $v(A) =$ *true*. If so, it is true that A follows from her explicit beliefs but the converse does not hold in general. If the explicit beliefs are jointly inconsistent, then their consequences will be everything, but the agent may still have implicit beliefs which are a reasonable set. The implicit beliefs will always include the explicit ones and will be *locally* closed under logical rules.

What about adjunction? If an agent implicitly believes A and B, does he also implicitly believe $A \wedge B$? If the answers to A and B are both *true* and $L(A \wedge B)$ has no more than l symbols, then the answer to $A \wedge B$ will also be *true*. However, it is easy to find cases where the longer formula $A \wedge B$ forces the agent to simultaneously consider several of his beliefs and may result in the detection of the underlying inconsistency. Let $L_1 = \{P, R\}$, $L_2 = \{Q, R\}$, $T_1 = Cn(P, R)$ and $T_2 = Cn(Q, \neg R)$. Then the answers to the queries 'P?' and 'Q?' will both be *yes*, but to '$P \wedge Q$?' it will be $\top$.

5 Base Refinement

A person who has explicit beliefs Γ in language L may organize these beliefs in smaller or larger chunks. Clearly if Γ is inconsistent, then it is un-workable to organize them in a single chunk. Moreover, organizing in larger chunks can make the computational problems harder. On the other hand it does have the advantage that more implicit beliefs can be derived in the sense described below.

Definition 4: A B-structure $\mathcal{B}$ *refines* another, $\mathcal{B}'$ if (i) every language L_i is a subset of another L'_j, (ii) every $L'_j = \bigcup L_i | L_i \subseteq L'_j$, and (iii) every $T'_j = Cn(\bigcup T_i | L_i \subseteq L'_j)$.

Example: Suppose that the B-structure $\mathcal{B}$ has the three languages $\{P, Q\}, \{Q, R\}, \{R, S\}$ and theories $T_i : i \leq 3$ generated by the formulae $P \rightarrow Q$, $\neg Q \wedge R$ and $R \wedge S$ respectively. the B-structure $\mathcal{B}'$ has the two languages $\{P, Q, R\}, \{R, S\}$. and theories $T'_j : j \leq 2$ generated by the formulae

$\neg P \wedge \neg Q \wedge R$ and $R \wedge S$ respectively. Clearly $\mathcal{B}$ refines $\mathcal{B}'$. Now to the query 'P?', $\mathcal{B}$ will give the answer $\perp$ and $\mathcal{B}'$ will give the answer 'yes' or *true*.

Theorem 1: Let $\mathcal{B}$ refine $\mathcal{B}'$ and A be a formula. Then $v(A) \leq_k v'(A)$ where $v(A), v'(A)$ respectively stand for $v_{\mathcal{B}}(A)$ and $v_{\mathcal{B}'}(A)$.

Proof: Given a formula A and B-structure $\mathcal{B}$ the query "A?" is answered by giving the value $v_{\mathcal{B}}(A) = v(A)$ which is defined using $\Gamma_A = \bigcup T_i : L_i \cap L(A) \neq \emptyset$, where $L(A)$ is the language of A. Similarly for $\mathcal{B}'$.

Now we note that if $L_i \subseteq L'_j$ and $L(A)$ intersects L_i then it also intersects L'_j. Hence $\Gamma_A \subseteq \Gamma'_A$. This immediately yields $v(A) \leq_k v'(A)$. $\Box$

In other words, $\mathcal{B}'$ always yields more information than $\mathcal{B}$. The downside is that $\mathcal{B}'$ may give the inconsistent answer $\top$ to a query A whereas $\mathcal{B}$ may have given a *true* or *false* answer to it. Thus the answer from $\mathcal{B}'$ can be more 'knowledgable' but useless! The partitioning of information into separate languages may on occasion miss some answers which might have been obtained without such partitioning, but it is much more likely to be consistent in particular queries.

6 B-Structure Revision

What will happen to a B-structure when an agent receives some information which overlaps two sub-languages? One possibility is that the agent accepts the consequences *within* each sub-language while still keeping them separate. If I learn that both Beijing and London had cold winters, I am not likely to merge all my other beliefs about the two cities. However, if I repeatedly receive information which overlaps two sub-languages, I may decide that the division is artificial and should be abandoned. Options A and B below correspond to these two different attitudes.

option A - the non-merging option: Assume that each of the languages L_i has its own AGM style revision operator. Given a belief structure $\mathcal{B}$ and a new input A, assume without loss of generality that A has been written in the *smallest language* in which it can be expressed [Parikh, 1996]. Now define, for each i, the *i-shadow* of A to be the set $\{B | B \in L_i \wedge A \vdash B\}$. This will be a theory T'_i in L_i. T'_i is what A has to say about the language L_i. Let $A_i \in L_i$ be such that $T'_i = Cn(A_i)$. Now define the theories T''_i by:

$T''_i = T_i \dot{+} A_i$ if A_i is consistent with T_i, and

$T_i'' = T_i *_i A_i$ otherwise where $*_i$ is a local revision operator for L_i. Then the revised B-structure $\mathcal{B} * A$ will equal $\{(L_1, T_1''), ..., (L_n, T_n'')\}$.

In practice, if $L(A) \cap L_i$ is empty, we can just leave T_i unchanged, saving computational time. We define $\mathcal{B} \dot{+} A$ analogously, except that we use the operator $\dot{+}$ on the various T_i.

option B - the merging option: This option generalizes the procedure of [Parikh, 1996]. Given a formula A as input and a belief base $\mathcal{B}$ let as before $\Gamma_A = \bigcup T_i : L_i \cap L(A) \neq \emptyset$, where $L(A)$ is the language of A. Let T_A be $Cn(\Gamma_A)$ and $L_A = \bigcup L_i | L_i \cap L(A) \neq \emptyset$. Now replace all those languages L_i such that $L_i \cap L(A) \neq \emptyset$ by the single language $\bigcup L_i : L_i \cap L(A) \neq \emptyset$, which is their union. At the same time, replace all the corresponding T_i by the theory $T_A * A$. This new way will specialize to the procedure in [Parikh, 1996] where the languages were all assumed to be disjoint but the receipt of information resulted in joining those theories whose languages overlapped the language of the new information.

6.1 Properties of B-structure revision

Both options A and B are *computationally efficient*. Generally we assume that each L_i has relatively small size, say under some fixed p, while the cardinality N of the whole language $L = \bigcup L_i$ might be quite large. Then given a k-partition and a formula A with at most l distinct symbols, the query answering procedures run in time which is exponential in $l \times k \times p$, but *linear* in n. (For option A, the procedure is linear in n, k and exponential only in $l \times p$). Thus if $k \times l \times p$ – the number of atomic propositions *relevant* to A – is small compared to n, as is usually the case, the computational cost will be much smaller than that of usual update procedures which are exponential in N.

The update procedures need not preserve refinements. If $\mathcal{B}$ refines $\mathcal{B}'$ and there is new information A, it may reveal an inconsistency in $\mathcal{B}'$, even though $\mathcal{B}$ is unfazed. It may seem foolish not to notice one's own inconsistencies, but these are often unavoidable, as with the well known Preface and Lottery paradoxes [Kyburg, 1961], [Kyburg, 1997]. The ability to retain a measure of consistency and act on the basis of one's implicit beliefs may have much to say in its favour.

Option B results in two formally distinct subject areas merging as the result of some new information which straddles them. In real life this is likely to happen only occasionally. Suppose I have a B-structure which keeps my beliefs about Turkey, Iraq and Iran separate. If I now receive a great many pieces of information about the Kurds who are scattered over these three countries then I may simply create the new subject *Asia Minor* and give up my attempt to deal with the three countries separately.

6.2 Analogues of the AGM axioms

Let $\mathcal{B} * A$ denote the revision of B-structure $\mathcal{B}$ by the formula A according to option B. The axioms B1-B5 below hold. The set theoretic notions $\in, \subseteq$ used in stating the AGM axioms are now replaced by more sophisticated generalizations which enter in with Belnap's four truth values. For option A, axiom **B2** needs the caveat that $A \in L_i$ for some i or, at least that $A \Leftrightarrow B_1 \wedge ... \wedge B_p$ where each B_k is in some L_i and p is small. However *disjunctive* information which straddles two of the L_i may be lost if we insist, as we do in option A, on keeping the L_i separate.

Let $\mathcal{B}$ be a belief structure, A a new piece of information, $*$ the revision operator. We then have the following axioms: (read $A \in \mathcal{B}$ as 'A is an implicit belief according to $\mathcal{B}$'.)

B1. $\mathcal{B} * A$, the revision of $\mathcal{B}$ by A, is a belief structure.

B2. $A \in \mathcal{B} * A$

B3. If $A \Leftrightarrow B$, then $\mathcal{B} * A = \mathcal{B} * B$

B4. $\mathcal{B} * A \subseteq \mathcal{B} \dot{+} A$. i.e. if $\mathcal{B} * A$ and $\mathcal{B} \dot{+} A$ give values v, v' to A, then $v <_k v'$.

B5. If A is consistent with $\mathcal{B}$, i.e it is not the case that $v_{\mathcal{B}}(A) \in \{false, \top\}$, then $\mathcal{B} *A = \mathcal{B} \dot{+} A$

Other Issues: Issues commonly raised in belief revision literature include, for instance, the axiom of recovery, revision by conjunctions, by sets of propositions etc. Most of these properties *will* hold at the local level provided that the original local operations have them. At the non-local level when several of the T_i interact in a particular case of a belief revision, more complex patterns of behaviour will emerge. These require further investigation.

6.3 Review of Previous Work

We briefly discuss how other authors have addressed the four issues mentioned at the outset.

Work on the minimal change model has concentrated on minimizing the number of beliefs given up during the *contraction* operation. To this end, operations such as partial meet contraction using selection functions were defined in [Alchourron *et. al*, 1985]. A motivation for the choice of beliefs to be dropped is given by the notion of *epistemic entrenchment* introduced by [Gärdenfors and Makinson, 1988] and refined for iterated belief change by [Nayak, 1994], [Darwiche and Pearl, 1994], [Lehmann, 1995]. [Williams, 1996] uses the concept of *maxi-adjustment* to achieve maximal inertia of information under iterated belief revision. [Georgatos, 1999] presents a generalization of entrenchment that serves as a representation of the AGM axioms. The notion of *epistemic relevance* is used for minimal contraction in [Hansson, 1992] and [Nebel, 1992].

The distinction between implicit and explicit beliefs, has been explored by the proponents of the *belief base* method such as [Fuhrmann, 1991], [Nebel, 1992], [Hansson 1991;1992]. [Rott, 1992] combines some intuitions in showing how epistemic entrenchment orderings can be carried out for *safe contractions* for belief bases.

Belief revision for inconsistent belief bases has been studied in an alternative approach by [Brewka, 1991]. Also the possibility of paraconsistent belief revision is explored by [Tanaka, 1997] while [Restall and Slaney, 1995] have developed a paraconsistent semantical representation based on the revision of models approach suggested in [Grove, 1988]. The work of [Schotch and Jennings, 1980] predates the AGM approach to belief revision. They consider an approach based on giving up the adjunction rule: from $\Gamma \vdash A$ and $\Gamma \vdash B$ conclude $\Gamma \vdash A \wedge B$. As we shall see, our treatment retains this rule at the local level.

The investigation of complexity procedures in [Nebel, 1992], via a fine-grained set of complexity classes, has shown that the complexity of base revision procedures satisfying AGM postulates is that of ordinary propositional derivability. Nebel's comparison of different revision methods shows that model-based revision methods such as those of [Dalal, 1988] have a complexity which exceeds both **NP** and **co-NP**.

Conclusion: We started this paper by indicating four desiderata which a framework for answering queries and for belief revision should try to meet. The B-structures framework meets all four. In future work we intend to carry out a thorough study of this interesting new model for belief representation and revision and to implement the query answering and revision procedures.

Acknowledgements: This research was supported in part by a grant from the Research Foundation of CUNY. We thank Melvin Fitting, Konstantinos Georgatos, Henry Kyburg and Graham Priest for helpful comments.

References:

[Alchourron *et. al*, 1985] Alchourron, C., Gärdenfors, P. and Makinson, D. (1985) On the logic of theory Change: partial meet functions for contraction and revision. *Journal of Symbolic Logic*, Vol 580, pp 510-530.

[Belnap, 1977] Belnap, N.D. A useful four-valued logic. *Modern Uses of Multiple-Valued Logic*, J.M Dunn and G. Epstein eds., D. Reidel.

[Brewka, 1991] Brewka, G., Belief revision in a framework for default reasoning. *The Logic of Theory Change*, A. Fuhrmann and M Morreau, editors, LNAI 465, Springer.

[Dalal, 1988] Dalal, M. Investigations into a theory of knowledge base revision. *Proceedings of the Seventh National Conference of the American Association for Artificial Intelligence*, Saint-Paul, MN, pp 475-479.

[Darwiche and Pearl, 1994] Darwiche, A, Pearl J. On the logic of iterated belief revision. *Proceedings of Theoretical Aspects of Rationality and Knowledge*, 1994, pages 5-23.

[Fitting, 1989] Fitting, Melvin. Bilattices and the theory of truth. *Journal of Philosophical Logic*, Vol 18, 225-256, 1989.

[Fuhrmann, 1991] Fuhrmann, Andre. Theory contraction through base revision. *Journal of Philosophical Logic*, Vol 20, pp175-203, 1991.

[Gärdenfors and Makinson, 1988] Gärdenfors P., Makinson, D. Revisions of knowledge systems using epistemic entrenchment. *Proceedings of Theoretical Aspects of Reasoning about Knowledge*, Moshe Vardi ed., Morgan-Kaufmann, pp 83-96, 1988.

[Gärdenfors, 1988] Gärdenfors, Peter. *Knowledge in Flux: Modeling the*

Dynamics of Belief States, Bradford Books, MIT Press, Cambridge, MA, 1988.

[Georgatos, 1999] Georgatos, Konstantinos. To preference via entrenchment. *Annals of Pure and Applied Logic*, to appear.

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

[Hansson, 1992] Hansson, S.O. A dyadic representation of belief. *Belief Revision*, Gärdenfors, P. ed., Cambridge, 1992.

[Hansson, 1991] Hansson, S.O. In defense of base contraction. *Synthese*, 91:239-245, 1992.

[Kyburg, 1961] Kyburg, Henry, *Probability and the Logic of Rational Belief*, Middletown, CT: Wesleyan, 1961

[Kyburg, 1997] Kyburg, Henry. The rule of adjunction and reasonable inference. *Journal of Philosophy*, **XCIV**, 3, 1997, pages 109-125.

[Lehmann, 1995] Lehmann, D. Belief revision, revised. *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence 1995*, pages 1534-1540.

[Minsky, 1986] Minsky, M., *The Society of Mind*, Simon and Schuster.

[Nayak, 1994] Nayak, Abhaya. Foundational belief change. *Journal of Philsophical Logic*, Vol. 23, pp 495-533, 1994

[Nebel, 1992] Nebel, B. Syntax based approaches to belief revision. *Belief Revision*, Gärdenfors, P. ed., Cambridge, 1992.

[Parikh, 1996] Parikh, Rohit. Beliefs, belief revision and splitting languages. *Preliminary Proceedings of Information Theoretic Approaches to Logic, Language and Computation*, 1996, editors L. Moss, M. de Rijke and J. Ginzburg. Final version to appear, CSLI, 1999.

[Restall and Slaney, 1995] Restall, G and Slaney, J. Realistic belief revision. *Technical Report: TR- ARP-2-95*, Automated Reasoning Project, Australian National University, 1995.

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

[Schotch and Jennings, 1980] Schotch P.K, Jennings R. Inference and necessity. *Journal of Philosophical Logic*, 9 (1980), 327-340.

[Tanaka, 1997] Tanaka, Koji. What does paraconsistency do? The case of belief revision. *The Logica Yearbook*, Timothy Childers ed., Praha, 1997, pp 188-197.

[Williams, 1996] Williams, Mary-Anne. A practical approach to belief revision: Reason-based change, in L. Aiello and S. Shapiro eds. *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference*, Morgan-Kaufmann, San Mateo, CA, 412-421, 1996.