A Unifying Semantics for Belief Change

C0300

Abstract. Many belief change formalisms employ plausibility orderings over the set of possible worlds to determine how the beliefs of an agent ought to be modified after the receipt of a new epistemic input. While most such possible world semantics rely on a single ordering, we look at using an extra ordering to aid in guiding the process of belief change. We show that this provides a unifying semantics for a wide variety of belief change operators. By varying the conditions placed on the second ordering, different families of known belief change operators can be captured, including AGM belief contraction and revision [1], severe withdrawal [14], systematic withdrawal [12], and the linear liberation and σ-liberation operators of [3]. Our approach also identifies novel classes of belief change operators that are worth further investigation.

1 INTRODUCTION

Current formalisms in belief change [5, 9] typically employ a plausibility ordering [6, 10] over the set of possible worlds or an epistemic entrenchment ordering over the set of sentences in an agent's belief set. Operators for change are then defined by manipulation of these orderings after receipt of a new epistemic input. There are many advantages to these approaches – foremost amongst them the guarantee that change will be effected in a principled manner, the provision of an intuitively plausible construction, and a formalism flexible enough to accommodate alternative change strategies and iteration. However there are some nuances that are not captured in such an approach. For instance, agents do not usually employ one fixed ordering throughout – often, different orderings might be used in different contexts such as those requiring greater caution or skepticism. Or different orderings might be used based on the source of the epistemic inputs. Such a critique is implicit in [4] where the notion of *eligibility* adds an extra dimension to belief change. A technical framework that provides tools for belief change operations based on multiple orderings appears in [2] where combination operations for a class of preference relations $\mathcal{P}$ are studied in terms of an additional guiding preference relation. In this study, the formalism for belief change – in particular for belief removal – that we will present can be considered a special case of [2] with $\leq$ – over the set of interpretations – being the single preference relation in $\mathcal{P}$, and $\preceq$ – our additional dimension – being the guiding relation.

An intuitive way to understand the second ordering on the set of worlds is to think of it as representing a more stringent assessment of the plausibility of states of affairs. Most rational agents are aware of certain contexts within which their reasoning plays out – certain contexts call for a different assessment of plausibility. For example, I enforce a certain amount of skepticism on verifying news reports – but will probably fall back on a more critical assessment when I'm trying to assess news reports in a different situation, say the impending declaration of a war. Such a treatment is reminiscent of contextualist assessments of epistemic statements – it is understood that the agent makes any knowledge claim relative to some implicit standard for assessing that claim and that different standards will induce differing assessments of the truth of epistemic claims. The contribution of the paper is the unification, in a single formal framework, of a large class of belief change operators by this method. It enables us to view belief change as the manipulation by the agent of assessments of plausibility of epistemic states of affairs in different contexts.

The plan of the paper is as follows. After laying down some technical preliminaries, in Sect. 2 we establish the foundations of our framework for removal with a semantic definition and an axiomatic characterisation. In Sect. 3 we study the class of belief removal operators obtained when the second ordering $\preceq$ is transitive. Sect. 4 builds up to a characterisation of *AGM contraction* [1] via sub-classes of belief removal operators satisfying the standard properties known as Vacuity, Inclusion and Recovery. Sect. 5 shows that important classes of *belief liberation* operators [3] can be captured in our framework. Sect. 6 isolates various classes of removal operators related to, and including, *systematic withdrawal* [12]. Sect. 7 shows that the limiting cases correspond to *AGM revision* [1] and *severe withdrawal* [14], while Sect. 8 concludes with some pointers to future work.

We assume a finitely generated propositional language L equipped with the usual constants, boolean operators and a classical Tarskian consequence relation Cn. $\mathcal{W}$ denotes the set of possible worlds/interpretations of L. Logical entailment is denoted by $\models$. For any set of sentences $A \subseteq L$, $[A]$ denotes the set of worlds satisfying all members of A (writing $[\phi]$ rather than $[\{\phi\}]$ for the singleton case). For a set $S \subseteq \mathcal{W}$, $Th(S)$ is the set of sentences true in all worlds in S. The object which undergoes change will be K, a consistent belief set (i.e., a deductively closed, consistent set of sentences). We take K to be arbitrary but fixed throughout. We assume that for all removal operators $\ominus$, $K \ominus \phi$ is only defined for non-tautologous propositions and refer to the set of non-tautologous members of L as L_*. The limiting case requires only a minor emendation. We make this choice for ease of technical presentation. Finally, given a total pre-order (i.e., a transitive, connected relation) $\leq$ on $\mathcal{W}$ and $S \subseteq \mathcal{W}$, $\min(S, \leq)$ will denote the set of $\leq$-minimal elements of S.

2 BASIC REMOVAL

We now set up our most general semantic construction of belief change operators. We refer to these as *removal* operators because the nett effect after being presented with an input ϕ is that ϕ is removed from the belief set. However, as we shall see in Sect. 7, the extreme case where the removal of a belief ϕ results in the addition of $\neg\phi$ is included in the framework.

Assume a total pre-order $\leq$ *anchored* on $[K]$. That is to say, $[K] = \min(\mathcal{W}, \leq)$. As usual we take $\leq$ to be an ordering of plausibility on the worlds, with worlds lower down in the ordering seen as more plausible. In what follows, $\sim$ will always denote the symmetric closure of $\leq$, i.e., $w_1 \sim w_2$ iff both $w_1 \leq w_2$ and $w_2 \leq w_1$. Now we assume that we are given a *second* binary relation $\preceq$ on $\mathcal{W}$, which we require to be a reflexive sub-relation of $\leq$. These two orderings provide the *context* in which an agent makes changes to its current beliefs.

Definition 1 *$(\leq, \preceq)$ is a K-*context *iff $\leq$ is a total pre-order (on $\mathcal{W}$) anchored on $[K]$, and $\preceq$ is a reflexive sub-relation of $\leq$.*

Given a belief set K and a K-context $(\leq, \preceq)$, we use $(\leq, \preceq)$ to define a *removal operator* $\ominus_{(\leq,\preceq)}$ for K by setting, for all $\phi \in L_*$,

$$K \ominus_{(\leq,\preceq)} \phi = Th(\{w \mid w \preceq w' \text{ for some } w' \in \min([\neg\phi], \leq)\})$$

That is, the models of the belief set resulting from a removal of ϕ are obtained by locating all the $\leq$-best models of $\neg\phi$, and adding to those, all worlds that are at least as $\preceq$-plausible.

Definition 2 *$\ominus$ is a* basic removal operator *(for K) iff $\ominus = \ominus_{(\leq,\preceq)}$ for some K-context $(\leq, \preceq)$.*

Basic removal is characterised by the following postulates:

(B1) $K \ominus \phi = Cn(K \ominus \phi)$
(B2) $\phi \notin K \ominus \phi$
(B3) If $\models \phi_1 \leftrightarrow \phi_2$ then $K \ominus \phi_1 = K \ominus \phi_2$
(B4) $K \ominus \bot = K$
(B5) $K \ominus \phi \subseteq Cn(K \cup \{\neg\phi\})$
(B6) If $\theta \in K \ominus (\theta \wedge \phi)$ then $\theta \in K \ominus (\theta \wedge \phi \wedge \psi)$
(B7) If $\theta \in K \ominus (\theta \wedge \phi)$ then $K \ominus \phi \subseteq K \ominus (\theta \wedge \phi)$
(B8) $(K \ominus \theta) \cap (K \ominus \phi) \subseteq K \ominus (\theta \wedge \phi)$
(B9) If $\phi \notin K \ominus (\theta \wedge \phi)$ then $K \ominus (\theta \wedge \phi) \subseteq K \ominus \phi$

Theorem 1 *Let K be a belief set and $\ominus$ an operator for K. Then $\ominus$ is a basic removal operator for K iff $\ominus$ satisfies* **(B1)–(B9)**.

All the rules above are already familiar from the belief change literature. Rules **(B1)–(B3)** belong to the six *basic AGM contraction postulates* [1]. Rules **(B4)** and **(B5)** are weakened versions – under our assumption that K is consistent – of another of the basic AGM postulates, namely the Vacuity rule:

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

As we will confirm in Sect. 4, basic removal operators do not generally satisfy **(Vacuity)**. The remaining two basic AGM contraction rules, neither of which are sound for basic removal, are:

(Inclusion) $K \ominus \phi \subseteq K$
(Recovery) $K \subseteq Cn((K \ominus \phi) \cup \{\phi\})$

(Inclusion) is questioned in [3], leading to the study of *belief liberation* operators, while **(Recovery)** has been questioned in many places in the literature (e.g. [7, 9]). Briefly, liberation operators cater to the intuition that removing a belief from an agent's corpus can remove the reasons for not holding others and hence lead to the inclusion of new beliefs. Of the other postulates for basic removal above, **(B8)** and **(B9)** are the two *supplementary* AGM contraction postulates [1], while **(B6)** and **(B7)** both follow from the AGM postulates (see [1, 7, 13]). The latter rule is closely related to the well-known rule *Cautious Monotony* from non-monotonic inference [11].

The completeness part of Theorem 1 is proved by using the following way to construct a pair of orderings from a given belief set and basic removal operator.

Definition 3 *The structure $(\leq, \preceq)$ obtained from a belief set K and a basic removal operator $\ominus$, and denoted by $\mathcal{C}(K, \ominus)$ is defined as follows (cf. [4]), for $w_1, w_2 \in \mathcal{W}$:*
$(\leq)$ $\quad w_1 \leq w_2$ *iff* $\neg\alpha_1 \notin K \ominus (\neg\alpha_1 \wedge \neg\alpha_2)$
$(\preceq)$ $\quad w_1 \preceq w_2$ *iff* $\neg\alpha_1 \notin K \ominus \neg\alpha_2$

where α_i is a sentence whose only model is w_i (for $i = 1, 2$).

In the theorem, $\mathcal{C}(K, \ominus)$ is used by checking that if $\ominus$ satisfies **(B1)–(B9)**, then $(\leq, \preceq)$ is a K-context and that $\ominus = \ominus_{(\leq,\preceq)}$. We employ this construction throughout the paper to prove that certain postulates are complete for certain sub-classes of basic removal.

We now proceed to investigate how different requirements on the second ordering of plausibility $\preceq$ and its interplay with $\leq$ can help us characterise different belief removal operations. We start with one of the simplest properties there is – transitivity.

3 TRANSITIVE REMOVAL

In this section we see what happens if we let the second order $\preceq$ be transitive, i.e., $\preceq$ becomes a pre-order. We'll call the K-context $(\leq, \preceq)$ transitive if $\preceq$ is transitive.

Definition 4 *We call $\ominus$ a* transitive removal operator *(for K) iff $\ominus = \ominus_{(\leq,\preceq)}$ for some transitive K-context $(\leq, \preceq)$.*

Transitive removal operators may be alternatively described as follows. As with any pre-order, the relation $\preceq$ partitions $\mathcal{W}$ into a set $\mathcal{W}/\equiv$ of equivalence classes via the relation $\equiv$ defined by $w_1 \equiv w_2$ iff both $w_1 \preceq w_2$ and $w_2 \preceq w_1$. The set $\mathcal{W}/\equiv$ is partially-ordered by the relation $\preceq^*$ defined by $[w_1]_\equiv \preceq^* [w_2]_\equiv$ iff $w_1 \preceq w_2$. Meanwhile, we can also define a relation $\leq^*$ on $\mathcal{W}/\equiv$ by $[w_1]_\equiv \leq^* [w_2]_\equiv$ iff $w_1 \leq w_2$. It is easy to check that $\leq^*$ is well-defined and that $\leq^*$ is a total pre-order on $\mathcal{W}/\equiv$ such that $\preceq^* \subseteq \leq^*$. Furthermore we have, for each $\phi \in L_*$, $K \ominus_{(\leq,\preceq)} \phi = Th(\bigcup \Upsilon)$, where

$$\Upsilon = \{X \in \mathcal{W}/\equiv \mid X \preceq^* Y \text{ for some } Y \in \min(\neg\phi, \leq^*)\},$$

and where $\min(\neg\phi, \leq^*)$ here denotes the set of $\leq^*$-minimal elements $Y \in \mathcal{W}/\equiv$ such that $Y \cap [\neg\phi] \neq \emptyset$. Note how worlds belonging to the same equivalence class are 'indistinguishable' to the agent using the K-context $(\leq, \preceq)$.

The next result shows how we can axiomatically characterise the class of transitive removal operators.

Theorem 2 (i). *If $(\leq, \preceq)$ is transitive then $\ominus_{(\leq,\preceq)}$ satisfies:*

(BT) *If $K \ominus \theta \not\subseteq K \ominus \phi$ then there exist $\psi, \lambda \in L_*$ such that $\phi \models \psi \models \lambda$ and $(K \ominus \theta) \cup (K \ominus \lambda) \models \psi$*

(ii). *If $\ominus$ satisfies* **(BT)** *then the relation $\preceq$ of $\mathcal{C}(K, \ominus)$ is transitive.*

So transitive removal operators may be characterised by **(B1)–(B9)** plus **(BT)**. **(BT)**, as might be noted, is a very weak requirement. One natural way to strengthen it is to require that $\psi = \phi$:

(BConserv) If $K \ominus \theta \not\subseteq K \ominus \phi$ then there exists $\lambda \in L_*$ such that $\phi \models \lambda$ and $(K \ominus \theta) \cup (K \ominus \lambda) \models \phi$

(BConserv) looks almost the same as the rules Conservativity and Weak Conservativity, which were proposed and argued-for in [8, 9] and used there to characterise operations of so-called *base-generated contraction*.

It turns out that, for basic removal operators, **(BConserv)** may be captured by requiring that, in addition to being transitive, $(\leq, \preceq)$ satisfies the following property:

(a) If $w_1 \sim w_2$ and $w_1 \preceq w_2$ then $w_2 \preceq w_1$

Theorem 3 (i). *If $(\leq, \preceq)$ is transitive and satisfies (a) then $\Bumpeq_{(\leq,\preceq)}$ satisfies* **(BConserv)**. (ii). *If $\Bumpeq$ satisfies* **(BConserv)** *then $\mathcal{C}(K, \Bumpeq)$ is transitive and satisfies (a).*

In terms of the alternative description of transitive removal given above in terms of equivalence classes, requiring (a) of $(\leq, \preceq)$ has the effect that the relations $\leq^*$ and $\preceq^*$ on $\mathcal{W}/\equiv$ satisfy, for all $X, Y \in \mathcal{W}/\equiv$: $X \preceq^* Y$ implies $X <^* Y$ or $X = Y$, where $<^*$ is the strict part of $\leq^*$. Thus any two distinct classes X, Y which are on the same 'level' according to $\leq^*$ (in that both $X \leq^* Y$ and $Y \leq^* X$) are incomparable according to $\preceq^*$.

As we will see, although **(BConserv)** is more restrictive than **(BT)**, the class of basic removals satisfying it remains general enough to include many other important sub-classes of basic removal.

By going a step further and identifying λ with ϕ in **(BConserv)** we arrive at a yet stronger postulate:

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

(BSConserv) is known as Strong Conservativity [8], and is used in [3] to help characterise the so-called *σ-liberation* operators (see Sect. 5). [3] also contains a detailed justification for the use of this rule. For basic removal, we can capture this property by requiring the following property, in conjunction with transitivity:

(b) If $w_1 \sim w_2$ then $w_1 \preceq w_2$

Theorem 4 (i). *If $(\leq, \preceq)$ is transitive and satisfies (b) then $\Bumpeq_{(\leq,\preceq)}$ satisfies* **(BSConserv)**. (ii). *If $\Bumpeq$ satisfies* **(BSConserv)** *then $\mathcal{C}(K, \Bumpeq)$ is transitive and satisfies (b).*

Condition (b) implies (a). In terms of the above construction in terms of $\mathcal{W}/\equiv$, having that $\preceq$ is transitive while strengthening (a) to (b) has the effect that the relation $\leq^*$ becomes a *total order* on $\mathcal{W}/\equiv$.

4 TOWARDS AGM CONTRACTION

It was mentioned in Sect. 2 that basic removal does not satisfy the three basic AGM contraction postulates **(Vacuity)**, **(Inclusion)** and **(Recovery)**. In Sect. 7 it is shown that the severe withdrawal operators, which are known not to satisfy **(Recovery)** [14], are all basic removal operators, thus proving that **(Recovery)** fails for basic removal. For the failure of the other two rules, suppose $K = Cn(\emptyset)$ and consider the K-context $(\leq, \preceq)$ where $\leq$ is the full relation $\mathcal{W} \times \mathcal{W}$ and $\preceq$ is just the equality relation. Then it is easy to check that, for any consistent $\phi \in L_*$, we get $K \Bumpeq_{(\leq,\preceq)} \phi = Th([\neg\phi]) = Cn(\neg\phi)$. Thus $K \Bumpeq_{(\leq,\preceq)} \phi \not\subseteq K$, even though $\phi \notin K$. 'One half' of **(Vacuity)**, however, *is* valid for basic removal:

Proposition 1 *Let $\Bumpeq$ be a basic removal operator for K, then $\Bumpeq$ satisfies: If $\phi \notin K$ then $K \subseteq K \Bumpeq \phi$*

The 'missing half' of **(Vacuity)** is: If $\phi \notin K$ then $K \Bumpeq \phi \subseteq K$. Clearly this rule doubles as a weakened version of **(Inclusion)**. Thus we see that, for basic removal operators, **(Inclusion)** actually implies **(Vacuity)**. Now let's verify under what conditions on $(\leq, \preceq)$ each of these postulates are satisfied by basic removal operators.

4.1 Vacuity

To ensure that $\Bumpeq_{(\leq,\preceq)}$ satisfy all of **(Vacuity)**, we require that all $\leq$-minimal elements (i.e., all elements of $[K]$) are $\preceq$-connected, i.e.,

(c) If (for each $i = 1, 2$) $w_i \leq w'$ for all w', then $w_1 \preceq w_2$

Theorem 5 (i). *If $(\leq, \preceq)$ satisfies (c) then $\Bumpeq_{(\leq,\preceq)}$ satisfies* **(Vacuity)**. (ii). *If $\Bumpeq$ satisfies* **(Vacuity)** *then $\mathcal{C}(K, \Bumpeq)$ satisfies (c).*

As is easily verified, (c) is implied by condition (b). Thus we see that any basic removal satisfying **(BSConserv)** satisfies **(Vacuity)**. However, our counter-example above shows that **(Vacuity)** is not valid for transitive removals satisfying (a).

Shouldn't **(Vacuity)** be a basic requirement for *any* rational removal operation? From a purely *minimal change* point of view it is certainly hard to contest, but we would nevertheless argue that there *are* plausible scenarios in which it can fail. Consider an agent who has equally good reasons to believe each of p and $\neg p$. In this situation the agent remains cautious and commits to believe neither p nor $\neg p$. But if this agent were then to receive information that undermines p then it seems plausible that it would come to believe (or assign significantly more plausibility to) $\neg p$.

Of course one could always try and *force* a given basic removal $\Bumpeq$ to satisfy **(Vacuity)** by defining a new operator $\div$ from $\Bumpeq$ by $K \div \phi = K$ if $\phi \notin K$, $K \div \phi = K \Bumpeq \phi$ otherwise. It is fairly straightforward to show that $\div$ so defined satisfies **(B1)–(B9)**, and so again forms a basic removal. However we run into difficulties in the case of transitive removal, for it turns out that rule **(BT)** is *not* preserved. Exploring ways out of this problem will be left for future work.

4.2 Inclusion

To obtain **(Inclusion)** we may add the following condition, stronger than (c):

(d) If $w_1 \leq w_2$ for all w_2 then $w_1 \preceq w_2$ for all w_2

So the $\leq$-minimum worlds are also the $\preceq$-minimum worlds.

Theorem 6 (i). *If $(\leq, \preceq)$ satisfies (d) then $\Bumpeq_{(\leq,\preceq)}$ satisfies* **(Inclusion)**. (ii). *If $\Bumpeq$ satisfies* **(Inclusion)** *then $\mathcal{C}(K, \Bumpeq)$ satisfies (d).*

Note that, even though basic removal operators do not satisfy **(Inclusion)** in general, it is always possible to *transform* a given basic removal $\Bumpeq$ into an operator which *does* satisfy that rule. We simply take the *incarceration* $\bumpeq$ of $\Bumpeq$ [3], i.e., the operator defined from $\Bumpeq$ by $K \bumpeq \phi = K \cap (K \Bumpeq \phi)$. It can be shown that the incarceration of a basic removal operator is always itself a basic removal, while furthermore if $\Bumpeq$ satisfies any of the three postulates from Sect. 3, then $\bumpeq$ will satisfy the same ones as well.

4.3 Recovery

To obtain **(Recovery)** it suffices to require the following condition:

(e) If $w_1 \preceq w_2$ then $w_1 = w_2$ or $w_1 \leq w'$ for all w'

So, apart from itself, nothing but $\leq$-minimal worlds may be below any world in $\preceq$.

Theorem 7 (i). *If $(\leq, \preceq)$ satisfies (e) then $\Bumpeq_{(\leq,\preceq)}$ satisfies* **(Recovery)**. (ii). *If $\Bumpeq$ satisfies* **(Recovery)** *then $\mathcal{C}(K, \Bumpeq)$ satisfies (e).*

The combination of (d) and (e) then states that the worlds below a world w in $\preceq$ are exactly w itself and the $\leq$-minimal worlds. And this gives us precisely AGM contraction (satisfying the basic plus supplementary AGM contraction postulates).

Theorem 8 *The following are equivalent:*
(i). *$\Bumpeq$ is a full AGM contraction operator.*
(ii). *$\Bumpeq$ satisfies* **(B1)–(B9)** *plus* **(Inclusion)** *and* **(Recovery)**.
(iii). *$\Bumpeq = \Bumpeq_{(\leq,\preceq)}$ for some $(\leq, \preceq)$ which satisfies (d) and (e).*

Observe that since (d)+(e) implies transitivity and (a), every full AGM contraction is a basic removal satisfying **(BConserv)**.

5 BELIEF LIBERATION

In [3] two models of belief liberation operators are presented, each in terms of finite sequences of sentences. The second model, *linear liberation*, is more general than the first, σ*-liberation*. The class of liberation operators it generates includes that generated by the first. The first construction employs a linearly ordered sequence of sentences and the second a set of candidate belief sets one of which corresponds to the agent's set after belief retraction. Axiomatic characterisations of each of these classes are also provided in [3]. Linear liberation is characterised by **(B1)–(B3)** plus **(Vacuity)** and the following rule:
(Hyperreg) If $\theta \notin K \circleddash (\theta \wedge \phi)$ then $K \circleddash (\theta \wedge \phi) = K \circleddash \theta$

This is the rule originally known as Hyperregularity from [8]. The first thing to note about **(Hyperreg)** is that, in the presence of **(B1)–(B4)**, it actually implies **(Vacuity)** *and* the remaining rules for basic removal **(B5)–(B9)**. Thus we see:

Proposition 2 *$\circleddash$ is a linear liberation operator iff it is a basic removal operator which satisfies* **(Hyperreg)**.

Is there a condition on $(\leq, \preceq)$ which corresponds exactly to **(Hyperreg)**? It turns out that the following condition does the trick:

(f) If $w_1 \sim w_2$ and $w_3 \preceq w_1$ then $w_3 \preceq w_2$

Rule (f) says that whether or not a world w_3 is below w_1 according to $\preceq$ depends only on the $\leq$-plausibility rank of w_1.

Theorem 9 (i). *If $(\leq, \preceq)$ satisfies (f) then $\circleddash_{(\leq,\preceq)}$ satisfies* **(Hyperreg)**. (ii). *If $\circleddash$ satisfies* **(Hyperreg)** *then $\mathcal{C}(K, \circleddash)$ satisfies (f).*

Thus we see that linear liberation operators may be represented by the class of K-contexts which satisfy (f).

In [3] it is shown that the σ-liberation operators are precisely those linear liberation operators which satisfy **(BSConserv)**. Using this fact together with Theorems 4 and 9 allows us to deduce:

Proposition 3 *$\circleddash$ is a σ-liberation operator iff $\circleddash = \circleddash_{(\leq,\preceq)}$ for some transitive $(\leq, \preceq)$ satisfying (b) and (f).*

However, we can simplify here, for as soon as $\preceq$ is transitive, conditions (b) and (f) become *equivalent*:

Proposition 4 *Let $(\leq, \preceq)$ be a transitive K-context. Then $(\leq, \preceq)$ satisfies (b) iff $(\leq, \preceq)$ satisfies (f).*

This means that in Prop. 3 it is unnecessary to require both (b) *and* (f) – just one of them will suffice. Depending on which one we choose to retain, we obtain two different characterisations of σ-liberation which provide alternatives to the one from [3]:

Theorem 10 *The following are equivalent:*
(i). *$\circleddash$ is a σ-liberation operator.*
(ii). *$\circleddash$ is a linear liberation operator which satisfies* **(BT)**.
(iii). *$\circleddash$ is a basic removal operator which satisfies* **(BSConserv)**.

The equivalence *(i)⇔(ii)* comes from combining Prop. 3 (retaining just (f)) with Theorems 2 and 9, while *(i)⇔(iii)* comes from combining Prop. 3 (retaining just (b)) with Theorem 4. Surprisingly, *(i)⇔(ii)* says that, in the axiomatisation of σ-liberation in [3], **(BSConserv)** may be replaced by the seemingly much weaker **(BT)**. Meanwhile, since *(i)⇔(iii)*, σ-liberation operators inherit the nice description in terms of $\mathcal{W}/\equiv$ given for the basic removals which satisfy **(BSConserv)** at the end of Sect. 3 (where $\leq^*$ is a total order on $\mathcal{W}/\equiv$).

Similar characterisations for sub-classes of liberation, such as the class of *dichotomous* liberation operators [3], exist. However, space considerations prevent us here from embroidering further on this theme.

6 SYSTEMATIC WITHDRAWAL

An interesting sub-class of basic removal operators, which includes both systematic [12] and severe withdrawal [14] (see below) is obtained by requiring the following condition on $(\leq, \preceq)$:

(g) If $w_1 < w_2$ then $w_1 \preceq w_2$

where $<$ is the strict part of $\leq$.

Theorem 11 (i). *If $(\leq, \preceq)$ satisfies (g) then $\circleddash_{(\leq,\preceq)}$ satisfies:*

(B10) *If $\theta \in K \circleddash (\theta \wedge \phi)$ then $\phi \notin K \circleddash \theta$*

(ii). *If $\circleddash$ satisfies* **(B10)** *then $\mathcal{C}(K, \circleddash)$ satisfies (g).*

The class of basic removal operators $\circleddash_{(\leq,\preceq)}$ such that $(\leq, \preceq)$ satisfies (g) still do not generally satisfy **(Inclusion)** or **(Vacuity)**, since condition (g) does not rule out that some $\leq$-minimal elements may be $\preceq$-unconnected. However they do come *mighty* close to satisfying **(Inclusion)**, in that the following is satisfied:

If $\theta \in K$ then $K \circleddash \theta \subseteq K$

Using this fact we can see that for *this* class of operators, **(Inclusion)** and **(Vacuity)** are equivalent.

The next condition on K-contexts is, essentially, a requirement for antisymmetry to hold:

(h) If $w_1 \preceq w_2$ then either $w_1 < w_2$ or $w_1 = w_2$

Theorem 12 (i). *If $(\leq, \preceq)$ satisfies (h) then $\circleddash_{(\leq,\preceq)}$ satisfies:*

(B11) *If $\models (\theta \vee \phi)$ and $\theta \notin K \circleddash \phi$ then $\phi \in K \circleddash (\theta \wedge \phi)$*

(ii). *If $\circleddash$ satisfies* **(B11)** *then $\mathcal{C}(K, \circleddash)$ satisfies (h).*

Clearly, by requiring (h) in combination with (g) (and reflexivity) we specify $\preceq$ uniquely:

(g)+(h) $w_1 \preceq w_2$ iff either $w_1 < w_2$ or $w_1 = w_2$

Note that $\preceq$ so defined will automatically be transitive and will satisfy the condition (a) from Sect. 3. Putting together Theorems 11 and 12, then, we have that the class of basic removal operators $\circleddash_{(\leq,\preceq)}$ where $\preceq$ is defined via (g)+(h) may be axiomatically characterised by **(B1)–(B11)**. This looks very much like the class of systematic withdrawals. A systematic withdrawal operator $\div$ can be defined in terms of $\leq$ as follows [12]:

$$K \div \phi = K \cap Th(\nabla_{\leq}(\min([\neg\phi], \leq)))$$

where $\nabla_{\leq}(X) = \{v \mid \exists w \in X \text{ s.t. } v < w\}$. Unlike systematic withdrawal, the class of removal operators defined by **(B1)–(B11)** *fails* to satisfy **(Inclusion)**/**(Vacuity)**, since all the $\leq$-minimal elements are necessarily *unconnected* according to $\preceq$. So in fact **(Vacuity)** will fail as soon as there is more than one $\leq$-minimal element. These operators satisfy instead:

If $\phi \notin K$ then $\neg\phi \in K \circleddash \phi$

That is, for these operators, we see that $K \circleddash \phi$ is an operation which 'demotes' the status of ϕ: if its current status is 'accepted', i.e., $\phi \in K$, then its status is 'demoted' to 'undecided' i.e., $\phi, \neg\phi \notin K \circleddash \phi$, while if its current status is 'undecided' then its status is 'demoted' to 'rejected'. If its status is already 'rejected' then no change occurs. However, if we take the incarcerations of these operators then we end up with precisely the class of systematic withdrawal operators.

Systematic withdrawal can also be obtained by weakening (h):

(j) If $w_1 \preceq w_2$ then $w_1 < w_2$, $w_1 = w_2$, or $w_1 \leq w' \; \forall w'$

So, unlike (h), (j) allows the models of K to be connected according to $\preceq$, although it does not force them to be.

Theorem 13 (i). *If $(\leq, \preceq)$ satisfies (j) then $\ominus_{(\leq,\preceq)}$ satisfies:*

(B12) *If $\phi \in K$, $\models (\theta \vee \phi)$ and $\theta \notin K \ominus \phi$ then $\phi \in K \ominus (\theta \wedge \phi)$*
(ii). *If $\ominus$ satisfies* **(B12)** *then $\mathcal{C}(K, \ominus)$ satisfies (j).*

Since the operators obtained from (g) and (h) form a sub-class of the operators obtained from (g) and (j), the latter class still does not satisfy **(Vacuity)**. But adding (a) (and therefore **(Vacuity)**) to (g) and (j) leads exactly to systematic withdrawal.

Theorem 14 *The following are equivalent:*
(i). *$\ominus$ is a systematic withdrawal.*
(ii). *$\ominus$ satisfies* **(B1)–(B9)** *plus* **(Vacuity)**, **(B10)** *and* **(B12)**.
(iii). *$\ominus = \ominus_{(\leq,\preceq)}$ for some $(\leq, \preceq)$ which satisfies (a), (g) and (j).*

As we shall see in the next section, the class of severe withdrawals can also be isolated in a similar manner.

7 LIMITING CASES

We have seen that the addition of the second ordering $\preceq$ provides us with considerable flexibility when defining removal operators. But what happens when we focus on the limits imposed on $\preceq$? In this section we consider the two cases where $\preceq$ is the *smallest* and the *largest* reflexive sub-relation of $\leq$. If we take $\preceq$ to be the smallest $\preceq$, the equality relation, then the operator $\ominus_{(\leq,\preceq)}$ reduces to:

$$K \ominus_{(\leq,\preceq)} \phi = Th(\min([\neg\phi], \leq)).$$

and we have the following result.

Theorem 15 (i). *If $\preceq$ is the equality relation then $\ominus_{(\leq,\preceq)}$ satisfies:*

(B13) *$\neg\phi \in K \ominus \phi$.*

(ii). *If $\ominus$ satisfies* **(B13)** *then $\preceq$ in $\mathcal{C}(K, \ominus)$ is the equality relation.*

Thus we see that removing ϕ here amounts to a *revision* by its negation, and in fact that $\ominus_{(\leq,\preceq)}$ essentially reduces to an AGM revision function (satisfying the full list of AGM revision postulates [1]). More precisely the operator $*_{(\leq,\preceq)}$ for K defined by

$$K *_{(\leq,\preceq)} \phi = K \ominus_{(\leq,\preceq)} \neg\phi$$

is an AGM revision operator. Moreover, *every* AGM revision operator can be obtained in this way. Note that in the above case, since $\phi \in K \ominus_{(\leq,\preceq)} \neg\phi$, the right-hand side here is equal to $Cn((K \ominus_{(\leq,\preceq)} \neg\phi) \cup \{\phi\})$. Thus what we have is just the Levi Identity [5]. In fact a result more general holds: *whenever* $(\leq, \preceq)$ is a K-context and $*_{(\leq,\preceq)}$ is defined from $\ominus_{(\leq,\preceq)}$ via the Levi Identity then $*_{(\leq,\preceq)}$ is an AGM revision operator.

By taking $\preceq$ to be the largest reflexive sub-relation of $\leq$ we get the full relation $\leq$, and the operator $\ominus_{(\leq,\preceq)}$ reduces to:

$$K \ominus_{(\leq,\preceq)} \phi = Th(\{w \mid w \leq w' \text{ for some } w' \in \min([\neg\phi], \leq)\}).$$

Thus, from the characterisation of severe withdrawal in terms of total pre-orders found in [14], we see clearly that setting $\preceq = \leq$ gives us the class of severe withdrawal operators. Note that $\preceq$ so defined will be transitive and satisfy condition (b) from Sect. 3 (and hence also (f) – see Prop. 4). From the results above it turns out that we can give an axiomatic characterisation of severe withdrawal which is different to the ones found in the literature (see [14]). To do this first note the following:

Proposition 5 *Let $(\leq, \preceq)$ be a K-context. Then $\preceq = \leq$ iff both (f) and (g) are satisfied.*

Using this fact with Theorems 9 and 11 then yields:

Theorem 16 *$\ominus$ is a severe withdrawal operator iff it satisfies* **(B1)–(B4)**, **(Hyperreg)** *and* **(B10)**.

8 CONCLUSION

In this study we have presented a unified framework for belief removal in terms of a possible world semantics which is distinctive in that it uses a pair of orderings over the set of worlds. We argued for the conceptual plausibility of this pair and showed how a large class of belief removal operators such as liberation, systematic and severe withdrawal operators could be characterised. This approach opens the door for identifying hitherto unstudied sub-classes of basic removal operators, such as those obtained by requiring of $\preceq$ to be a total pre-order and a partial order.An obvious generalisation to consider in future work is the extension to propositional languages with a countably infinite number of propositional variables. Also, a detailed study of the connection between basic removal, base-generated contraction, and sequence-based retraction is of interest. Finally, as in any formalism for belief change, we need to consider iterated removal and how this affects the adjustment of worlds in both $\leq$ and $\preceq$, as well as the interplay between $\leq$ and $\preceq$.

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**, 510–530, (1985).

[2] H Andreka, M. Ryan, and P. Y. Schobbens, 'Operators and laws for combining preference relations', *Journal of Logic and Computation*, **12**(1), 13–53, (2002).

[3] R. Booth, S. Chopra, A. Ghose, and T. Meyer, 'Belief liberation (and retraction)', in *Proceedings of the Ninth Conference on Theoretical Aspects of Rationality and Knowledge (TARK'03)*, pp. 159–172, (2003).

[4] J. Cantwell, 'Eligible contraction', *Studia Logica*, **73**, 167–182, (2003).

[5] P. Gärdenfors, *Knowledge in Flux*, MIT Press, 1988.

[6] A. Grove, 'Two modellings for theory change', *Journal of Philosophical Logic*, **17**, 157–170, (1988).

[7] S. O. Hansson, 'Changes of disjunctively closed bases', *Journal of Logic, Language and Information*, **2**(4), 255–284, (1993).

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

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

[10] H. Katsuno and A.O. Mendelzon, 'Propositional knowledge base revision and minimal change', *Artificial Intelligence*, **52**, 263–294, (1991).

[11] S. Kraus, D. Lehmann, and M. Magidor, 'Nonmonotonic reasoning, preferential models and cumulative logics', *Artificial Intelligence*, **44**, 167–207, (1991).

[12] T. Meyer, J. Heidema, W. Labuschagne, and L. Leenen, 'Systematic withdrawal', *Journal of Philosophical Logic*, **31**, 415–443, (2002).

[13] H. Rott, 'Preferential belief change using generalized epistemic entrenchment', *Journal of Logic, Language and Information*, **1**, 45–78, (1992).

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