Local Change: A preliminary report

Sven Ove Hansson and Renata Wassermann

Department of Philosophy
Uppsala University
sven_ove.hansson@filosofi.uu.se

ILLC
University of Amsterdam
renata@wins.uva.nl

Abstract

An agent can usually hold a very large amount of beliefs. However, only a small part of these beliefs is used at a time. Efficient operations for belief change should affect the beliefs of the agent locally, that is, the changes should be performed only in the relevant part of the beliefs. In this paper we generalize the operations for belief change defined in [Han97a]. We obtain representation theorems for the operations based on a generic consequence operator C that does not need to be classic. We define a local consequence operator that only considers the relevant part of a belief base and show that this operator can be used to define local versions of the operations for belief change.

1 Introduction

An agent living in the world is confronted with new information all the time, be it through observation or through communication. The agent must be able to decide what to do with each piece of information that comes in. Depending on the reliability of the source of information and on previous beliefs held by the agent, new information can be accepted or rejected. In most models of belief revision, including the now classic AGM model [AGM85], the new information is always accepted. Recently, several models have been developed with the more realistic feature that inputs can be either accepted or rejected [Han97b].

In most belief revision models, belief states are assumed to be consistent. The set of beliefs held by an agent is assumed to be closed under logical consequence. It follows from this that an agent cannot have contradictory beliefs without believing in everything. Inconsistencies can never be local, but always propagate to the global level and corrupt the entire belief state.

However, this is not true for actual agents in the real world. Arguably, we all have inconsistent elements in our belief states, but these do not induce us to believe in everything. It is, to take just one example, quite feasible to believe both that Jesus was a human being and that Jesus was not a human being without believing that the moon is made of cheese. In other words, actual belief systems are capable of containing *local* inconsistencies, that do not corrupt the entire belief system. Since inconsistencies can be local, different inconsistent belief systems may exist, that differ in terms of the beliefs actually endorsed by the respective agents.

The purpose of this paper is to develop a theory of belief change in which local inconsistencies are accepted. We will adopt the general framework of sentential input-assimilating modelling of belief change [Han98] that dominates this area. Hence, belief states will be represented as sets of sentences in a language L closed under truth-functional operations. The belief state is assumed to change only in response to external inputs, and sentences in L will be used to represent inputs.

In the now classic AGM theory and most of its successors, belief states are represented by beliefs sets, i.e. logically closed sets of sentences [AGM85]. In formal terms, a subset K of L is a belief set if and only if $K = Cn(K)$, where Cn is the (supraclassical) consequence operator that represents the logic. A major alternative is to use a belief base, i.e. a set not closed under logical consequence, as the representation of the belief set [Fuh91, Han89, Neb92]. For every belief base B, its closure $Cn(B)$ is a belief set that represent the beliefs held by the doxastic agent. The elements of B are assumed to be in a sense more basic beliefs, from which the elements of $Cn(B) \setminus B$ are derived. Belief bases have substantial advantages in terms of computability [Neb96], and their increased expressive power as compared to belief sets can be used to represent important features of actual belief systems [Han92a]. On the other hand, belief sets have important advantages in terms of simplicity.

For our present purpose, the most important difference between belief sets and belief bases is that the latter allow for a simple and direct representation of the fact referred to above, namely that there can be different inconsistent belief states. Consider, for instance, the following two belief bases:

$B_1 = \{p, \neg p, q_1, q_2, q_3, q_4, q_5\}$

$B_2 = \{p, \neg p, \neg q_1, \neg q_2, \neg q_3, \neg q_4, \neg q_5\}$

These are distinct belief bases. We can suppose q_1 but not $\neg q_1$ to be endorsed according to B_1, whereas the opposite holds for B_2. However, on the belief set level these distinctions are lost since $Cn(B_1) = Cn(B_2) = L$. More generally speaking, belief bases but not belief sets allow us to distinguish between different inconsistency-containing belief states. Therefore, belief base representation will be used in what follows.

In AGM theory [AGM85], three forms of belief change are identified: contraction, expansion, and revision. Contraction consists of retracting a specified sentence from the belief set. Expansion consists of adding new information set-theoretically to the belief set. If the old and the new information are not logically compatible, then the new belief state after expansion will be inconsistent. Revision is consistency-preserving incorporation of new information, i.e. if the input sentence is consistent, then the new belief set will be consistent. If necessary, consistency is obtained by deleting parts of the original belief set. These operations have also been defined for belief bases [Fuh91, Han98].

Two additional operations of change on belief bases were introduced in [Han97a]: consolidation and semi-revision. Consolidation consists in making an inconsistent belief base consistent. Semi-revision is an operation that may, depending on the input sentence, either accept or reject new information.

One of these five operations, namely expansion, cannot easily be restricted to only a part of the belief state. To expand a belief base B by a sentence α means to perform the simple operation $B + \alpha = B \cup \{\alpha\}$. There is no way to restrict this addition of α to only a part of B, unless we introduce an extralogical division of B into compartments. As will be seen below, so much can be achieved with the division of B that can be derived from logic alone that we do not wish at this stage to add this further complication.

The other four operations can, however, readily be localized as follows:

1. *Local Contraction*: A belief is removed from a certain part of the belief base. I may, for instance, give up my believe that Fermat's theorem can be proved, without thinking about the consequences of this change for my beliefs about various non-mathematical matters.

2. *Local Consolidation*: This consists in removing inconsistencies from some part of the belief base. The rest of the agent's belief may well be inconsistent. For instance, I can make my beliefs about biological evolution consistent, while retaining global inconsistency between biological and religious beliefs.

3. *Local Revision*: This consists in adding a belief to the belief base in such a way that a certain part of the resulting base is consistent. If I

see, for example, that it is a sunny day in Amsterdam, this contradicts my belief that it is always raining in Holland and leads to revision. This can be done without checking whether my beliefs about Brazilian politics are consistent with the new belief.

4. *Local Semi-revision*: This consists of adding a belief to the belief base and making a part of the resulting base consistent, possibly rejecting the new belief. If I hear that it is snowing in Rio de Janeiro this contradicts my beliefs about the climate there. I may either accept or reject this new information. In both cases, my beliefs about Latin grammar (and most other subjects) will be unaffected by the operation.

Our first task is to identify the compartments in a belief base in terms of which local change can be defined. This will be done in Section 2. In Sections 3-6 we define local operations of belief change and present some representation results.

The following informal example motivates our definitions:

> **Example:** Suppose I am at home and I hear on the radio that my friend Carol has been murdered yesterday night and that there were no traces of doors or windows having been forced. I talked to her yesterday on the phone and she was home with her flatmates Ann and Bill. I know that noone else, except for Ann, Bill and Carol has the keys to their apartment. I conclude that Ann or Bill must have done it. But I have known Ann for quite some time and cannot believe that she would be able to murder anyone. I believe she did not do it. For similar reasons, I believe Bill did not do it. This is clearly inconsistent with my belief that one of them did it. So I decide to visit my friend Paul to ask what he thinks. In front of his place I see the lights are on. I know that if the lights are on, then Paul is home. I get out of the car and Paul's neighbour, inferring that I am coming to visit Paul tells me he is not home. This is all very confusing, but I am sure of one thing: I do not believe I am asleep!
>
> This illustrates the fact that inconsistencies are local, that is, the fact that I have inconsistent beliefs does not cause me to believe anything.

I have expanded my belief base with the information given by the neighbour and reached a local inconsistency. I am interested now in whether Paul is at home or not. For a moment, I forget the murder and think of the reasons I have to believe that Paul is at home and that he is not. In order to eliminate the local inconsistency, I have to give up at least one of the beliefs. Suppose I ring the bell and Paul answers the door. Then I reject the neighbour's information that Paul is not home. On the other hand, suppose that he does not answer the door. Since I see the lights on, I give up my previous belief that if the lights are on then he is at home. He must have forgotten the lights on when he left. In both cases, I eliminated the local inconsistency around the fact that Paul is at home, but note that I am still inconsistent about the crime! This is an example of local consolidation.

The whole operation, that is, adding the information given by the neighbour and then locally consolidating the beliefs illustrates the operation of local semi-revision.

In the rest of this paper we consider L to be a propositional language closed under the usual truth-functional connectives and containing a constant $\perp$ denoting falsum. We call *inference operation* any total function taking sets of formulas to sets of formulas, that is, any function from $\mathcal{P}(L)$ to $\mathcal{P}(L)$. We use C to denote inference operators and for any such operator, the relation $\vdash_C$ is defined as follows:

$A \vdash_C \alpha$ iff $\alpha \in C(A)$.

If C is a classical (tarskian) consequence operator, i.e. if it satisfies monotony, inclusion, and iteration, then it can also be denoted Cn.

Due to space limitation, all the proofs have been omitted. The interested reader can contact one of the authors for the full version of the paper.

2 Defining compartments and local implication

In this section we define the notions of compartments that will be used for defining local inference operations. Local inference operations are restricted to relevant compartments of a belief base.

2.1 Defining compartments

There are two major ways to introduce compartments of minds or databases. First, compartments may be introduced as an addition to the logic, so that one and the same belief base can be divided in different ways into compartments. Secondly, they may be derived from the logic. The second method is the more economical, requiring no extra entities, and should be tried out first. We are going to use it here.

The concept we will use is that of the compartment "around" a sentence or set of sentences. Intuitively, the compartment around a sentence α in B is the subset of B that is relevant to α, and the compartment around a set A in B is the subset of B that is relevant to at least one element of A. Since one and the same sentence β may be relevant to both α_1 and α_2, compartments thus defined will typically be overlapping.

Letting $c(A, B)$ denote the compartment for A in B, we will therefore assume that:

(1) $c(A, B) = \bigcup_{\alpha \in A} c(\alpha, B)$.

It follows from this that $c(\emptyset, B) = \emptyset$ for all B.

Due to (1), we can content ourselves with defining c for single sentences.

We are only interested in defining compartments around contingent expressions, since intuitively, no formula should be relevant for tautologies or contradictions. For these, we define: $c(\alpha, B) = \emptyset$ if $\alpha \in Cn(\emptyset)$ or $\neg\alpha \in Cn(\emptyset)$.

Since compartments are based on the underlying logic, they will depend on the inference operator. A useful tool in the definition of logical compartment is that of a *kernel set* [Han94]. Given a set A and an inference operator C, the set D of sentences is an *α-kernel* if and only if it is a minimal set implying α. The *kernel operator* $\perp\!\!\!\perp_C$, based on C, is the operator that, given a belief base B and a sentence α, selects all α-kernels that are subsets of B. Hence:

Definition 2.1 *Let C be an inference operation on L. Then the kernel operation $\perp\!\!\!\perp_C$ is the operation such that for all subsets B and elements α of L, $X \in B \perp\!\!\!\perp_C \alpha$ if and only if:*

1. $X \subseteq B$

2. $\alpha \in C(X)$

3. for all Y*, if* $Y \subset X$ *then* $\alpha \notin C(Y)$

The elements of $B \perp\!\!\!\perp_C \alpha$ *are called kernels.*

Observation 2.2 *Let* C *be an inference operation on the language* L *and* $\perp\!\!\!\perp_C$ *its associated kernel operation. If* C *satisfies compactness, then* $B \perp\!\!\!\perp_C \alpha \neq \emptyset$ *for all* $B \subseteq L$ *and all* $\alpha \in C(B)$.

We use $\perp\!\!\!\perp$ to denote the kernel operation associated with Cn, the classic consequence operator on L.

A first attempt to define the compartment for α in B is: (α contingent)

(2) $c_1(\alpha, B) = \bigcup(B \perp\!\!\!\perp \alpha)$

This definition is unsatisfatory since inconsistent kernels will be included, e.g. $c_1(\{p\}, \{p, q, \neg q\}) = \{p, q, \neg q\}$, since $\{q, \neg q\} \in B \perp\!\!\!\perp \alpha$. This problem can be solved by leaving out inconsistent kernels:

(3) $c_2(\alpha, B) = \bigcup((B \perp\!\!\!\perp \alpha) \setminus (B \perp\!\!\!\perp \bot))$

But this is insufficient, since negations are relevant. We would like to have, for example: $c(\{q\}, \{p, p \to \neg q, r, r \to s, s\}) = \{p, p \to \neg q\}$. This leads to a modification:

(4) $c_3(\alpha, B) = \bigcup((B \perp\!\!\!\perp \alpha) \cup (B \perp\!\!\!\perp \neg\alpha) \setminus (B \perp\!\!\!\perp \bot))$

Note that the compartments so defined may well be inconsistent.

We have now a definition of $c(\alpha, B)$ for arbitrary sentences α. Using (1) we extend this definition for logical compartments around sets of sentences and get our final definition:

Definition 2.3 *The* A*-compartment of* B*, where* A *and* B *are sets of sentences is defined as:*

$c(A, B) = \bigcup_{\alpha \in A}(\bigcup((B \perp\!\!\!\perp \alpha) \cup (B \perp\!\!\!\perp \neg\alpha) \setminus (B \perp\!\!\!\perp \bot))$

We call c *a compartmentalization function.*

Observation 2.4 *If* c *is the compartmentalization function defined above, then for all sets of sentence* A *and* B, $c(A, B) = c(A, c(A, B))$.

2.2 Defining local implication

Using the definition of logical compartments presented above, we can define relevant implication as the inference relation that considers only the relevant part of the belief set. In other words, we can use logical compartments as defined above to derive, given a (global) consequence operator, a local inference operator for each logical compartment. This can be done as follows:

Definition 2.5 *Let Cn be a consequence operation on L and let c be the compartmentalization function derived from Cn. Then for any set A, the A-restriction of Cn is the inference operation C_A such that for all sets B of sentences: $C_A(B) = Cn(c(A, B))$.*

A set B is A-restrictedly consistent if and only if $\perp \notin C_A(B)$

$B \vdash_A \alpha$ is an abbreviation of $\alpha \in C_A(B)$, and $B \vdash_\beta \alpha$ of $\alpha \in C_{\{\beta\}}(B)$.

> **Example revisited:** Let p stand for the proposition "Paul is at home", q for "The lights are on", a for "Ann is the murderer", b for "Bill is the murderer" and r for "I am asleep".
>
> My belief base B after talking to Paul's neighbour contains: $\{q, q \rightarrow p, \neg p, a \vee b, \neg a, \neg b, \neg r\}$. I am interested in whether Paul is at home, that is, the relevant beliefs are $c(p, B) = \{q, q \rightarrow p, \neg p\}$. Even though this set is inconsistent we have that $r \notin C_r(B) = Cn(c(r, B)) = Cn(\{\neg r\})$.

From the definition above it follows that:

Corollary 2.6 *1. $B \vdash_\alpha \beta$ iff $B \vdash_{\neg\alpha} \beta$*

2. if B is consistent, then:

- *$B \vdash_\alpha \alpha$ iff $B \vdash \alpha$*
- *$B \vdash_B \alpha$ iff $B \vdash \alpha$*

3. $B \vdash_\top \alpha$ iff $B \vdash_\perp \alpha$ iff $\vdash \alpha$

Observation 2.7 *Let C_A be the A-restriction of the consequence operation Cn. Then it holds for all sets B of sentences that $C_A(B) = Cn(C_A(B))$, that is, $C_A(B)$ is a theory.*

A number of authors [Gab85, Mak89, Lin91] have studied inference operations through the formal properties they have. We show now what kind of property the local inference operator C_A has.

Theorem 2.8 *For any set of sentences A, C_A, the A-restriction of the classical consequence operation Cn, satisfies the following properties:*

1. *Monotony - If $B \subseteq D$, then $C_A(B) \subseteq C_A(D)$;*
2. *Compactness - If $\alpha \in C_A(B)$ then there is some finite $D \subseteq B$ such that $\alpha \in C_A(D)$;*
3. *Idempotency (iteration) - $C_A(C_A(B)) = C_A(B)$;*
4. *Consistency preservation - If $\perp \notin Cn(B)$, then $\perp \notin C_A(B)$;*
5. *Cut - If $B \subseteq D \subseteq C_A(B)$, then $C_A(D) = C_A(B)$;*
6. *Path independence - $C_A(C_A(B) \cap C_A(D)) = C_A(Cn(B) \cap Cn(D))$; and*
7. *Sen - $C_A(B) \cup C_A(D) \subseteq C_A(B \cup D)$.*

We now present a list of properties which are not satisfied by C_A.

Theorem 2.9 *For some set of sentences A, C_A, the A-restriction of the consequence operation Cn, does not satisfy the following properties:*

1. *Inclusion: $B \subseteq C_A(B)$;*
2. *Supraclassicality: $Cn(B) \subseteq C_A(B)$;*
3. *Deduction theorem: If $q \in C_A(B \cup \{p\})$, then $p \to q \in C_A(B)$;*
4. *Reductio ad absurdum: If $\perp \in C_A(B \cup \{\neg p\})$, then $p \in C_A(B)$;*
5. *Falsity: For all p, $p \in C_A(\{\perp\})$;*
6. *Distributivity: $C_A(B) \cap C_A(D) \subseteq C_A(Cn(B) \cap Cn(D))$;*
7. *Chernov: $C_A(B \cup D) \subseteq Cn(C_A(B) \cup D)$;*
8. *Arrow: If $\perp \notin Cn(C_A(B) \cup D)$, then $C_A(B \cup D) = Cn(C_A(B) \cup D)$; and*

9. *Explosiveness: For all α and β, $\beta \in C_A(\{\alpha, \neg\alpha\})$.*

The classical consequence operator Cn satisfies:
(1) inclusion: $A \subseteq Cn(A)$
(2) idempotence (iteration): $Cn(A) = Cn(Cn(A))$
(3) monotony: if $A \subseteq B$, then $Cn(A) \subseteq Cn(B)$.
Nonmonotonic logics leave property (3) out, resource logics [Gab94] leave property (2) out. Our framework leaves (1) out [1].

The inclusion and supraclassicality properties are clearly not wanted, since the purpose of the operator C_A is to ignore irrelevant data.

That explosiveness does not hold agrees with the intuition that some inconsistencies are "less harmful" than others. If the inconsistency is completely irrelevant to the current reasoning, it should not trivialize the belief state causing the agent to believe every possible sentence. On the other hand, if the inconsistency turns out to be relevant, that is, if given a belief base B and a set of sentences A it is the case that $\perp \in Cn(c(A, B))$, then the relation $\vdash_A$ becomes explosive, giving an indication that something has to be repaired. This can also be interpreted as not letting inconsistencies the agent is not aware of play a rôle in reasoning.

3 The local operators

In this section we present the local operators that will be presented individually in more detail in further sections.

We will attempt to find local operators that are as close to AGM-style operators as possible, but only change parts (compartments) of the belief base.

- Local expansion is not possible in this framework, since new sentences are added to B as a whole. They cannot be added to some compartment(s) of B and not to others.

- Local contraction ($B \dot{-}_A \alpha$): The A-compartment of B is made not imply α. B as a whole may still imply α.

[1] Thanks due to Frans Voorbraak

- Local consolidation ($B!_A$): The A-compartment of B is made consistent. This may leave B as a whole inconsistent.

- Local revision: AGM-style revision by α consists of two suboperations, expansion by α and contraction by $\neg\alpha$. In the classic AGM operation of revision (first developed for belief sets), contraction by $\neg\alpha$ takes place first and is then followed by expansion by α. In [Han92b] this was called "internal revision" and an alternative procedure, "external revision" was proposed. It consists in first expanding the belief set by α and after that contracting by $\neg\alpha$.

  As we have seen, in the framework we have chosen, the expansive suboperation of revision cannot be made local. Therefore, we will assume that local revision consists of (non-local) expansion combined with local contraction.

  Local internal revision: The A-compartment of B is contracted by $\neg\alpha$ and the result is expanded by α. ($B \mp_A \alpha = (B \dot{-}_A \neg\alpha) + \alpha$)

  Local external revision: B receives α and the A-compartment of B is contracted by $\neg\alpha$. B as a whole need not be made consistent. ($B \pm_A \alpha = (B + \alpha) \dot{-}_A \neg\alpha$)

- Local semi-revision ($B?_A\alpha$): Just as for revision, the expansive part cannot be made local.

  B receives α and the A-compartment of B is consolidated. ($=(B+\alpha)!_A$)

When the A-compartment of B is empty, the contraction and consolidation operations above leave B unchanged, while internal/external revision and semi-revision all coincide with expansion.

4 Local contraction

We call *local contraction* an operation that, given two sets of sentences A and B and a sentence α returns a subset of B that does not A-restrictedly imply α[2].

[2]Harman [Har86] presents a principle for belief contraction that says that a belief should not be given up if it can be "easily reinferred". In our framework, we interpret "easily reinferred" as "inferred from the relevant beliefs".

4.1 Local kernel contraction

The idea behind kernel contraction is that, if we remove from the belief base at least one element of each α-kernel (minimal subset of the base that implies α), we obtain a belief base that does not imply α. In the case of the local operation, we are interested in obtaining a belief base that does not A-restrictedly imply α, where A is a set of sentences. We consider then only those kernel sets that A-restrictedly imply α and use an incision function to select the formulas from the kernel sets to be removed. An incision function is a function defined from sets of sets of sentences into sets of sentences, selecting at least one sentence from each set of the argument.

Definition 4.1 *An incision function is any function $\sigma : \mathcal{P}(\mathcal{P}(L)) \to \mathcal{P}(L)$ such that for any subset S of $\mathcal{P}(L)$:*

1. $\sigma(S) \subseteq \bigcup S$
2. *If $\emptyset \neq X \in S$, then $X \cap \sigma(S) \neq \emptyset$*

Definition 4.2 *Let C be an inference operation on L and σ an incision function. The local kernel contraction of B determined by C and σ is the operation $\dot{-}_{C,\sigma}$ such that for all sets of sentences B:*
$B \dot{-}_{C,\sigma} \alpha = B \setminus \sigma(B \perp\!\!\!\perp_C \alpha)$

The following theorem characterizes the operation of local kernel contraction. The intended interpretation is that C should be a local restriction C_A of a consequence operator Cn, as defined above. However, the formal result is more general and is therefore stated in terms of an operation C of a more general type.

Theorem 4.3 *Let C be an inference operation satisfying monotonicity and compactness. An operation $\dot{-}$ is an operation of local kernel contraction determined by C and some incision function if and only if for all sets of sentences B:*

- *If $\alpha \notin C(\emptyset)$, then $\alpha \notin C(B \dot{-} \alpha)$ (success)*
- *$B \dot{-} \alpha \subseteq B$ (inclusion)*

- *If $x \in B \setminus B \dot{-} \alpha$, then there is some $B' \subseteq B$ such that $\alpha \notin C(B')$ and $\alpha \in C(B' \cup \{x\})$ (core-retainment)*

- *If for all subsets B' of B $\alpha \in C(B')$ if and only if $\beta \in C(B')$, then $B \dot{-} \alpha = B \dot{-} \beta$ (uniformity)*

As a special case of the operation described above we can consider contracting with respect to the sentence being contracted, that is, taking $B \dot{-}_{\alpha} \alpha$.

If the belief base is consistent, this turns out to be equivalent to non-local contraction, as the following results show:

Lemma 4.4 *If B is consistent, then $B \perp\!\!\!\perp_{C_\alpha} \alpha = B \perp\!\!\!\perp \alpha$.*

Corollary 4.5 *If B is consistent and $B \dot{-} \alpha$ is an operation of kernel contraction defined as: $B \dot{-} \alpha = B \setminus \sigma(B \perp\!\!\!\perp \alpha)$, where σ is an incision function. then $B \dot{-}_{\alpha} \alpha = B \dot{-} \alpha$.*

This shows that when contracting a consistent belief base B by a sentence α we can disconsider all sentences in B that are not relevant for α, that is, all sentences that are not in $c(\alpha, B)$.

4.2 Local partial meet contraction

Another way of constructing a contraction operation is starting from, instead of the minimal subsets of the base implying the sentence to be contracted, the maximal subsets not implying the sentence. There may be several such sets. A selection function is applied to them, and the intersection of the selected sets is taken as the outcome of the operation.

A remainder operator $\perp_C$ selects for every set of sentences B and every sentence α the maximal subsets of B that do not imply α according to C. Formally:

Definition 4.6 *Let C be an inference operation on L. The remainder operation $\perp_C$ is the operation such that for all subsets B and elements α of L, $X \in B \perp_C \alpha$ if and only if:*

1. *$X \subseteq B$,*

2. *$\alpha \notin C(X)$, and*

3. $\alpha \in C(Y)$ *for all* Y *such that* $X \subset Y \subseteq B$.

Definition 4.7 *A selection function for* L *is a function* $g : \mathcal{P}(\mathcal{P}(L)) \rightarrow \mathcal{P}(\mathcal{P}(L))$ *such that for all subsets* S *of* $\mathcal{P}(L)$*,*

1. $g(S) \subseteq S$
2. *If* $S \neq \emptyset$ *then* $g(S) \neq \emptyset$

Definition 4.8 *The local partial meet contraction operator based on a consequence operator* C *and a selection function* g *is the operator* $\dot{-}_{C,g}$ *such that for all sets of sentences* B*, and sentences* α*:*

1. *If* $B \perp_C \alpha \neq \emptyset$*, then* $B \dot{-}_{C,g} \alpha = \bigcap g(B \perp_C \alpha)$.
2. *If* $B \perp_C \alpha = \emptyset$*, then* $B \dot{-}_{C,g} \alpha = B$.

Theorem 4.9 *Let* C *satisfy compactness and monotonicity. An operator* $\dot{-}$ *is an operator of local partial meet contraction based on* C *if and only if for all sets* B *of sentences and sentences* α*:*

- *If* $\alpha \notin C(\emptyset)$*, then* $\alpha \notin C(B \dot{-} \alpha)$ *(success)*
- $B \dot{-} \alpha \subseteq B$ *(inclusion)*
- *If* $\delta \in B \setminus (B \dot{-} \alpha)$*, then there is some* X *such that* $B \dot{-} \alpha \subseteq X \subseteq B$*,* $\alpha \notin C(X)$ *and* $\alpha \in C(X \cup \{\delta\})$ *(relevance)*
- *If for all subsets* B' *of* B $\alpha \in C(B')$ *if and only if* $\beta \in C(B')$*, then* $B \dot{-} \alpha = B \dot{-} \beta$ *(uniformity)*

Theorem 4.10 *1. If* $\dot{-}$ *is a local kernel contraction operator with respect to an operator* C*, then it is local partial meet contraction operator with respect to the operator* C.

2. It does not hold in general that if $\dot{-}$ *is a local partial meet contraction operator with respect to an operator* C*, then it is local kernel contraction operator with respect to an operator* C.

5 Local consolidation

We will call *local consolidation* an operation that, given two sets of formulas A and B returns a subset of B that is A-restrictedly consistent.

5.1 Local kernel consolidation

In local kernel consolidation an incision function is used to select at least one element of every minimal inconsistent subset of the relevant beliefs. The selected elements are removed from the base.

Definition 5.1 *Let C be an inference operation on the language L and σ an incision function. Then the local kernel consolidation operation determined by C and σ is the operation $!_{C,\sigma}$ such that for all subsets B of L:*
$B!_{C,\sigma} = B \setminus \sigma(B \perp\!\!\!\perp_C \perp)$

Theorem 5.2 *Let C be an inference operation satisfying compactness, monotonicity and $\perp \notin C(\emptyset)$. An operation $!$ is an operation of local kernel consolidation determined by C and some incision function if and only if for all sets B of sentences:*

- $\perp \notin C(B!)$ *(consistency)*
- $B! \subseteq B$ *(inclusion)*
- *If $\alpha \in B \setminus (B!)$, then there is some X such that $X \subseteq B$, $\perp \notin C(X)$ and $\perp \in C(X \cup \{\alpha\})$ (core-retainment)*

Note that the result of kernel consolidation is not always a maximal consistent subset of the given belief base.

5.2 Local partial meet consolidation

As in the case of local contraction, a local consolidation operation can be defined alternatively by considering the maximal consistent subsets of the base, instead of the minimal inconsistent subsets.

Definition 5.3 *The local partial meet consolidation operator based on an inference operator C and a selection function g is the operator $!_{C,g}$ such that for all subsets B of L, $B!_{C,g} = \bigcap g(B \perp_C \perp)$.*

Theorem 5.4 *Let C satisfy monotonicity, compactness, and $\perp \notin C(\emptyset)$. An operation ! is an operation of local partial meet consolidation based on C and some selection function if and only if for all sets B of sentences:*

- $\perp \notin C(B!)$ *(consistency)*
- $B! \subseteq B$ *(inclusion)*
- *If $\alpha \in B \backslash (B!)$, then there is some X such that $B! \subseteq X \subseteq B$, $\perp \notin C(X)$ and $\perp \in C(X \cup \{\alpha\})$ (relevance)*

6 Local semi-revision

The operation of semi-revising a belief base B by a sentence α consists of revising the base in a way that does not assign the highest priority to the incoming information, that is, α may be rejected. As is the case with AGM revision, that can be defined in terms of contraction and expansion, semi-revision can be defined in terms of consolidation and expansion via the identity [Han97a]:

$B?\alpha = (B + \alpha)!$

Semi-revision consists then of two phases: first the belief α is added to the base, and then the resulting base is consolidated.

Using local kernel consolidation we obtain a local kernel semi-revision operation:

$B?_{C,\sigma}\alpha = B \cup \{\alpha\} \setminus \sigma((B \cup \{\alpha\}) \perp\!\!\!\perp_C \perp)$

From local partial meet consolidation we define a local partial meet semi-revision operation:

$B?_{C,g}\alpha = \bigcap g((B \cup \{\alpha\}) \perp_C \perp)$.

Example re-revisited: Before talking to Paul's neighbour, my belief base B contained: $\{q, q \to p, a \vee b, \neg a, \neg b, \neg r\}$. The neighbour says that Paul is not home, so I locally semi-revise my belief base by $\neg p$. This means first adding $\neg p$ set-theoretically to B and then locally consolidating. Let B' be B after the expansion, that is, $B' = B \cup \{\neg p\}$. We have $c(p, B') = \{q, q \to p, \neg p\}$ and $c(p, B') \perp_{Cn} \perp = \{\{q \to p, q\}, \{q \to p, \neg p\}, \{q, \neg p\}\}$. Local partial meet consolidation gives us then, with the choice function $g(c(p, B') \perp_{Cn} \perp) = \{\{q \to p, q\}\}$, $B'' = \{q, q \to p, a \vee b, \neg a, \neg b, \neg r\} = B$, that is, the new information is rejected (with another choice function the information can be accepted).

7 Related Work

A logic is said to be paraconsistent if its inference relation is not explosive. Technically, when a logic has more than one inference relation (as is the case here, since for every set A there is an associated inference relation $\vdash_A$), all of them have to be non-explosive for the logic to be called paraconsistent. If we take A to be the whole language, $\vdash_A$ is explosive, so the logic of local reasoning cannot be called paraconsistent in this sense.

One of the main characteristics of Relevance Logic [Dun86] is that from $\{\alpha, \neg\alpha \vee \beta\}$ one cannot derive β. Depending on which set we use to limit the compartments, we can obtain the same property. We have, for example, $\{p, \neg p \vee q\} \not\vdash_p q$.

There are in the literature several attempts to define the concepts of relevance and dependence. [Lug96] classifies approaches to relevant implication in two groups: those which impose conditions on the contents of implied formulas and those which impose conditions on the deduction of the implied formulas. In the first group she mentions Parry, whose implications are valid only if no new variables are introduced in the consequent, Anderson and Belnap, who impose that the antecedent and the consequent share variables and Epstein, who associates themes to the formulas and imposes conditions on the themes of the antecedent and the consequent. In the second group she mentions approaches requiring that the antecedent is used in an "essential way" on the derivation of the consequent. Among these she mentions Myhill and again Anderson and Belnap.

The way our relevant compartments are defined, through kernel sets, emphasizes the second approach, that is, we consider relevant to a formula α the formulas that appear in a minimal derivation of α or its negation.

In [dCH96] a set of nine postulates is presented that any dependence relation should satisfy. They show then how a contraction operation satisfying the Gärdenfors postulates can be obtained from a dependence relation.

Given a dependence relation $\rightsquigarrow$, where $\alpha \rightsquigarrow \beta$ should be read as "β depends on α", Fariñas del Cerro and Herzig define a contraction operation by (their definition is for theories):

(Def $\dot{-}$) $\gamma \in K\dot{-}\alpha$ if and only if either $\vdash \gamma$ or $\gamma \rightsquigarrow \gamma$ and $\alpha \not\rightsquigarrow \gamma$

They show that, given a dependence relation $\rightsquigarrow$ satisfying their set of postulates, the contraction operation defined by (Def $\dot{-}$) satisfies the eight Gärdenforsian postulates for contraction. But they do not show where the dependence relation comes from.

Given our definition of compartments and a belief base B, we can define a dependence relation by:

(*) $\alpha \rightsquigarrow \beta$ if and only if $\alpha \in c(\beta, B)$

Our dependence relation satisfies only five of the nine dependence postulates, but this suffices for defining a contraction operation satisfying several of the contraction postulates. The problem is that such a contraction operation removes more than what is needed, since the fact that β depends on α does not mean that there is no way of obtaining β without using α.

In [RS95], an approach is taken where the logic under which belief sets are closed is paraconsistent. The motivation behind it is very close to ours, to have a way in which to work with inconsistent beliefs without trivializing the belief sets.

They use first degree entailment, a modification of (propositional) classical logic where each formula can have as its truth value any subset of {`true`, `false`}. A formula α is entailed by a set of formulas X if and only if any valuation that that makes every element of X contain at least `true`.

8 Conclusions

We have shown that the operations for belief change defined in [Han97a] can be generalized to other notions of inference that are not classical.

We have defined an inference operation that considers only the relevant parts of a belief base and shown that this inference operation can be used to define local versions of the operations of consolidation, contraction and semi-revision.

These local operations give the desired result when applied to examples of the type that motivated our work and that cannot be treated satisfactorily in the AGM model.

Acknowledgments: We would like to thank André Fuhrmann, Hans Rott, Frans Voorbraak and two anonymous referees for providing valuable comments. The second author is supported by a grant from the Brazilian funding agency CAPES. This paper was started during a visit of the second author to Sweden, sponsored by the Dutch Research Council (NWO) through the PIONIER-project Reasoning with Uncertainty.

References

[AGM85] Carlos Alchourron, Peter Gärdenfors, and David Makinson. On the logic of theory change. *Journal of Symbolic Logic*, 50, 1985.

[dCH96] Luis Fariñas del Cerro and Andreas Herzig. Belief change and dependence. In Yohav Shoham, editor, *Proceedings of TARK 96*, 1996.

[Dun86] Michael Dunn. *Relevance Logic and Entailment*, chapter III.3. Handbook of Philosophical Logic. D. Reidel, 1986.

[Fuh91] André Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20:175–203, 1991.

[Gab85] Dov Gabbay. Theoretical foundations for nonmonotonic reasoning in expert systems. In K. Apt, editor, *Expert Systems, Logics and Models of Concurrent Systems*, pages 439–459, 1985.

[Gab94] Dov Gabbay. Labelled deductive systems. Technical Report MPI-I-94-223, Max-Planck-Institute für Informatik, 1994.

[Han89] Sven Ove Hansson. New operators for theory change. *Theoria*, 55:114–132, 1989.

[Han92a] Sven Ove Hansson. In defense of base contraction. *Synthese*, 91:239–245, 1992.

[Han92b] Sven Ove Hansson. Reversing the levi identity. *Journal of Philosophical Logic*, 22:637–639, 1992.

[Han94] Sven Ove Hansson. Kernel contraction. *Journal of Symbolic Logic*, 59:845–859, 1994.

[Han97a] Sven Ove Hansson. Semi-revision. *Journal of Applied Non-Classical Logic*, 7(1-2):151–175, 1997.

[Han97b] Sven Ove Hansson, editor. *Theoria - special issue on non-prioritized belief revision*, 1997.

[Han98] Sven Ove Hansson. *A Textbook of Belief Dynamics*. Kluwer Academic Press, 1998.

[Har86] Gilbert Harman. *Change in View - Principles of Reasoning*. MIT Press, 1986.

[Lin91] Sten Lindström. A semantic approach to nonmonotonic reasoning: Inference operations and choice. (Unpublished Manuscript), 1991.

[Lug96] Valérie Lugardon. *Sur le Fondaments de la Notion de Dépendance et de son Application a la Theorie de l'Action*. PhD thesis, Université Paul Sabatier de Toulouse, 1996.

[Mak89] David Makinson. General theory of cumulative inference. In M. Reinfrank, J. de Kleer, M. L. Ginsberg, and E. Sandewall, editors, *Proceedings of the Second International Workshop on Non-Monotonic Reasoning*, LNAI. Springer-Verlag, 1989.

[Neb92] Bernhard Nebel. Syntax based approaches to belief revision. In Peter Gärdenfors, editor, *Belief Revision*. Cambridge University Press, 1992.

[Neb96] Bernhard Nebel. How hard is it to revise a belief base? Technical Report Report 83, Albert-Ludwigs-Universität Freiburg - Institut für Informatik, 1996.

[RS95] Greg Restall and John Slaney. Realistic belief revision. Technical Report TR-ARP-2-95, Automatic Reasoning Project, Australian National University, 1995.