

A General Family of Preferential Belief Removal Operators

Richard Booth¹, Thomas Meyer², and Chattrakul Sombattheera¹

¹ Mahasarakham University, Faculty of Informatics, Mahasarakham 44150 Thailand
{richard.b, chattrakul.s}@msu.ac.th

² Meraka Institute, CSIR and School of Computer Science, University of
Kwazulu-Natal, South Africa
tommie.meyer@meraka.org.za

Abstract. Most belief change operators in the AGM tradition assume an underlying plausibility ordering over the possible worlds which is transitive and *complete*. A unifying structure for these operators, based on supplementing the plausibility ordering with a second, guiding, relation over the worlds was presented in [5]. However it is not always reasonable to assume completeness of the underlying ordering. In this paper we generalise the structure of [5] to allow incomparabilities between worlds. We axiomatise the resulting class of belief removal functions, and show that it includes an important family of removal functions based on *finite prioritised belief bases*.

1 Introduction

The problem of *belief removal* [1,5,19], i.e., the problem of what an agent, hereafter \mathcal{A} , should believe after being directed to remove some sentence from his stock of beliefs, has been well studied in philosophy and in AI over the last 25 years. During that time many different families of removal functions have been studied. A great many of them are based on constructions employing *total preorders* over the set of possible worlds which is meant to stand for some notion \leq of relative *plausibility* [12]. A unifying construction for these families was given in [5], in which a general construction was proposed which involved supplementing the relation \leq with a second, guiding, relation \preceq which formed a subset of \leq . By varying the conditions on \preceq and its interaction with \leq many of the different families can be captured as instances.

The construction in [5] achieves a high level of generality, but one can argue it fails to be general enough in one important respect: the underlying plausibility order \leq is *always* assumed to be a total preorder which by definition implies it is *complete*, i.e., for any two worlds x, y , we have either $x \leq y$ or $y \leq x$. This implies that agent \mathcal{A} is *always* able to decide which of x, y is more plausible. This is not always realistic, and so it seems desirable to study belief removal based on plausibility orderings which allow *incomparabilities*. A little work has been done on this ([3,8,9,12,17], and especially the choice-theoretic approach to belief change advocated in [18]) but not much. This is in contrast to work in nonmonotonic

reasoning (NMR), the research area which is so often referred to as the “other side of the coin” to belief change. In NMR, semantic models based on incomplete orderings are the norm, with work dating back to the seminal papers on *preferential models* of [13,20]. Our aim in this paper is to relax the completeness assumption from [5] and to investigate the resulting, even more general class of removal functions.

The plan of the paper is as follows. In Sect. 2 we give our generalised definition of the construction from [5], which we call (*semi-modular*) *contexts*. We describe their associated removal functions, as well as mention the characterisation from [5]. Then in Sect. 3 we present an axiomatic characterisation of the family of removal functions generated by semi-modular contexts. Then, in Sect. 4 we mention a couple of further restrictions on contexts, leading to two corresponding extra postulates. In Sect. 5 we mention an important subfamily of the general family, i.e., those removals which may be generated by a finite prioritised base of *defaults*, before moving on to AGM style removal in Sect. 6. We conclude in Sect. 7.

Preliminaries: We work in a finitely-generated propositional language L . The set of non-tautologous sentences in L is denoted by L_* . The set of propositional worlds/models is W . For any set of sentences $X \subseteq L$, the set of worlds which satisfy every sentence in X is denoted by $[X]$. Classical logical consequence and equivalence are denoted by \vdash and \equiv respectively. As above, we let \mathcal{A} denote some agent whose beliefs are subject to change. A *belief set* for \mathcal{A} is represented by a single sentence which is meant to stand for all its logical consequences. A *belief removal function* (hereafter just *removal function*) belonging to \mathcal{A} is a unary function $*$ which takes any non-tautologous sentence $\lambda \in L_*$ as input and returns a new belief set $*(\lambda)$ for \mathcal{A} such that $*(\lambda) \not\vdash \lambda$. For any removal function $*$ we can always derive an associated belief set. It is just the belief set obtained by removing the contradiction, i.e., $*(\perp)$.

The following definitions about orderings will be useful in what follows. A binary relation R over W is:

- *reflexive* iff $\forall x : xRx$
- *transitive* iff $\forall x, y, z : xRy \ \& \ yRz \rightarrow xRz$
- *complete* iff $\forall x, y : xRy \vee yRx$
- a *preorder* iff it is reflexive and transitive
- a *total preorder* iff it is a complete preorder

The above notions are used generally when talking of “weak” orderings, where xRy is meant to stand for something like “ x is at least as good as y ”. However in this paper, following the lead of [17], we will find it more natural to work under a *strict* reading, where xRy denotes “ x is strictly better than y ”. In this setting, the following notions will naturally arise. R is:

- *irreflexive* iff $\forall x : \text{not}(xRx)$
- *modular* iff $\forall x, y, z : xRy \rightarrow (xRz \vee zRy)$
- a *strict partial order (spo)* iff it is both irreflexive and transitive
- the *strict part of* another relation R' iff $\forall x, y : xRy \leftrightarrow (xR'y \ \& \ \text{not}(yR'x))$
- the *converse complement* of R' iff $\forall x, y : xRy \leftrightarrow \text{not}(yR'x)$

We have that R is a modular spo iff it is the strict part of a total preorder [15]. So in terms of *strict* relations, much of the previous work on belief removal, including [5], assumes an underlying strict order which is a modular spo. It is precisely the modularity condition which we want to relax in this paper.

Given any ordering R and $x \in W$, let $\nabla_R(x) = \{z \in W \mid zRx\}$ be the set of all worlds below x in R . Then we may define a new binary relation \sqsubseteq^R from R by setting $x \sqsubseteq^R y$ iff $\nabla_R(x) \subseteq \nabla_R(y)$. That is, $x \sqsubseteq^R y$ iff every element below x in R is also below y in R . It is easy to check that if R is a modular spo then $x \sqsubseteq^R y$ iff not (yRx) , i.e., \sqsubseteq^R is just the converse complement of R .

2 Contexts, Modular Contexts and Removals

In this section we set up our generalised definition of a context, show how each such context yields a removal function and vice versa, and recap the main results from [5].

2.1 Contexts

We assume our agent \mathcal{A} has in his mind *two* binary relations ($<$, \prec) over the set W . The relation $<$ is a *strict* plausibility relation which forms the basis for \mathcal{A} 's actionable beliefs, i.e., $x < y$ means that, to \mathcal{A} 's mind, and on the basis of all available evidence, *world x is strictly more plausible than y* . We assume $<$ is a strict partial order. In addition to this there is a second binary relation \prec . This relation is open to several different interpretations, but the one we attach is as follows: $x \prec y$ means “ \mathcal{A} has **an explicit reason** to hold x more plausible than y (or to treat x more favourably than y)”. We will use \preceq to denote the converse complement of $<$, i.e., $x \preceq y$ iff $y \not< x$. Thus $x \preceq y$ iff \mathcal{A} has no reason to treat y more favourably than x . Note \preceq and \prec are interdefinable, and we find it convenient to switch between them freely.

Note the equivalence “ $x \prec y$ iff both $x \preceq y$ and $y \not\preceq x$ ” holds only if \prec is asymmetric, which might not hold in general, since it is perfectly possible for \mathcal{A} to have one explicit reason to hold x more plausible than y , and another to hold y more plausible than x . In this case both these reasons will compete with each other, with at most one of the pairs $\langle x, y \rangle$ or $\langle y, x \rangle$ making it into \mathcal{A} 's plausibility relation $<$.

What are the properties of \prec ? We assume only two things, at least to begin with: (i) an agent can never possess a reason to hold a world strictly more plausible than itself, and (ii) an agent does not hold a world x to be more plausible than another world y , i.e., $x < y$, *without* being in possession of some reason for doing so. (Note this latter property lends a certain “foundationalist” flavour to our construction.) All this is formalised in the following definition:

Definition 1. A context \mathcal{C} is a pair of binary relations ($<$, \prec) over W such that:

- (C1) $<$ is a strict partial order
- (C2) \prec is irreflexive
- (C3) $< \subseteq \prec$

If $<$ is modular then we call \mathcal{C} a *modular context*. We will later have grounds for strengthening (C3). How does \mathcal{A} use his context \mathcal{C} to construct a removal function $\ast_{\mathcal{C}}$? In terms of models, the set $[\ast_{\mathcal{C}}(\lambda)]$ of models of his new belief set, when removing a sentence λ , *must* include some $\neg\lambda$ -worlds. Following the usual practice in belief revision, he should take the most plausible ones according to $<$, i.e., the $<$ -minimal ones. But which, if any, of the λ -worlds should be included? The following principle was proposed by Rott and Pagnucco [19]:

Principle of Weak Preference

If one object is held in equal or higher regard than another, the former should be treated no worse than the latter.

Rott and Pagnucco use this principle to argue that the new set of worlds following removal should contain all worlds x which are not less plausible than a $<$ -minimal $\neg\lambda$ -world y , i.e., $y \not\prec x$. We propose to apply a tempered version of this principle using the second ordering \prec . We include x if there is *no explicit reason to believe* that y is more plausible than x , i.e., if $y \not\prec x$.

Definition 2. (\ast from \mathcal{C}) Given a context \mathcal{C} we define the removal function $\ast_{\mathcal{C}}$ by setting, for each $\lambda \in L_{\ast}$, $[\ast_{\mathcal{C}}(\lambda)] = \bigcup \{ \nabla_{\prec}(y) \mid y \in \min_{<}([\neg\lambda]) \}$.

It can be shown that different contexts give rise to different removal functions, i.e., the mapping $\mathcal{C} \mapsto \ast_{\mathcal{C}}$ is injective. The case of modular contexts was the one which was studied in detail in [5], where it was shown how, by placing various restrictions on the interaction between $<$ and \prec , this family captures a wide range of removal operations which have been previously studied, for example both AGM contraction *and* AGM revision [1]¹, severe withdrawal [19], systematic withdrawal [16] and belief liberation [4]. For the general family in that paper the following representation result was proved.

Theorem 1. [5,6] Let \mathcal{C} be a modular context. Then $\ast_{\mathcal{C}}$ satisfies the following rules:

- ($\ast 1$) $\ast(\lambda) \not\prec \lambda$
- ($\ast 2$) If $\lambda_1 \equiv \lambda_2$ then $\ast(\lambda_1) \equiv \ast(\lambda_2)$
- ($\ast 3$) If $\ast(\lambda \wedge \chi) \vdash \chi$ then $\ast(\lambda \wedge \chi \wedge \psi) \vdash \chi$
- ($\ast 4$) If $\ast(\lambda \wedge \chi) \vdash \chi$ then $\ast(\lambda \wedge \chi) \vdash \ast(\lambda)$
- ($\ast 5$) $\ast(\lambda \wedge \chi) \vdash \ast(\lambda) \vee \ast(\chi)$
- ($\ast 6$) If $\ast(\lambda \wedge \chi) \not\prec \lambda$ then $\ast(\lambda) \vdash \ast(\lambda \wedge \chi)$

Furthermore if \ast is any removal function satisfying the above 6 rules, there exists a unique modular context \mathcal{C} such that $\ast = \ast_{\mathcal{C}}$.

All these rules are familiar from the belief removal literature. ($\ast 1$) is the Success postulate while ($\ast 2$) is a syntax-irrelevance property. ($\ast 3$) is sometimes known as Conjunctive Trisection [11,17]. It says if χ is believed after removing the conjunction $\lambda \wedge \chi$, then it should also be believed when removing the longer

¹ The fact that basic removal also covers AGM revision is what motivated our choice of the contraction-revision “hybrid” symbol \ast to denote removal functions.

conjunction $\lambda \wedge \chi \wedge \psi$. Rule **(\ast 4)** is closely-related to the rule Cut from non-monotonic reasoning [13], while **(\ast 5)** and **(\ast 6)** are the two AGM supplementary postulates for contraction [1].

Note the non-appearance in this list of the AGM contraction postulates Vacuity ($\ast(\perp) \not\vdash \lambda$ implies $\ast(\lambda) \equiv \ast(\perp)$), Inclusion ($\ast(\perp) \vdash \ast(\lambda)$) and Recovery ($\ast(\lambda) \wedge \lambda \vdash \ast(\perp)$), none of which are valid in general for removal functions generated from modular contexts. Vacuity has been argued against as a general principle of belief removal in [5,6]. Inclusion has been questioned in [4], while Recovery has long been regarded as controversial (see, e.g., [10]). Nevertheless we will see in Sect. 6 how each of these three rules may be captured within our general framework.

The second part of Theorem 1 was proved using the following construction.

Definition 3. (*\mathcal{C} from \ast*) Given any removal function \ast we define the context $\mathcal{C}(\ast) = (\prec, \prec)$ as follows: $x \prec y$ iff $y \notin [\ast(\neg x \wedge \neg y)]$ and $x \prec y$ iff $y \notin [\ast(\neg x)]$.²

[5] showed that if \ast satisfies **(\ast 1)**-**(\ast 6)** then $\mathcal{C}(\ast)$ is a modular context and $\ast = \ast_{\mathcal{C}(\ast)}$.

3 Characterising the General Family

Now we want to drop the assumption that \prec is modular and assume only it is a strict partial order. How can we characterise the resulting class of removal functions? We focus first on establishing which of the postulates from Theorem 1 are sound for the general family, modifying our initial construction as and when necessary. Clearly we cannot expect that all the rules remain sound. In particular rule **(\ast 6)** is known to depend on the modularity of \prec and so might be expected to be the first to go. However we might hope to retain weaker versions of it, for instance:

(\ast 6a) If $\ast(\lambda \wedge \chi) \vdash \chi$ then $\ast(\lambda) \vdash \ast(\lambda \wedge \chi)$

(\ast 6b) $\ast(\lambda) \wedge \ast(\chi) \vdash \ast(\lambda \wedge \chi)$

These two rules appear respectively as (-8c) and (-8r) in [18] (see also [9]). **(\ast 6b)** follows from **(\ast 6)** given **(\ast 1)**.

Proposition 1. *If \mathcal{C} is a general context then $\ast_{\mathcal{C}}$ satisfies **(\ast 1)**, **(\ast 2)**, **(\ast 4)**, **(\ast 5)** and **(\ast 6a)** but not **(\ast 6b)** (hence also **(\ast 6)**) in general.*

Surprisingly, we lose **(\ast 3)**, as the following counterexample shows:

Example 1. Assume $L = \{p, q\}$ and let the 4 valuations of L be $W = \{00, 11, 01, 10\}$, where the first and second numbers denote the truth-values of p, q respectively. Let $\leq = \{(00, 10)\}$ and $\preceq = \{(10, 01)\}$ (strictly speaking the reflexive closure of this). We have $[\ast_{\mathcal{C}}(p \wedge q)] = \{00, 10, 01\}$ and $[\ast_{\mathcal{C}}(q)] = \{00\}$. Hence $10 \in [\neg q \wedge \ast_{\mathcal{C}}(p \wedge q)]$ but $10 \notin [\ast_{\mathcal{C}}(q)]$.

² When a world appears in the scope of a propositional connective, it should be understood as denoting any sentence which has that world as its only model.

This leaves us with a problem, since whereas $(\ast 6)$ is to be considered somewhat dispensible, $(\ast 3)$ is a very reasonable property for removal functions. Is there some way we can capture it? It turns out we can capture it if we strengthen the basic property $(C3)$ to:

$$(C3a) \preceq \subseteq \sqsubseteq <$$

In other words if $z < x$ and $x \preceq y$ then $z < y$. $(C3a)$ is a *coherence* condition between \prec and $<$. It is saying that if there is a world z which \mathcal{A} judges to be more plausible than x but not to y then \mathcal{A} has a reason to treat y more favourably than x . Note that for modular contexts $(C3)$ and $(C3a)$ are equivalent, but in general they are not.

Proposition 2. *If \mathcal{C} satisfies $(C3a)$ then $\ast_{\mathcal{C}}$ satisfies $(\ast 3)$.*

Thus $(C3a)$ seems necessary. Note rule $(C3a)$ may also be interpreted as a restricted form of modularity for $<$, since it may be re-written as $\forall x, y, z (z < x \rightarrow (y \prec x \vee z < y))$. For this reason we make the following definition:

Definition 4. *A semi-modular context is any context \mathcal{C} satisfying $(C3a)$.*

In the rest of the paper we will work only with semi-modular contexts. It can be shown that $\ast_{\mathcal{C}}$ still fails in general to satisfy $(\ast 6b)$ even for semi-modular contexts.

So far we have a list of sound properties for the removal functions defined from semi-modular contexts. They are the same as the rules which characterise modular removal, but with $(\ast 6)$ replaced by the weaker $(\ast 6a)$. It might be hoped that this list is complete, i.e., that *any* removal function \ast satisfying these 6 rules is equal to $\ast_{\mathcal{C}}$ for some semi-modular context \mathcal{C} . Indeed we might expect to be able to show $\ast = \ast_{\mathcal{C}(\ast)}$, where $\mathcal{C}(\ast)$ is the context defined via Definition 3. The following result gives us a good start.

Proposition 3. *Let \ast be any removal function satisfying $(\ast 1)$ - $(\ast 5)$ and $(\ast 6a)$. Then $\mathcal{C}(\ast)$ is a context, i.e., satisfies $(C1)$ - $(C3)$.*

However to get $(C3a)$ it seems an extra property is needed:

$$(\ast C) \text{ If } \ast(\lambda) \wedge \neg\lambda \vdash \ast(\chi) \wedge \neg\chi \text{ then } \ast(\lambda) \vdash \ast(\chi)$$

We can rephrase this using the *Levi Identity* [14]. Given any removal function \ast we may define a *revision function* \ast^R by setting, for each consistent sentence $\lambda \in L$, $\ast^R(\lambda) = \ast(\neg\lambda) \wedge \lambda$. Then rule $(\ast C)$ may be equivalently written as:

$$(\ast C') \text{ If } \ast^R(\neg\lambda) \vdash \ast^R(\neg\chi) \text{ then } \ast(\lambda) \vdash \ast(\chi)$$

Thus $(\ast C')$ is effectively saying that if revising by $\neg\lambda$ leads to a stronger belief set than revising by $\neg\chi$, then removing λ leads to a stronger belief set than removing χ . The next result confirms that this rule is sound for the removal functions generated by semi-modular contexts, and that this property is enough to show that $\mathcal{C}(\ast)$ satisfies $(C3a)$.

Proposition 4. *Let \mathcal{C} be a semi-modular context. Then $\ast_{\mathcal{C}}$ satisfies $(\ast\mathbf{C})$. Furthermore if \ast is any removal function satisfying $(\ast\mathbf{C})$ then the context $\mathcal{C}(\ast)$ satisfies $(\mathbf{C3a})$.*

Rule $(\ast\mathbf{C})$ is actually quite strong. In the presence of $(\ast\mathbf{3})$ it can be shown to imply $(\ast\mathbf{4})$. This means that, in the axiomatisation of $\ast_{\mathcal{C}}$ we can replace $(\ast\mathbf{4})$ with $(\ast\mathbf{C})$. To show that the list of rules is complete, it remains to prove $\ast = \ast_{\mathcal{C}(\ast)}$. It turns out that here we need the following weakening of $(\ast\mathbf{6b})$:

$$(\ast\mathbf{E}) \quad \neg(\lambda \wedge \chi) \wedge \ast(\lambda) \wedge \ast(\chi) \vdash \ast(\lambda \wedge \chi)$$

This rule may be reformulated as “ $\ast(\lambda) \wedge \ast(\chi) \vdash (\lambda \wedge \chi) \vee \ast(\lambda \wedge \chi)$ ”. In this reformulation, the right hand side of the turnstile may be thought of as standing for all those consequences of the conjunction $\lambda \wedge \chi$ which are *believed* upon its removal. The rule is saying that any such surviving consequence must be derivable from the *combination* of $\ast(\lambda)$ and $\ast(\chi)$.

Proposition 5. *Let \mathcal{C} be a semi-modular context. Then $\ast_{\mathcal{C}}$ satisfies $(\ast\mathbf{E})$.*

Theorem 2. *Let \ast be any removal function satisfying $(\ast\mathbf{1}), (\ast\mathbf{2}), (\ast\mathbf{3}), (\ast\mathbf{C}), (\ast\mathbf{5}), (\ast\mathbf{6a})$ and $(\ast\mathbf{E})$. Then $\ast_{\mathcal{C}(\ast)} = \ast$.*

Thus, to summarise, the family of removal functions defined from semi-modular contexts is completely characterised by $(\ast\mathbf{1})$ – $(\ast\mathbf{3})$, $(\ast\mathbf{C})$, $(\ast\mathbf{5})$, $(\ast\mathbf{6a})$ and $(\ast\mathbf{E})$.

4 Transitivity and Priority

In this section we look at imposing an extra couple of properties on semi-modular contexts $\mathcal{C} = (\prec, \prec)$, both of which were investigated in the case of modular contexts in [5]. There it was shown how the resulting classes of removal functions still remain general enough to include a great many of the classes which have been previously proposed in the context of modular removal.

The first property is the transitivity of \preceq , thus making \preceq a preorder. (Recall \preceq is the converse complement of \prec , so this is equivalent to making \prec modular.) According to our above interpretation of \preceq this means *if there is no reason to treat y more favourably than x , and no reason to treat z more favourably than y then there is no reason to treat z more favourably than x .*

Proposition 6. (i). *If \preceq is transitive then $\ast_{\mathcal{C}}$ satisfies the following strengthening of $(\ast\mathbf{C})$:*

$$(\ast\mathbf{C+}) \quad \text{If } \ast(\lambda) \wedge \neg\lambda \vdash \ast(\chi) \text{ then } \ast(\lambda) \vdash \ast(\chi)$$

(ii). *If \ast satisfies $(\ast\mathbf{C+})$ then the relation \preceq in $\mathcal{C}(\ast)$ is transitive.*

Note this property is a great deal simpler than the one used to characterise transitivity of \preceq in the modular context in [5]. It can be re-written as: If $\ast^R(\neg\lambda) \vdash \ast(\chi)$ then $\ast(\lambda) \vdash \ast(\chi)$. It says that if the belief set following removal of χ is contained in the belief set following the *revision* by $\neg\lambda$, then it must be contained also in the belief set following the removal of λ . This seems like a reasonable property.

Corollary 1. *For any removal function $*$, the following are equivalent:*

(i). $*$ is generated by a semi-modular context $\mathcal{C} = (\prec, \prec)$ such that \preceq is transitive. (ii). $*$ satisfies the list of rules given at the end of Sect. 3, with $(*\mathbf{C})$ replaced by $(*\mathbf{C}+)$.

Now consider the following property of a context $\mathcal{C} = (\prec, \prec)$:

(**CP**) If $x \prec y$ and $y \not\prec x$ then $x < y$

This, too, looks reasonable: if \mathcal{A} has an explicit reason to hold x more plausible than y , but not vice versa, then in the final reckoning he should hold x to be strictly more plausible than y . Consider the following property of removal functions:

(***P**) If $*(\lambda) \vdash \chi$ and $*(\chi) \not\vdash \lambda$ then $*(\lambda \wedge \chi) \vdash \chi$

This property is briefly mentioned as *Priority* in [3], and is also briefly mentioned right at the end of [7]. It can be read as saying that if λ is excluded following removal of χ , but not vice versa, then χ is strictly more entrenched than λ . **For the case of modular removal**, we can obtain the following exact correspondence between (**CP**) and (***P**):

Proposition 7. (i). *If \mathcal{C} is a modular context satisfying (**CP**) then $*_{\mathcal{C}}$ satisfies (***P**).* (ii). *If $*$ satisfies (***P**) then $\mathcal{C}(*_{\mathcal{C}})$ satisfies (**CP**).*

The proof of Proposition 7(i) makes critical use of the modularity of $<$. It turns out that (***P**) is *not* sound for general semi-modular contexts, even if we insist on (**CP**).

Example 2. Suppose $L = \{p, q\}$ and that $< = \{(01, 11)\}$ while $\preceq = \{(01, 11)\}$ (strictly speaking the reflexive closure of this). One can verify that \mathcal{C} is a semi-modular context and that (**CP**) is satisfied. Now let $\lambda = p \vee \neg q$ and $\chi = \neg p$. Then $[*_{\mathcal{C}}(\lambda)] = \{01\}$, $[*_{\mathcal{C}}(\chi)] = \{11, 01, 10\}$ and $[*_{\mathcal{C}}(\lambda \wedge \chi)] = \{01, 10\}$ and we have $*_{\mathcal{C}}(\lambda) \vdash \chi$, $*_{\mathcal{C}}(\chi) \not\vdash \lambda$, and $*_{\mathcal{C}}(\lambda \wedge \chi) \not\vdash \chi$. Hence (***P**) is not satisfied.

The question now is, which postulate corresponds to (**CP**) for general semi-modular contexts? Here is the answer:

Proposition 8. (i). *If \mathcal{C} is a semi-modular context which satisfies (**CP**), then $*_{\mathcal{C}}$ satisfies the following rule:*

(***P'**) *If $*(\lambda) \vdash \chi$ and $*(\chi) \vdash *(\lambda \wedge \chi)$ then $*(\chi) \vdash \lambda$*

(ii). *If $*$ satisfies (***P'**), plus $(*\mathbf{C})$ and $(*\mathbf{1})$, then $\mathcal{C}(*_{\mathcal{C}})$ satisfies (**CP**).*

It is straightforward to see (***P'**) is weaker than (***P**) given (***1**), while it implies (***P**) given (***6**).

5 Finite Base-Generated Removal

In this section we mention a concrete and important subfamily of our general family of removal functions, the ideas behind which can be seen already throughout the literature on nonmonotonic reasoning and belief change (see in particular [3] for a general treatment in a belief removal context). Given any, possibly inconsistent, set Σ of sentences, let $\text{cons}(\Sigma)$ denote the set of all consistent subsets of Σ . We assume agent \mathcal{A} is in possession of a finite set Σ of sentences which are possible *assumptions* or *defaults*, together with a strict preference ordering \Subset on $\text{cons}(\Sigma)$ (with sets “higher” in the ordering assumed more preferred). We assume the following two properties of \Subset :

- (**Σ 1**) \Subset is a strict partial order
- (**Σ 2**) If $A \subset B$ then $A \Subset B$

(**Σ 2**) is a monotonicity requirement stating a given set of defaults is strictly preferred to all its proper subsets.

Definition 5. *If $\Sigma \subseteq L$ is a finite set of sentences and \Subset is a binary relation over $\text{cons}(\Sigma)$ satisfying (**Σ 1**) and (**Σ 2**). Then we call $\mathbb{S} = \langle \Sigma, \Subset \rangle$ a prioritised default base. If in addition \Subset is modular then we call \mathbb{S} a modular prioritised default base.*

How does the agent use a prioritised default base $\mathbb{S} = \langle \Sigma, \Subset \rangle$ to remove beliefs? For $\Sigma \subseteq L$ and $\lambda \in L_*$ let $\text{cons}(\Sigma, \lambda) \stackrel{\text{def}}{=} \{S \in \text{cons}(\Sigma) \mid S \not\vdash \lambda\}$. Then from \mathbb{S} we may define a removal function $*_{\mathbb{S}}$ by setting, for each $\lambda \in L_*$,

$$*_{\mathbb{S}}(\lambda) = \bigvee \left\{ \bigwedge S \mid S \in \max_{\Subset} \text{cons}(\Sigma, \lambda) \right\}.$$

In other words, after removing λ , \mathcal{A} will believe precisely those sentences which are consequences of *all maximally preferred* subsets of Σ which do not imply λ .

We will now show how the family of removal functions generated from prioritised default bases fits into our general family. From a given $\mathbb{S} = \langle \Sigma, \Subset \rangle$ we may define a context $\mathcal{C}(\mathbb{S}) = (\prec, \prec)$ as follows. Let $\text{sent}_{\Sigma}(x) \stackrel{\text{def}}{=} \{\alpha \in \Sigma \mid x \in [\alpha]\}$. Then

- $x \prec y$ iff $\text{sent}_{\Sigma}(y) \Subset \text{sent}_{\Sigma}(x)$
- $x \prec y$ iff $\text{sent}_{\Sigma}(x) \not\subseteq \text{sent}_{\Sigma}(y)$

Thus we define x to be more plausible than y iff the set of sentences in Σ satisfied by x is more preferred than the set of sentences in Σ satisfied by y . Meanwhile we have the natural interpretation for \prec that \mathcal{A} has a reason to hold x to be more plausible than y precisely when one of the sentences in Σ is satisfied by x but not y .

Theorem 3. (i). $\mathcal{C}(\mathbb{S})$ defined above forms a semi-modular context (which is modular if \Subset is modular). (ii). \preceq is transitive and the condition (**CP**) from Sect. 4 holds. (iii). $*_{\mathbb{S}} = *_{\mathcal{C}(\mathbb{S})}$.

Thus we have shown that every removal function generated by a prioritised default base may *always* be generated by a semi-modular context which furthermore satisfies the two conditions on contexts mentioned in the previous section. By the results of the previous sections, this means we automatically obtain a list of sound postulates for the default base-generated removals.

Corollary 2. *Let Σ be any prioritised default base. Then \ast_{Σ} satisfies all the rules listed at the end of Sect. 3, as well as $(\ast\mathbf{C}+)$ and $(\ast\mathbf{P}')$ from the last section.*

Note we have shown how every prioritised default base gives rise to a semi-modular context satisfying \preceq -transitivity and (\mathbf{CP}) . An open question is whether *every* such context arises in this way.

6 AGM Preferential Removal

Recall that three of the basic AGM postulates for contraction do not hold in general for the removal functions generated by semi-modular contexts, namely Inclusion, Recovery and Vacuity. In this section we show how each of these rules can be captured. In [5] it was shown already how they may be captured within the class of modular context-generated removal.

The Inclusion rule is written in our setting as follows:

$$(\ast\mathbf{I}) \ast(\perp) \vdash \ast(\lambda)$$

To capture $(\ast\mathbf{I})$ for any removal generated from any semi-modular context $\mathcal{C} = (<, \prec)$, we need only to require the following condition on \mathcal{C} :

$$(\mathbf{CI}) \min_{<}(W) \subseteq \min_{\prec}(W)$$

According to our interpretation of \prec , (\mathbf{CI}) is stating that, for any world x , if \mathcal{A} has some explicit reason favour some world y over x (i.e., $y \prec x$) then in the final reckoning \mathcal{A} must hold *some* world z (not necessarily the same as y) more plausible than x (i.e., $z < x$).

Proposition 9. (i). *If \mathcal{C} satisfies (\mathbf{CI}) then $\ast_{\mathcal{C}}$ satisfies $(\ast\mathbf{I})$.* (ii). *If \ast satisfies $(\ast\mathbf{I})$ then $\mathcal{C}(\ast)$ satisfies (\mathbf{CI}) .*

The Recovery rule is written as follows:

$$(\ast\mathbf{R}) \ast(\lambda) \wedge \lambda \vdash \ast(\perp)$$

The corresponding property on contexts $\mathcal{C} = (<, \prec)$ is:

$$(\mathbf{CR}) \text{ If } y \notin \min_{<}(W) \text{ and } x \neq y \text{ then } x \prec y$$

Thus the only worlds $\nabla_{\preceq}(x)$ contains, other than x itself, are worlds in $\min_{<}(W)$.

Proposition 10. (i). *If \mathcal{C} satisfies (\mathbf{CR}) then $\ast_{\mathcal{C}}$ satisfies $(\ast\mathbf{R})$.* (ii). *If \ast satisfies $(\ast\mathbf{R})$ then $\mathcal{C}(\ast)$ satisfies (\mathbf{CR}) .*

Note the combination of **(CI)** and **(CR)** specifies \prec , equivalently \preceq , uniquely in terms of $<$, viz. $x \preceq_{agm} y$ iff $x = y$ or $x \in \min_{<}(W)$, and we obtain the removal recipe of AGM contraction, in which removal of λ boils down to just adding the $<$ -minimal $\neg\lambda$ -worlds to the $<$ -minimal worlds:

$$[*_{agm}(\lambda)] = \min_{<}(W) \cup \min_{<}([\neg\lambda]).$$

It is easy to check that the resulting context \mathcal{C} satisfies condition **(C3a)** and thus forms a semi-modular context. It is also easy to check **(CP)** is satisfied and that the above-defined \preceq_{agm} is transitive. Thus the above $*_{agm}$ also satisfies **(*C+)** and **(*P')** from Sect. 4. It can also be shown to satisfy **(*6b)**.

The Vacuity rule is written as follows:

(*V) If $*(\perp) \not\models \lambda$ then $*(\lambda) \equiv *(\perp)$

Unlike in the modular case, where Vacuity is known to follow from Inclusion for modular removal functions [5], **(*V)** does not even hold in general for the above preferential AGM contraction $*_{agm}$. This was essentially noticed, in a revision context, in [2].

Example 3. Let $L = \{p, q\}$ and $< = \{(11, 01)\}$. So $[\ast_{agm}(\perp)] = \{00, 11, 10\}$. Let $\lambda = p$. Then we have $\ast_{agm}(\perp) \not\models \lambda$ (because $00 \in [\ast_{agm}(\perp)]$), but $\min_{<}([\neg\lambda]) = \{00, 01\}$, so $[\ast_{agm}(\lambda)] = \min_{<}(W) \cup \min_{<}([\neg\lambda]) = W \neq [\ast_{agm}(\perp)]$.

In order to ensure $*_{agm}$ satisfies **(*V)** it is necessary, as is done in [12], to enforce the following property on $<$.

(<V) $\forall x, y ((x \in \min_{<}(W) \wedge y \notin \min_{<}(W)) \rightarrow x < y)$.

In other words all $<$ -minimal worlds can be compared with, and are below, every world which is not $<$ -minimal. For general semi-modular contexts $\mathcal{C} = (<, \prec)$ we also require the following condition, which is weaker than **(CI)**:

(CV) If $x, y \in \min_{<}(W)$ then $x \not\prec y$

This property says that for any two of his $<$ -minimal worlds, \mathcal{A} will not have explicit reason to hold one to be more plausible than the other.

Proposition 11. (i). If \mathcal{C} satisfies **(CV)** and **(<V)** then $*_{\mathcal{C}}$ satisfies **(*V)**.
(ii). If $*$ satisfies **(*V)** then $\mathcal{C}(\ast)$ satisfies **(CV)**.

7 Conclusion

In this paper we introduced a family of removal functions, generalising the one given in [5] to allow for incomparabilities in the plausibility relation $<$ between possible worlds. Removal is carried out using the plausibility relation in combination with a second relation \prec which can be thought of as indicating “reasons” for holding one world to be more plausible than another. We axiomatically characterised this general family as well as certain subclasses, and we showed how this

family includes some important and natural families of belief removal, specifically those which may be generated from prioritised default bases and the preferential counterpart of AGM contraction. Our results show the central construct used in this paper, i.e., semi-modular contexts, to be a very useful tool in the study of belief removal functions.

For future work we would like to employ semi-modular contexts in the setting of *social belief removal* [6], in which there are several agents, each assumed to have their own removal function, and in which all agents must remove some belief to become consistent with each other. [6] showed that, under the assumption that each agent uses a removal function generated from a *modular* context, certain *equilibrium points* in the social removal process are guaranteed to exist. An interesting question would be whether these results generalise to the *semi-modular* case.

Acknowledgements

Thanks are due to two anonymous reviewers for some helpful comments.

References

1. Alchourrón, C., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50(2), 510–530 (1985)
2. Benferhat, S., Lagrue, S., Papini, O.: Revision of partially ordered information: Axiomatization, semantics and iteration. In: Pack Kaelbling, L., Saffiotti, A. (eds.) *IJCAI*, pp. 376–381. Professional Book Center (2005)
3. Bochman, A.: *A Logical Theory of Nonmonotonic Inference and Belief Change*. Springer, Heidelberg (2001)
4. Booth, R., Chopra, S., Ghose, A., Meyer, T.: Belief liberation (and retraction). *Studia Logica* 79(1), 47–72 (2005)
5. Booth, R., Chopra, S., Meyer, T., Ghose, A.: A unifying semantics for belief change. In: *Proceedings of ECAI 2004*, pp. 793–797 (2004)
6. Booth, R., Meyer, T.: Equilibria in social belief removal. In: *KR*, pp. 145–155 (2008)
7. Cantwell, J.: Relevant contraction. In: *Proceedings of the Dutch-German Workshop on Non-Monotonic Reasoning, DGNMR 1999* (1999)
8. Cantwell, J.: Eligible contraction. *Studia Logica* 73, 167–182 (2003)
9. Arló Costa, H.: Rationality and value: The epistemological role of indeterminate and agent-dependent values. *Philosophical Studies* 128(1), 7–48 (2006)
10. Hansson, S.O.: Belief contraction without recovery. *Studia Logica* 50(2), 251–260 (1991)
11. Hansson, S.O.: Changes on disjunctively closed bases. *Journal of Logic, Language and Information* 2, 255–284 (1993)
12. Katsuno, H., Mendelzon, A.O.: Propositional knowledge base revision and minimal change. *Artif. Intell.* 52(3), 263–294 (1992)
13. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44, 167–207 (1991)

14. Levi, I.: *The Fixation of Belief and Its Undoing*. Cambridge University Press, Cambridge (1991)
15. Maynard-Zhang, P., Lehmann, D.: Representing and aggregating conflicting beliefs. *Journal of Artificial Intelligence Research* 19, 155–203 (2003)
16. Meyer, T., Heidema, J., Labuschagne, W., Leenen, L.: Systematic withdrawal. *Journal of Philosophical Logic* 31(5), 415–443 (2002)
17. Rott, H.: Preferential belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information* 1, 45–78 (1992)
18. Rott, H.: *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford University Press, Oxford (2001)
19. Rott, H., Pagnucco, M.: Severe withdrawal (and recovery). *Journal of Philosophical Logic* 28, 501–547 (1999)
20. Shoham, Y.: A semantical approach to nonmonotonic logics. In: *LICS*, pp. 275–279 (1987)