Belief Liberation (and Retraction)

Richard Booth
Department of Computer Science
University of Leipzig
04109 Leipzig, Germany
booth@informatik.uni-leipzig.de

Samir Chopra
Department of Computer Science
Brooklyn College of CUNY
Brooklyn, New York
schopra@sci.brooklyn.cuny.edu

Aditya Ghose
School of IT and Computer Science
University of Wollongong
Wollongong, Australia
aditya@uow.edu.au

Thomas Meyer
National ICT Australia
School of Computer Science and Engineering
University of New South Wales
tmeyer@cse.unsw.edu.au

Abstract

We provide a formal study of belief retraction operators that do not necessarily satisfy the (Inclusion) postulate. Our intuition is that a rational description of belief change must do justice to cases in which dropping a belief can lead to the inclusion, or 'liberation', of others in an agent's corpus. We provide a few possible weakenings of the (Inclusion) postulate and then provide two models of liberation via retraction operators, σ-liberation and linear liberation. We show that the class of σ-liberation operators is included in the class of linear ones and provide axiomatic characterisations for each class. We also show how any given retraction operator (including the liberation operators) can be 'converted' into either a withdrawal operator (i.e., satisfying (Inclusion)) or a revision operator via (a slight variant of) the Harper Identity and the Levi Identity respectively.

1 Introduction

Formal modelings of rational belief change are inevitably interested in plausible descriptions of the process of dropping beliefs. The *AGM framework*, named after its originators Alchourrón, Gärdenfors and Makinson [1, 7], characterises belief contraction via a set of postulates. One of these, (Inclusion), states that the belief set that is the result of contraction must be included in the belief set prior to contraction. Justifications for (Inclusion) are hard to find - it is usually just taken for granted. But, there are situations in which *the removal of a belief might lead to the inclusion of new ones.* Consider an agent that keeps track of information received and which has received both $\neg\phi$ and then ϕ over a period of time. When it draws inferences from this set of information, it prioritises more recent information and hence does not infer $\neg\phi$. But information that causes it to retract ϕ can be viewed as also leading to either an increase in the plausibility of $\neg\phi$ or even to a belief in $\neg\phi$ and other beliefs that were blocked by ϕ. A similar situation occurs in settings involving default reasoning [16]. If an agent was committed to a default rule that sanctioned belief in ϕ provided it was consistent to assume ψ, and also believed $\neg\psi$ to be true, it would be unable to apply the default rule and would consequently not believe ϕ. Retraction of the belief $\neg\psi$ makes the default rule applicable, thus sanctioning belief in ϕ.

We believe that the overriding messages from examples like these is that removing one belief might remove the grounds for withholding another. That is, when a 'blocking' belief is removed from an agent's belief corpus, so are the reasons or arguments against other beliefs which the agent had not previously entertained. Such a model is in the spirit of a *foundational* approach to belief change [4] and this is, we argue, as it should be, since an agent's corpus is most plausibly viewed as a set of beliefs along with the reasons for holding them. Thus, belief retractions can be 'liberating': beliefs which were blocked are 'set free'. In this paper we start from a set of basic postulates for *retraction* which *excludes* (Inclusion) and also the much-debated postulate of (Recovery). The broad class of operators so defined is designed to include the 'traditional' operators of AGM contraction and *withdrawal* [12], but our main focus is to study those retractions which can be viewed as liberation operators. We do not aim to jettison the Principle of Minimal Change in this study - the intuitions there are certainly worth retaining. Doing justice to that particular methodological principle while not ignoring other equally important ones[1], and rejecting (Inclusion) will be an objective of ours. A formal argument which supports our pre-theoretic intuitions above is that it is well-known that, when defining a *revision* operator $*$ from an AGM contraction operator $\div$ via the Levi Identity ([7], see also Section 4.2 of this paper), $\div$ isn't required to satisfy (Recovery) to ensure $*$ satisfies the AGM revision postulates. Less widely acknowledged is the fact that $\div$ doesn't have to satisfy (Inclusion) either. That is, if $\div$ is a retraction and $*$ is defined from $\div$ via Levi then $*$ is a partial meet revision.

We begin in Section 2 by formally defining retraction operators and by looking at a few possible weakenings of the (Inclusion) postulate. In Section 3 we provide two models of liberation via retraction operators, *σ-liberation* and *linear liberation.* Each of these utilises a finite *sequence* of sentences which guides the operation of belief removal. Though they differ in the *way* they utilise the sequence, we will show that the class of σ-liberation operators is included in the class of linear liberation operators and provide axiomatic characterisations for each class. We also axiomatise a number of subclasses of linear liberation. In Section 4 we show how a given retraction operator can be 'converted' into either a withdrawal operator (satisfying (Inclusion)) or a revision operator using (a slight variant of) the Harper Identity and the Levi Identity respectively. Section 5 is given over to a discussion of one of the postulates which characterises σ-liberation, namely (Strong Conservativity). We briefly conclude in Section 6 before finishing off with some ideas for further work in Section 7.

We assume a propositional language L generated by finitely many propositional variables. We use $\models$ to denote classical entailment and Cn to denote the classical logical consequence operator; $\top, \bot$ have their usual meanings. We assume that the object of change is a *consistent* belief set K i.e., a deductively closed set of sentences. We take K to be arbitrary and fixed throughout. As is usual we use $K + \phi$ to denote $Cn(K \cup \{\phi\})$. In this paper we do not consider the more difficult problem of iterated retraction.

2 Postulates for retraction

We first present the *basic AGM postulates*, which characterise *partial meet contraction* [1]. We use $K \mathbin{\overset{\circ}{\underset{\circ}{-}}} \phi$ to denote the result of removing the sentence ϕ from K.

(L1) $K \mathbin{\overset{\circ}{\underset{\circ}{-}}} \phi = Cn(K \mathbin{\overset{\circ}{\underset{\circ}{-}}} \phi)$ (Closure)

[1]Such an approach is explicit in the work of Rott and Pagnucco[17], and Meyer et al.[13] where the Principle of Minimal Change gives way to other methodological principles.

(L2) If $\not\models \phi$ then $\phi \notin K \ominus \phi$ (Success)

(L3) If $\phi \notin K$ then $K \ominus \phi = K$ (Vacuity)

(L4) If $\models \phi_1 \leftrightarrow \phi_2$ then $K \ominus \phi_1 = K \ominus \phi_2$ (Extensionality)

(L5) $K \ominus \phi \subseteq K$ (Inclusion)

(L6) $K \subseteq (K \ominus \phi) + \phi$ (Recovery)

(Recovery) has already been seen as problematic (see for example [8, 12]). Following [12], we will call any operator which satisfies **(L1)**–**(L5)** a *withdrawal* operator. We want now to go a step further and shed (Inclusion) from the list as well. However we keep the following basic condition, which follows from **(L1)**, **(L5)** and **(L6)**:

$$K \ominus \top = K \qquad \text{(Failure)[6]}$$

Definition 1 *Let K be a belief set and $\ominus$ be an operator for K. Then $\ominus$ is a* retraction operator *(for K) iff $\ominus$ satisfies* **(L1)**–**(L4)** *and (Failure).*

2.1 Weaker versions of (Inclusion)

While we reject (Inclusion), we intend to do justice to the Principle of Minimal Change. Thus it behooves us to look for weaker versions which disallow gratuitous addition of new beliefs. We now consider some potential weakenings. Later we will check whether our proposed liberation operators satisfy these weakenings. The first is the following:

(w1) If $\theta \in K \ominus \phi$ and $\theta \notin K$ then $\neg\theta \in K$

This rule stipulates when a sentence may be introduced into the new belief set during an operation of removal: a new sentence θ may be introduced *only if* its negation was present before the removal operation (and, since $K \ominus \phi$ is always consistent,[2] has necessarily been given up during the removal). This formalises the intuition that a new sentence is introduced only if there was previously something present in the belief set which had kept it out but which is now no longer there. Though **(w1)** looks reasonable at first glance, the following indicates it is *too* strong for our purposes.

Proposition 1 *In the presence of (Closure), the rule* **(w1)** *is equivalent to:*

(w1′) *If K is not complete then $K \ominus \phi \subseteq K$*

Since this result may seem somewhat surprising, let us briefly give its proof. To show **(w1)** implies **(w1′)** suppose K is not complete[3]. Then there exists some $\lambda \in L$ such that both $\lambda \notin K$ and $\neg\lambda \notin K$. If there existed $\theta \in (K \ominus \phi) \setminus K$ then $\theta \notin K$ would give us either $\theta \vee \lambda \notin K$ or $\theta \vee \neg\lambda \notin K$ (since K is deductively closed). Suppose the former holds. Then, since $\theta \vee \lambda \in K \ominus \phi$ (which follows from $\theta \in K \ominus \phi$ and (Closure)), we may apply **(w1)** to obtain $\neg(\theta \vee \lambda) \in K$ and so $\neg\lambda \in K$. In a similar way if we suppose $\theta \vee \neg\lambda \in K \ominus \phi$ then we obtain $\lambda \in K$. Either way we get a contradiction and so there can be no $\theta \in (K \ominus \phi) \setminus K$, i.e., $K \ominus \phi \subseteq K$ as required. To show **(w1′)** implies **(w1)** suppose $\theta \in (K \ominus \phi) \setminus K$. Then $K \ominus \phi \not\subseteq K$ and so, applying **(w1′)**, we must have that K is complete. Hence, from $\theta \notin K$ we get $\neg\theta \in K$ as required.

[2] This is already ensured by the (Success) postulate.

[3] A belief set K is complete iff for all $\lambda \in L$ either $\lambda \in K$ or $\neg\lambda \in K$.

The rule $(\mathbf{w1'})$ says that (Inclusion) holds whenever the prior belief set K is not complete. Since the prior belief set K typically will *not* be complete, $(\mathbf{w1'})$ isn't much of a weakening of (Inclusion) and we should not be too disappointed when a suggested operation of retraction does not satisfy it (or the equivalent $(\mathbf{w1})$)[4]. A relaxed version of $(\mathbf{w1})$ is:

$(\mathbf{w2})$ If $\theta \in K \ominus \phi$ and $\theta \notin K$ then there exists $\psi \in L$ such that $\psi \models \theta$, $\neg\psi \in K$ and $\psi \in K \ominus \phi$

That is, every $\theta \in (K \ominus \phi) \setminus K$ can be 'traced back' to, i.e., is a *logical consequence* of, a sentence whose negation was in K but which is now included in $K \ominus \theta$. Equivalently:

Proposition 2 *In the presence of (Closure), the rule* $(\mathbf{w2})$ *is equivalent to:*

$(\mathbf{w2'})$ *If* $(K \ominus \phi) \cup K$ *is consistent then* $K \ominus \phi \subseteq K$

$(\mathbf{w2'})$ says: the new belief set is either included in the old one, or the agent now believes the negation of a sentence it previously held to be true. That is, if the agent does not weaken its belief set, it has made a complete about-turn regarding some beliefs.

Our last weakening of (Inclusion) is a property often held to be characteristic of withdrawal operators. Recall that, when one removes a sentence θ from K using an operation $\ominus$ of withdrawal, one does so without thereby including its negation $\neg\theta$ in the new belief set. There is just one possible situation when $\neg\theta \in K \ominus \theta$, and that is if $\neg\theta \in K$ (in which case – assuming as we do that K is consistent – $\theta \notin K$ and so $K \ominus \theta = K$ by (Vacuity)). That is, the following rule is taken to hold:

$(\mathbf{w3})$ If $\neg\theta \notin K$ then $\neg\theta \notin K \ominus \theta$

3 Models of liberation

We now present two models of liberation operators; each will be presented in terms of finite sequences of sentences. The second model is more general than the first: the class of liberation operators it generates includes that generated by the first.

3.1 σ-liberation

In our first model, the central intuition is that both the agent's set of beliefs and the way it removes beliefs from its belief set are formed on the basis of the information that it has received over the course of its intellectual career. We assume the agent has at its disposal a given *belief sequence* σ which is just a finite sequence $(\alpha_1, \ldots, \alpha_n)$ of sentences, with α_n being the most recent information the agent has received[5]. What beliefs is the agent committed to on the basis of σ, i.e., what is the belief set K_σ *associated with* σ? An obvious answer would be to take the set $[\![\sigma]\!]$ of all the sentences appearing in σ and to then close under Cn. The problem with this answer, of course, is that we would like K_σ to be consistent, and it could well be that $[\![\sigma]\!]$ is *inconsistent*. Instead we use the priority of information encoded in σ to help us - initially - pick out consistent subsets of $[\![\sigma]\!]$. We define the increasing sequence of sets $\Gamma_i(\sigma)$ inductively by setting $\Gamma_0(\sigma) = \emptyset$ and then, for each $i = 0, 1, \ldots, n-1$,

$$\Gamma_{i+1}(\sigma) = \begin{cases} \Gamma_i(\sigma) \cup \{\alpha_{n-i}\} & \text{if } \Gamma_i(\sigma) \cup \{\alpha_{n-i}\} \not\models \bot \\ \Gamma_i(\sigma) & \text{otherwise} \end{cases}$$

[4]Note that this equivalence depends on the presence of (Closure), i.e., on the new belief set being deductively closed. It could well be a different story if we worked instead with logically open belief *bases* [10].

[5]The sentences can stand for anything, not just a record of observations. The main thing is that we have a linearly ordered/prioritised set of sentences. Such a treatment is reminiscent of [3]. See also [15].

That is, starting with α_n, we work our way backwards through the sequence, adding each sentence as we go, provided it is consistent with the sentences collected up to that point. Note $\Gamma_n(\sigma)$ forms a maximal consistent subset of $[\![\sigma]\!]$. (In particular if $[\![\sigma]\!]$ is consistent then $\Gamma_n(\sigma) = [\![\sigma]\!]$.) We then take $Cn(\Gamma_n(\sigma))$ to be the belief set associated with σ.

Definition 2 *Let K be a belief set and $\sigma = (\alpha_1, \dots, \alpha_n)$ a belief sequence. We say σ is a belief sequence* relative to K *iff* $K = Cn(\Gamma_n(\sigma))$.

Example 1 Suppose $\sigma = (\neg p \wedge \neg q, p, p \to q)$ where p and q are distinct propositional variables. Then $\Gamma_0(\sigma) = \emptyset$, $\Gamma_1(\sigma) = \{p \to q\}$, $\Gamma_2(\sigma) = \{p, p \to q\} = \Gamma_3(\sigma)$. Hence the belief set K associated with this σ is given by $K = Cn(\Gamma_3(\sigma)) = Cn(p \wedge q)$. Note how belief in the first/oldest sentence $\neg p \wedge \neg q$ in σ is suppressed in particular by the more recent sentence p.

Given a belief sequence σ relative to K, we want to use σ to define an operation $\ominus_\sigma$ for K such that $K \ominus_\sigma \phi$ represents the result of removing ϕ from K. If ϕ is a tautology we just set $K \ominus_\sigma \phi = K$. Otherwise we introduce sequences of sets $\Gamma_i(\sigma, \phi)$ inductively by setting $\Gamma_0(\sigma, \phi) = \emptyset$ and then, for each $i = 0, 1, \dots, n-1$,

$$\Gamma_{i+1}(\sigma, \phi) = \begin{cases} \Gamma_i(\sigma, \phi) \cup \{\alpha_{n-i}\} & \text{if } \Gamma_i(\sigma, \phi) \cup \{\alpha_{n-i}\} \not\vdash \phi \\ \Gamma_i(\sigma, \phi) & \text{otherwise} \end{cases}$$

That is, starting at the end with α_n, we work our way backwards through the sequence, adding each sentence as we go, provided adding it to the sentences collected up to that point does not lead to the inference of ϕ. Note that $\Gamma_i(\sigma) = \Gamma_i(\sigma, \bot)$ and that $\Gamma_n(\sigma, \phi)$ is set-inclusion maximal amongst the subsets of $[\![\sigma]\!]$ which do not imply ϕ. We then set

$$K \ominus_\sigma \phi = \begin{cases} Cn(\Gamma_n(\sigma, \phi)) & \text{if } \not\vdash \phi \\ K & \text{otherwise} \end{cases}$$

Definition 3 *Let K be a belief set and $\ominus$ be an operator for K. Then $\ominus$ is a σ-*liberation *operator (for K) iff $\ominus = \ominus_\sigma$ for some belief sequence σ relative to K.*

Example 2 Suppose $K = Cn(p \wedge q)$ and let σ from Example 1 be the belief sequence relative to K. Suppose we wish to remove p. We first compute $\Gamma_3(\sigma, p)$. We have $\Gamma_0(\sigma, p) = \emptyset$, $\Gamma_1(\sigma, p) = \{p \to q\}$ $= \Gamma_2(\sigma, p)$ and $\Gamma_3(\sigma, p) = \{\neg p \wedge \neg q, p \to q\}$. Hence $K \ominus_\sigma p = Cn(\Gamma_3(\sigma, p)) = Cn(\neg p \wedge \neg q)$. Note how, at the second stage, p is *nullified*, which leads to the reinstatement, or liberation, of $\neg p \wedge \neg q$.

As the above example clearly shows, σ-liberation operators do not necessarily satisfy (Inclusion). In fact this example shows that σ-liberation does not satisfy the weaker version **(w3)** since we have $\neg p \notin K$ but $\neg p \in K \ominus_\sigma p$. Hence σ-liberation can result in the addition of the negation of the sentence being removed. This example also shows that σ-liberation doesn't satisfy **(w1)** (or, therefore, **(w1′)**), since we have $(\neg p \wedge \neg q) \vee r \in (K \ominus_\sigma p) \setminus K$ but $\neg((\neg p \wedge \neg q) \vee r) \notin K$ for any propositional variable r distinct from p, q. What properties *are* satisfied by σ-liberation? Well, first of all, we can confirm that σ-liberation is indeed a retraction operator according to our basic definition:

Proposition 3 *Every σ-liberation operator satisfies the basic retraction postulates—***(L1)***–***(L4)** *and (Failure)— and so is a retraction operator.*

We can also show that σ-liberation operators satisfy the weak inclusion postulate **(w2)**, but for this we will wait until Section 3.3 where we provide an axiomatic characterisation of σ-liberation.

3.2 Linear liberation

We now use a different way of using a sequence of sentences to define a retraction operator. These sequences are different from the σ used before, and will be employed in a simpler fashion. Intuitively, the agent has in mind several different candidate belief sets. We assume that the agent can order these candidate belief sets linearly according to preference, with the agent's actual current belief set identified with the most preferred belief set in this ordering. Since we work in a *finite* propositional language, every belief set can be identified with a single sentence. Therefore, we represent the agent's epistemic state as a sequence $\rho = (\beta_1, \ldots, \beta_m)$ of sentences, where each β_i stands for the belief set $Cn(\beta_i)$. $Cn(\beta_1)$ is the most preferred belief set, $Cn(\beta_2)$ is the next most preferred belief set, and so on[6].

Definition 4 *Let K be a belief set and $\rho = (\beta_1, \ldots, \beta_m)$ a finite sequence of sentences. Then ρ is a* K-sequence *iff we have $K = Cn(\beta_1)$.*

Now to remove a sentence ϕ from K using a K-sequence ρ we just take our new belief set to be the one generated by the most preferred sentence – according to ρ – which does not imply ϕ. If no such sentence exists, equivalently, if $\bigvee_k \beta_k \models \phi$, then we just take our new belief set to be K if ϕ is a tautology, and $Cn(\emptyset)$ otherwise. More precisely, from a given K-sequence ρ we define the operator $\ominus_\rho$ for K by

$$K \ominus_\rho \phi = \begin{cases} Cn(\beta_i) \text{ where } i = \min\{k \mid \beta_k \not\models \phi\} & \text{if } \bigvee_k \beta_k \not\models \phi \\ K & \text{if } \models \phi \\ Cn(\emptyset) & \text{otherwise} \end{cases}$$

Definition 5 *Let K be a belief set and $\ominus$ be an operator for K. Then $\ominus$ is a* linear liberation *operator (for K) iff $\ominus = \ominus_\rho$ for some K-sequence ρ.*

K-sequences essentially correspond to the 'linear' variety of the type of general epistemic state considered by Alexander Bochman [2]. (Unlike us, Bochman also considers infinite languages.)

It turns out that linear liberation operators do not generally satisfy (Inclusion) either. For a very simple counterexample let $K = Cn(p)$ and consider the K-sequence $\rho = (p, \neg p)$. Then clearly we get $K \ominus_\rho p = Cn(\neg p)$, so $\neg p$ has entered the belief set. Thus linear liberation operators also fail to satisfy **(w3)**. We can characterise the class of linear liberation operators as follows (cf. Representation Theorem 5 in [2]):

Proposition 4 *Let K be a belief set and $\ominus$ an operator for K. Then $\ominus$ is a linear liberation operator iff $\ominus$ is a retraction operator that satisfies*

If $\theta \notin K \ominus (\theta \wedge \phi)$ then $K \ominus \theta = K \ominus (\theta \wedge \phi)$ *(Hyperregularity)*

The name (Hyperregularity) comes from [9]. When added to the basic retraction postulates this rule allows us to derive some extra properties:

Proposition 5 *Let $\ominus$ be a retraction operator which satisfies (Hyperregularity). Then $\ominus$ also satisfies the following two properties:*

- *Either $K \ominus (\theta \wedge \phi) = K \ominus \theta$ or $K \ominus (\theta \wedge \phi) = K \ominus \phi$*

[6] Since ρ is a *sequence* this means that the same sentence may appear more than once in ρ. However, for the results in this paper, it turns out that this feature can be ignored if desired.

- *If* $\theta \notin K \ominus \phi$ *and* $\phi \notin K \ominus \theta$ *then* $K \ominus \theta = K \ominus \phi$

The first property above is the postulate known as (Decomposition) from [1]. The second property gives a condition for when removing two different sentences yields the same result. We now look at some subclasses of the class of linear liberation operators.

3.3 Special cases of linear liberation

Note that, in the definition of a K-sequence, there need not be *any* relationship between the sentences β_i. Other, more restricted classes of liberation operators can now be found by placing restrictions on the β_i. We consider four here. First it is natural to ask: when does a linear liberation operator $\ominus_\rho$ satisfy (Inclusion)? It is quite easy to see that this will happen if and only if each sentence in ρ is a logical consequence of β_1, i.e.,

(A) For each $i = 1, \ldots, m$ we have $\beta_1 \models \beta_i$

Proposition 6 *Let* $\ominus$ *be a linear liberation operator for* K. *Then* $\ominus$ *satisfies (Inclusion) iff* $\ominus = \ominus_\rho$ *for some* K*-sequence* ρ *which satisfies condition (A).*

Next consider the following, stronger, condition on a K-sequence $\rho = (\beta_1, \ldots, \beta_m)$:

(B) For $i < j$ we have $\beta_i \models \beta_j$

(B)—which says that sentences get progressively logically weaker through ρ[7]— leads to an important class of withdrawal operators – the class of *severe withdrawal* operators [17] which, as is shown in [17], may be characterised by the basic retraction postulates plus (Inclusion) and the following two rules:

- If $\not\models \theta$ then $K \ominus \theta \subseteq K \ominus (\theta \wedge \phi)$
- If $\theta \notin K \ominus (\theta \wedge \phi)$ then $K \ominus (\theta \wedge \phi) \subseteq K \ominus \theta$ (Conjunctive Inclusion)

Note that the second rule above corresponds to "one half" of (Hyperregularity). It is in fact one of the two AGM *supplementary postulates* for contraction [1]. The first rule above is essentially a strengthened version of the other supplementary postulate "$(K \ominus \theta) \cap (K \ominus \phi) \subseteq K \ominus (\theta \wedge \phi)$".

Proposition 7 *Let* K *be a belief set and* $\ominus$ *an operator for* K. *Then* $\ominus = \ominus_\rho$ *for some* K*-sequence* ρ *which satisfies condition (B) iff* $\ominus$ *is a severe withdrawal operator.*

An example of a condition that doesn't lead to the satisfaction of (Inclusion) is the following:

(C) For $i \neq j$, $\beta_i \wedge \beta_j$ is inconsistent

Thus (C) says that the sentences in ρ represent mutually incompatible points of view.

Proposition 8 *Let* K *be a belief set and* $\ominus$ *an operator for* K. *Then* $\ominus = \ominus_\rho$ *for some* K*-sequence* ρ *which satisfies condition (C) iff* $\ominus$ *is a linear liberation operator that satisfies*

If $(K \ominus \theta) \cup (K \ominus \phi)$ *is consistent then* $K \ominus \theta = K \ominus \phi$ *(Dichotomy)*

[7]Such sequences are also studied in [5] from the perspective of qualitative utility in economics.

A justification for this postulate is provided by the condition on the sequence ρ. An agent has a sequence of mutually incompatible belief sets. Its way of dealing with changes will necessarily have to be dichotomous. When is such a mode of reasoning sensible? When the agent has become quite sophisticated through a process of refinement and ironing out differences in its belief corpus. The theories in ρ, then, are most plausibly viewed as the end products of a period of making small changes and converging on a cluster of (incompatible) alternatives. Therefore, we refer to this type of liberation as *dichotomous* liberation. (Dichotomy) can also be seen as describing belief change that lies between contraction and revision - a view confirmed by the discussion in Section 4.1.

Finally we have the following condition:

(D) For $i < j$ we have either $\beta_i \models \beta_j$ or $\beta_i \wedge \beta_j \models \bigvee_{k<i} \beta_k$

Each of (B) and (C) implies (D), a condition which leads us to the following subclass of linear liberation:

Proposition 9 *Let K be a belief set and $\ominus$ an operator for K. Then $\ominus = \ominus_\rho$ for some K-sequence ρ which satisfies condition (D) iff $\ominus$ is a linear liberation operator that satisfies*

If $(K \ominus \theta) \cup (K \ominus \phi) \not\models \phi$ *then* $K \ominus \theta \subseteq K \ominus \phi$ *(Strong Conservativity)*

The postulate (Strong Conservativity) has been studied in [9], where it is shown to be a characteristic postulate for *base-generated maxichoice contraction* operators. We shall discuss this postulate in Section 5. The significance of this particular subclass of linear liberation operators is that it is equivalent to none other than the class of σ-liberation operators from Section 3.1. The correspondence is proved in the following result:

Proposition 10 *Let K be a belief set. Then for each belief sequence σ relative to K there exists a K-sequence ρ satisfying (D) such that $\ominus_\sigma = \ominus_\rho$. Conversely for each K-sequence ρ satisfying (D) there exists a belief sequence σ relative to K such that $\ominus_\rho = \ominus_\sigma$.*

Hence we may state:

Corollary 1 *Let K be a belief set and let $\ominus$ be an operator for K. Then $\ominus$ is a σ-liberation operator iff $\ominus$ is a linear liberation operator that satisfies (Strong Conservativity).*

So σ-liberation may be axiomatically characterised by the basic retraction postulates plus (Hyperregularity) and (Strong Conservativity). Furthermore the results of this section allow us to say more. For instance we can now immediately see that every severe withdrawal operator is also a σ-liberation operator (as is every dichotomous liberation operator). Also, since it is known that severe withdrawal doesn't satisfy (Recovery), σ-liberation doesn't satisfy (Recovery) either. Finally, note that the weak inclusion postulate (**w2′**) is just a special instance of (Strong Conservativity) (since $K \ominus \bot = K$ for any retraction operator for K, a fact that follows from (Vacuity) and our assumption that K is consistent). Hence we can now see that every σ-liberation operator satisfies (**w2′**) (and the equivalent (**w2**)).

4 From retraction to withdrawal and revision

In this section we consider the relationship between retraction operators and the two more traditional belief change operators of withdrawal and revision. In particular we show how retraction operators can be 'converted' into either withdrawal or revision operators.

4.1 Retraction to withdrawal

What distinguishes retraction operators from withdrawal operators is that removing beliefs using the former may lead to the introduction of *new* beliefs into the belief set, while using the latter *always* leads to a new belief set which is a subset of the prior belief set. However, there is a simple way in which a given retraction operator may be *transformed* into a withdrawal operator. After retraction is performed, we simply discard all sentences which were not originally elements of K, i.e., from each retraction operator $\ominus$ for K we can define the new operator $\dot{-}$ for K by setting for each $\phi \in L$,

$$K \dot{-} \phi = K \cap (K \ominus \phi)$$

Obviously $\dot{-}$ is guaranteed to satisfy (Inclusion). This is strongly reminiscent of the *Harper Identity* [7]. A formal difference is the appearance of "ϕ" rather than "$\neg\phi$" on the right-hand side. A more crucial difference is that while the Harper Identity is usually employed as a means of obtaining a withdrawal operation from a given *revision* operator, here we use a slight variant of it to obtain a withdrawal operator from a *retraction* operator. Continuing with our liberation metaphor, we make the following definition:

Definition 6 *Let K be a belief set and let $\ominus$ be an operator for K. If the operator $\dot{-}$ for K is defined from $\ominus$ as above then we call $\dot{-}$ the* incarceration[8] *of $\ominus$.*

It is fairly easy to see that, as well as (Inclusion), the incarceration of a retraction operator also satisfies (Closure), (Success), (Vacuity) and (Extensionality), and thus that

Proposition 11 *The incarceration of a retraction operator is a withdrawal operator.*

What about our subclasses of liberation operators? What happens, for instance, when we take the incarceration of a linear liberation operator? Suppose $\ominus$ is a linear liberation operator. Then by definition we have $\ominus = \ominus_\rho$ for some K-sequence ρ. Now we can perform a modification to ρ to get a new sequence $f(\rho)$ as follows. Given $\rho = (\beta_1, \ldots, \beta_m)$ we just replace each β_i by $\beta_i \vee \beta_1$ (for $i > 1$), i.e., we define

$$f(\rho) = (\beta_1, (\beta_2 \vee \beta_1), (\beta_3 \vee \beta_1), \ldots, (\beta_n \vee \beta_1))$$

Clearly, since β_1 is unchanged, $f(\rho)$ is again a K-sequence. Furthermore we have:

Proposition 12 *Let ρ be a K-sequence and $\dot{-}$ be the incarceration of $\ominus_\rho$. Then $\dot{-} = \ominus_{f(\rho)}$.*

Thus the incarceration of a linear liberation operator is again a linear liberation operator which furthermore satisfies (Inclusion). Also, every linear liberation operator satisfying (Inclusion) arises as the incarceration of some linear liberation operator, namely itself. Note too, that the postulates for linear liberation together with (Inclusion) characterise the first special case of linear liberation (i.e., the sequences which satisfy (A)).

What happens when we take the incarceration of a σ-liberation operator? From Proposition 9 and Corollary 1 we know that $\ominus$ forms a σ-liberation operator iff $\ominus = \ominus_\rho$ for some K-sequence ρ which satisfies the condition (D). Thus we know from Proposition 12 that if $\dot{-}$ is the incarceration of a σ-liberation operator then $\dot{-} = \ominus_{f(\rho)}$ for some K-sequence ρ which satisfies (D). We can show the following:

Proposition 13 *Let ρ be a K-sequence. If ρ satisfies condition (D) then so too does $f(\rho)$.*

[8] We are grateful to David Makinson for suggesting this terminology.

Thus the condition (D) on K-sequences remains *invariant* under the modification f. (We remark that the same cannot be said of condition (C).) This result tells us then that every incarceration of a σ-liberation operator has the form $\Bumpeq_\rho$ for some ρ satisfying condition (D). Hence as a corollary we may state:

Corollary 2 *The incarceration of a σ-liberation operator is again a σ-liberation operator which furthermore satisfies (Inclusion). Also, every σ-liberation operator satisfying (Inclusion) arises as the incarceration of some σ-liberation operator, namely itself.*

4.2 Retraction to revision

To *revise* a belief set K by a given sentence ϕ means to modify K so that it includes ϕ, while preserving consistency. From each retraction operator $\Bumpeq$ for K we can define the revision operator $*$ for K via the Levi Identity:

$$K * \phi = (K \Bumpeq \neg\phi) + \phi$$

The Levi Identity is usually employed to define a revision operator from a given *withdrawal* operator. A central result in the AGM theory of belief change [1, 12] shows that if $\Bumpeq$ is a withdrawal operator then $*$ satisfies all the basic AGM postulates for revision[9]. The next result confirms that it is not necessary for $\Bumpeq$ to satisfy (Inclusion) for this result to go through.

Proposition 14 *Let $\Bumpeq$ be a retraction operator for K and let $*$ be defined from $\Bumpeq$ via the Levi Identity. Then $*$ satisfies the basic AGM revision postulates (relative to K). Furthermore, for* every *operator $*$ for K which satisfies the basic AGM revision postulates there exists a retraction operator $\Bumpeq$ for K such that $*$ may be obtained from $\Bumpeq$ via the Levi Identity.*

The first part of this proposition follows from the proof of the AGM result for withdrawal operators, and by noticing that in the only place in that proof where (Inclusion) is applied, namely in showing that the revision postulate "$K * \phi \subseteq K + \phi$" holds, it can be replaced with a use of (Vacuity) (in fact (Weak Vacuity 1) – see below). The second part follows from the well-known result in AGM theory that every operator $*$ satisfying the basic AGM revision postulates may be obtained via the Levi Identity from a partial meet contraction operator for K (i.e., satisfying **(L1)**–**(L6)**). Clearly every partial meet contraction operator is a retraction operator according to our definition. The above result shows us, then, that retraction operators are equally as suitable as withdrawal operators when it comes to using them as a stepping-stone for revision. For linear liberation operators we can say more:

Proposition 15 *Let $\Bumpeq$ and $*$ be as in the previous proposition. Then if $\Bumpeq$ additionally satisfies (Hyperregularity) then $*$ will satisfy both supplementary AGM revision postulates. Furthermore, for* every *operator $*$ for K which satisfies* all *the AGM revision postulates (basic plus supplementary) there exists a retraction operator $\Bumpeq$ for K satisfying (Hyperregularity) such that $*$ may be obtained from $\Bumpeq$ via the Levi Identity.*

The second part of this proposition is shown by observing that every severe withdrawal operator is a retraction operator satisfying (Hyperregularity). From results in [17, Sect. 7] we know that, given any operator $*$ for K satisfying the full list of AGM revision postulates, there is a severe withdrawal operator which, when the Levi Identity is applied to it, yields $*$.

[9][10] points out that $\Bumpeq$ is not required to satisfy (Closure) for this result. For the full list of (basic plus supplementary) AGM revision postulates we refer the reader to, e.g., [7, 10].

For a given retraction operator $\Bumpeq$ it is natural to wonder what happens if, instead of applying the Levi Identity directly to $\Bumpeq$, we first take its *incarceration* $\bumpeq$ and then apply the Levi Identity to $\bumpeq$. The next result shows that this has no effect on the resulting revision operator, i.e., that $\Bumpeq$ and $\bumpeq$ are *revision-equivalent* [12].

Proposition 16 *Let $\Bumpeq$ be a retraction operator for K and let $\bumpeq$ be the incarceration of $\Bumpeq$. Then, for all $\phi \in L$, $(K \Bumpeq \neg\phi) + \phi = (K \bumpeq \neg\phi) + \phi$.*

The above result may seem surprising. Since it is perfectly possible that $K \Bumpeq \phi \supset K \bumpeq \phi$, it might be expected that revision based on $\Bumpeq$ could sometimes lead to a strictly larger belief set than revision based on just $\bumpeq$. However, since every retraction operator satisfies (Weak Vacuity 1) (see below), we have $(K \Bumpeq \neg\phi) + \phi \subseteq (K + \phi) + \phi = K + \phi$, and so $(K \Bumpeq \neg\phi) + \phi = (K + \phi) \cap ((K \Bumpeq \neg\phi) + \phi)$ $= (K \cap (K \Bumpeq \neg\phi)) + \phi = (K \bumpeq \neg\phi) + \phi$ as claimed.

Overall, the results of this section have shown that it is possible to get a long way without (Inclusion). However we end this section by remarking that it is possible to get away with even less of the AGM contraction postulates, for propositions 11, 14, 15 and 16 do not even need the full power of (Vacuity); they can be derived using both of its following two weakenings, the first of which also doubles as another weakening of (Inclusion):

- $K \Bumpeq \phi \subseteq K + \neg\phi$ (Weak Vacuity 1)
- If $\phi \notin K$ then $K \subseteq K \Bumpeq \phi$ (Weak Vacuity 2)

5 The postulate (Strong Conservativity)

As we saw in Section 3.3, the key postulate which characterises the class of σ-liberation operators within the larger class of linear liberation operators is (Strong Conservativity). In Section 4.1 we saw that this property also holds for the class of withdrawal operators which are the incarcerations of the σ-liberation operators. In this section we want to take a closer look at this postulate. We give two arguments for its reasonableness: one based on its interplay with the other postulates, in particular (Recovery), and one showing how it can be given an interpretation which squares well with the Principle of Minimal Change.

5.1 (Strong Conservativity) and (Recovery)

First of all, the connections between (Strong Conservativity) and some of the other postulates seem to indicate that it's a reasonable requirement, at least for *withdrawal* operators:

Proposition 17 *Let $\Bumpeq$ be an operator for K which satisfies (Inclusion), (Closure), (Extensionality), (Success) and (Conjunctive Inclusion). Then the following are equivalent:*
(i). *$\Bumpeq$ satisfies both (Recovery)* **and** *(Strong Conservativity).*
(ii). *$\Bumpeq$ satisfies the following property:*

If $\theta \in K$ and $\theta \notin K \Bumpeq \phi$ then $\theta \rightarrow \phi \in K \Bumpeq \phi$ *(Fullness)*

(Fullness) is the characteristic postulate of the so-called *maxichoice* contraction operators [1] and as such represents 'pure' minimal change. The problem with maxichoice contraction is that it can often lead to unintuitive results (see [7]). This prompts the search for weakened versions of (Fullness) which embody less stringent forms of minimal change. The above direction *(ii)⇒(i)* shows that (in the presence of those other rules) both (Strong Conservativity) and (Recovery) fit the bill equally well here. However the direction *(i)⇒(ii)* shows that we can't have *both* properties without again getting (Fullness). The path chosen by AGM is to go with (Recovery) and leave out (at least implicitly) (Strong Conservativity). But if we agree with the many authors who have called (Recovery) into question and decide to relax it, (Strong Conservativity) ought then to come back into serious consideration. To put it another way, (Strong Conservativity) gets liberated by the removal of (Recovery)!

5.2 (Strong Conservativity) and the Principle of Minimal Change

One way to formalise the Principle of Minimal Change is to consider the relationship between the belief sets obtained when removing different sentences from a belief set K. We will look specifically at what can be said about the union of these resulting belief sets. Consider *any* two sentences θ and ϕ. If $(K \ominus \theta) \cup (K \ominus \phi)$ is inconsistent, there is no relationship at all between these two belief sets. For retraction operators satisfying (Inclusion) this can never occur, but for those satisfying (Dichotomy) it will necessarily be the case. If $K \ominus \phi$ is strictly included in $K \ominus \theta$ then $(K \ominus \theta) \cup (K \ominus \phi) = K \ominus \theta$. By (Success) it follows that $K \ominus \phi \not\models \theta$ and that $(K \ominus \theta) \cup (K \ominus \phi) \not\models \theta$. So we cannot insist that θ be contained in $K \ominus \phi$, nor can we require that the union of $K \ominus \phi$ and $K \ominus \theta$ give us θ. On the other hand, we *can* insist that ϕ should follow from $(K \ominus \theta) \cup (K \ominus \phi)$ since it would be compatible with the axioms of retraction. Moreover, such a requirement is equivalent to forcing ϕ to be in $K \ominus \theta$. That is, $K \ominus \phi$ removes so little from K that *every* belief set stronger than $K \ominus \phi$ which is obtained by removing some sentence from K, will include ϕ. If $K \ominus \theta$ is strictly included in $K \ominus \phi$, the same argument holds, but with the roles of θ and ϕ reversed. The remaining case to consider is where $K \ominus \theta \not\subset K \ominus \phi$, $K \ominus \phi \not\subset K \ominus \theta$, and $(K \ominus \theta) \cup (K \ominus \phi)$ is consistent. Then $K \ominus \theta$ should remove so little from K, and $K \ominus \phi$ should remove so little from K, that $K \ominus \theta$ and $K \ominus \phi$, when put together should yield both θ and ϕ. So:

1. If $(K \ominus \theta) \cup (K \ominus \phi)$ is inconsistent then $(K \ominus \theta) \cup (K \ominus \phi) \models \phi$ (by classical logic).
2. If $K \ominus \phi \subset K \ominus \theta$ then $K \ominus \theta \models \phi$ or, if $K \ominus \phi \subset K \ominus \theta$ then $(K \ominus \theta) \cup (K \ominus \phi) \models \phi$.
3. If $K \ominus \theta \not\subset K \ominus \phi$, $K \ominus \phi \not\subset K \ominus \theta$ and $(K \ominus \theta) \cup (K \ominus \phi)$ is consistent, then $(K \ominus \theta) \cup (K \ominus \phi) \models \theta \wedge \phi$, or, $(K \ominus \theta) \cup (K \ominus \phi) \models \phi$, since the roles of θ and ϕ are symmetrical.

Now observe that the three antecedents in 1, 2 and 3 above make up the three cases that together are equivalent to the condition that $K \ominus \theta \not\subseteq K \ominus \phi$. That is, $K \ominus \theta \not\subseteq K \ominus \phi$ iff $(K \ominus \theta) \cup (K \ominus \phi)$ is inconsistent, or $K \ominus \phi \subset K \ominus \theta$, or $K \ominus \theta \not\subset K \ominus \phi$, $K \ominus \phi \not\subset K \ominus \theta$ and $(K \ominus \theta) \cup (K \ominus \phi)$ is consistent. So 1, 2 and 3 can jointly be restated as the following property:

If $K \ominus \theta \not\subseteq K \ominus \phi$ then $(K \ominus \theta) \cup (K \ominus \phi) \models \phi$

which is the contrapositive of (Strong Conservativity). So (Strong Conservativity) is the Principle of Minimal Change expressed in terms of the relationship between belief sets resulting from the removal of different sentences from a given belief set K.

6 Conclusion

We have provided a formal study of belief change operators that do not satisfy (Inclusion), to do justice to the intuition that dropping a belief may lead to the inclusion of others in the agent's corpus. We provided two models of liberation via retraction operators, *σ-liberation* and *linear liberation*, both of which utilised a finite *sequence* of sentences to guide the operation of belief removal. We showed that the class of σ-liberation operators is included in the class of linear liberation operators, and provided axiomatic characterisations for each class. We also characterised a number of subclasses of linear liberation, including severe withdrawal. We showed how a given retraction operator can be transformed into either a withdrawal operator (satisfying (Inclusion)) or a *revision* operator. Finally we discussed a couple of justifications for the central postulate of σ-liberation, namely (Strong Conservativity).

7 Future Work

For future work we would like to generalise the σ-liberation model. In this paper, the belief sequences σ consisted of sentences which, intuitively, represented previous *revision* inputs which the agent has received. Any previous *retraction* steps which might have taken place are *not* represented. This means that, in this model, we are essentially restricting the domain of σ-liberation to those belief sets K which are formed by a process of revision *alone*. One natural way to record retraction steps would be to allow σ to include so-called *disbeliefs* $\overline{\gamma}$ (where $\gamma \in L$), as seen in [14], where $\overline{\gamma}$ indicates a retraction of γ. This would also open the way for a sequence-based model of *iterated* retraction[10]: when retracting ϕ we can obtain a new sequence by appending $\overline{\phi}$ to the end of σ. This new sequence is then ready for the next input. We intend a full investigation of the properties of such a model. Other directions for further research are to consider more general models that do not satisfy (Vacuity) as well as (Inclusion), and also to find other sequence-based constructions which are able to model operations, such as AGM contraction and systematic withdrawal [13], that cannot be handled with our current ones. Finally, note that in this paper we restricted ourselves to working within a *finitely generated* propositional language L. This choice brought with it certain representational advantages, such as being able to identify any belief set with a single sentence. We would like to consider also the more generalised case involving a countable number of propositional variables.

Acknowledgements

Booth's work is supported by the DFG project "Computationale Dialektik". Thanks are due to David Makinson and to several anonymous referees for some useful comments and suggestions.

References

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

[2] A. Bochman, Belief contraction as nonmonotonic inference, *Journal of Symbolic Logic*, **65** (2000) 605–626.

[3] S. Chopra, K. Georgatos and R. Parikh, Relevance sensitive non-monotonic inference for belief sequences, *Journal of Applied Non-Classical Logics*, **11** (1-2) (2001) 131–150.

[10] A model of iterated *revision* based on sequences may be found in [11].

[4] J. Doyle, Reason maintenance and belief revision: Foundations versus coherence theories, in: P. Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, (1992) 29–51.

[5] M. Freund, On rational preferences, *Journal of Mathematical Economics* **30** (1998) 215–228.

[6] A. Fuhrmann and S. O. Hansson, A survey of multiple contractions, *Journal of Logic, Language and Information*, **3** (1994) 39–76.

[7] P. Gärdenfors, *Knowledge in Flux*, MIT Press, (1988).

[8] S. M. Glaister, Recovery recovered, *Journal of Philosophical Logic*, **29** (2000) 171–206.

[9] S. O. Hansson, Theory contraction and base contraction unified, *Journal of Symbolic Logic*, **58** (1993) 602–625.

[10] S. O. Hansson, *A Textbook of Belief Dynamics*, Kluwer Academic Publishers, (1999).

[11] D. Lehmann, Belief revision, revised, in: *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI'95)* (1995) 1534-1540.

[12] D. Makinson, On the status of the postulate of recovery in the logic of theory change, *Journal of Philosophical Logic*, **16** (1987) 383–394.

[13] T.A Meyer, J. Heidema, W. Labuschagne and L. Leenen, Systematic Withdrawal, *Journal of Philosophical Logic*, **31** 5 (2002) 415-443.

[14] T. A. Meyer, A. K. Ghose and S. Chopra, Non-prioritised ranked belief change, in: *Proceedings of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge (TARK'01)* (2001) 151–162.

[15] B. Nebel, How hard is it to revise a belief base?, in: D. Dubois and H. Prade (eds.), *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Vol. 3: Belief Change*, Kluwer Academic Publishers, (1998) 77–145.

[16] R. Reiter, A logic for default reasoning, *Artificial Intelligence*, **13** (1980) 81–132.

[17] H. Rott and M. Pagnucco, Severe withdrawal (and recovery), *Journal of Philosophical Logic*, **28** (1999) 501–547.