On the Logic of Iterated Non-prioritised Revision

Richard Booth

Department of Computer Science, University of Leipzig,
Augustusplatz 10/11, 04109 Leipzig, Germany
booth@informatik.uni-leipzig.de

Abstract. We look at iterated non-prioritised belief revision, using as a starting point a model of non-prioritised revision, similar to Makinson's *screened revision*, which assumes an agent keeps a set of *core* beliefs whose function is to block certain revision inputs. We study postulates for the iteration of this operation. These postulates generalise some of those which have previously been proposed for iterated AGM ("prioritised") revision, including those of Darwiche and Pearl. We then add a second type of revision operation which allows the core itself to be revised. Postulates for the iteration of this operator are also provided, as are rules governing mixed sequences of revisions consisting of both regular and core inputs. Finally we give a construction of both a regular and core revision operator based on an agent's *revision history.* This construction is shown to satisfy most of the postulates.

1 Introduction and Preliminaries

The most popular basic framework for the study of belief revision has been the one due to Alchourrón, Gärdenfors and Makinson (AGM) [1, 10]. This framework has been subjected in more recent years to several different extensions and refinements. Two of the most interesting of these have been the study of so-called *non-prioritised* revision [2, 12, 13, 19], i.e., revision in which the input sentence is not necessarily accepted, and of *iterated* revision [3, 5, 7, 8, 17, 21], i.e., the study of the behaviour of an agent's beliefs under a *sequence* of revision inputs. However, most of the extensions in the former group are concerned only with single-step revision. Similarly, most of the contributions to the area of iterated AGM revision are in the setting of normal, "prioritised" revision in which the input sentence is always accepted. However the question of *iterated non-prioritised revision* is certainly an interesting one, as can be seen from the following example.[1]

Example 1. Your six-year-old son comes home from school and tells you that today he had lunch at school with King Gustav. Given your expectations of the

[1] Based on an example given in [9–Ch. 7] to illustrate non-prioritised revision.

G. Kern-Isberner, W. Rödder, and F. Kulmann (Eds.): WCII 2002, LNAI 3301, pp. 86–107, 2005.

King's lunching habits, you dismiss this information as a product of your son's imagination, i.e., you reject this information. But then you switch on the TV news and see a report that King Gustav today made a surprise visit to a local school. Given *this* information, your son's story doesn't seem quite so incredible as it did. Do you now believe your son's information?

What this example seems to show is that information which is initially rejected (such as your son's information) may still have an influence on the results of *subsequent* revisions. In particular if subsequent information lends support to it, then this could cause a *re-evaluation* of the decision to reject, possibly even leading the input to be accepted *retrospectively.* The main purpose of this paper is to study patterns of iterated non-prioritised revision such as these.

We will use as a starting point one particular model of non-prioritised revision, the idea behind which first appeared behind Makinson's *screened revision* [19], and then again as a special case of Hansson et al.'s *credibility-limited revision* [13]. It is that an agent keeps, as a subset of his beliefs, a set of *core beliefs* which he considers "untouchable". This set of core beliefs then acts as the determiner as to whether a given revision input is accepted or not: if an input ϕ is consistent with the core beliefs then the agent accepts the input and revises his belief set by ϕ using a normal AGM revision operator. On the other hand if ϕ contradicts the core beliefs then the agent rejects ϕ rather than give up any of the core beliefs. In this case his belief set is left undisturbed. We will see that this quite simple model will already give us a flavour of some of the interesting issues at stake. For a start, to be able to iterate this operator we need to say not only what the new belief set is after a revision, but also what the new *core* belief set is.

The explicit inclusion in an agent's epistemic state of a second set of beliefs to represent the agent's core beliefs invites the question of what would happen if this set too were to be subject to revision by external inputs, just like the normal belief set. This question will also be taken up in this paper. Thus we will have *two* different types of revision operator existing side-by-side: the usual operators described above, which we shall call *regular revision* operators, and *core revision* operators. Again both single-step and iterated core revision will be looked at. We also look at the particularly interesting possibility of performing *mixed* sequences of revisions consisting of both regular and core revisions.

The plan of the paper is as follows. We start in Sect. 2 by briefly describing the revision operators of AGM and introducing our primitive notions of *epistemic state* and *epistemic frame.* Then, in Sect. 3, we look at regular revision. We consider postulates for both the single-step and the iterated case. The latter will involve adapting some well-known postulates from the literature on iterated AGM revision – principally those proposed by Darwiche and Pearl [8] – to our non-prioritised situation. In Sect. 4 we look at some postulates for single-step and iterated core revision. Some possible rules for mixed sequences of revision inputs will be looked at in Sect. 5. Then in Sect. 6, using a particular representation of an agent's epistemic state, we provide a construction of both a regular and a core revision operator. These operators are shown to display most of the behaviour described by our postulates. We conclude in Sect. 7.

1.1 Preliminaries

We assume a propositional language generated from finitely many propositional variables. Let L denote the set of sentences of this language. Cn denotes the classical logical consequence operator. We write $Cn(\theta)$ rather than $Cn(\{\theta\})$ for $\theta \in L$ and use L^+ to denote the set of all classically consistent sentences. Formally, a *belief set* will be any set of sentences $K \subseteq L$ which is (i) consistent, i.e., $Cn(K) \neq L$, and (ii) deductively closed, i.e., $K = Cn(K)$. We denote the set of all belief sets by $\mathcal{K}$. Given $K \in \mathcal{K}$ and $\phi \in L$, we let $K + \phi$ denote the *expansion* of K by ϕ, i.e., $K + \phi = Cn(K \cup \{\phi\})$.

We let $\mathcal{W}$ denote the set of propositional worlds associated to L, i.e., the set of truth-assignments to the propositional variables in L. For any set $X \subseteq L$ of sentences we denote by $[X]$ the set of worlds in $\mathcal{W}$ which satisfy all the sentences in X (writing $[\phi]$ rather than $[\{\phi\}]$ for the case of singletons). Given a set $S \subseteq \mathcal{W}$ of worlds we write $Th(S)$ to denote the set of sentences in L which are satisfied by all the worlds in S. A *total pre-order* on $\mathcal{W}$ is any binary relation $\leq$ on $\mathcal{W}$ which is reflexive, transitive and connected (for all $w_1, w_2 \in \mathcal{W}$ either $w_1 \leq w_2$ or $w_2 \leq w_1$). For each such order $\leq$ we let $<$ denote its strict part and $\sim$ denote its symmetric part, i.e., we have $w_1 < w_2$ iff both $w_1 \leq w_2$ and $w_2 \not\leq w_1$, and $w_1 \sim w_2$ iff both $w_1 \leq w_2$ and $w_2 \leq w_1$. Given a total pre-order $\leq$ on $\mathcal{W}$ and given $S \subseteq \mathcal{W}$ we will use $\min(S, \leq)$ to denote the set of worlds which are minimal in S under $\leq$, i.e., $\min(S, \leq) = \{w \in S \mid w \leq w' \text{ for all } w' \in S\}$. We will say that a total pre-order $\leq$ on $\mathcal{W}$ is *anchored on* S if S contains precisely the minimal elements of $\mathcal{W}$ under $\leq$, i.e., if $S = \min(\mathcal{W}, \leq)$.

2 AGM and Epistemic Frames

The original AGM theory of revision is a theory about how to revise a fixed generic belief set K by any given sentence. In this paper we simplify by assuming all revision input sentences are consistent. (For this reason the usual, but for us vacuous, pre-condition "if ϕ is consistent" is absent from our formulation of AGM postulate **(K*5)** below.) At the centre of this theory is the list of *AGM revision postulates (relative to K)* which seek to rationally constraint the outcome of such a revision. Using $K * \phi$ as usual to denote the result of revising K by $\phi \in L^+$, the full list of these postulates is:

(K*1) $K * \phi = Cn(K * \phi)$
(K*2) If $\phi_1 \leftrightarrow \phi_2 \in Cn(\emptyset)$ then $K * \phi_1 = K * \phi_2$
(K*3) $\phi \in K * \phi$
(K*4) If $\neg\phi \notin K$ then $K * \phi = K + \phi$
(K*5) $K * \phi$ is consistent
(K*6) If $\neg\phi \notin K * \theta$ then $K * (\theta \wedge \phi) = (K * \theta) + \phi$

Note the presence of **(K*3)** – the "success" postulate – which says that the input sentence is always accepted. For a given belief set K, we shall call any function $*$ which satisfies the above postulates a *simple AGM revision function for K*. It is well-known that requiring these postulates to hold is equivalent to requiring that, when performing an operation of revision on his belief set K, an agent acts as though he has a total pre-order $\leq$ on the set of worlds $\mathcal{W}$ representing some subjective assessment of their relative *plausibility*, with the worlds in $[K]$ being the most plausible, i.e., $\leq$-minimal. Given the input sentence ϕ, the agent then takes as his new belief set the set of sentences true in all the most plausible worlds satisfying ϕ. Precisely we have:

Theorem 1 ([11, 15]). *Let $K \in \mathcal{K}$ and $*$ be an operator which, for each $\phi \in L^+$, returns a new set of sentences $K * \phi$. Then $*$ is a simple AGM revision function for K iff there exists some total pre-order $\leq$ on $\mathcal{W}$, anchored on $[K]$, such that, for all $\phi \in L^+$, $K * \phi = Th(\min([\phi], \leq))$.*

In this paper we will make extensive use of the above equivalence.

2.1 Epistemic Frames

One of the morals of the work already done on attempting to extend the AGM framework to cover iterated revision (see, e.g. [8, 14, 18, 21]) is that, in order to be able to formally say anything interesting about iterated revision, it is necessary to move away from the AGM representation of an agent's epistemic state as a simple belief set, and instead assume that revision is carried out on some more comprehensive object of which the belief set is but one ingredient. We will initially follow [8] in taking an abstract view of epistemic states. As in that paper, we assume a set Ep of epistemic states as primitive and assume that from each such state $\mathbb{E} \in Ep$ we can extract a belief set $\triangle(\mathbb{E})$ representing the agent's regular beliefs in $\mathbb{E}$. Unlike in [8] however, we also explicitly assume that we can extract a second belief set $\blacktriangle(\mathbb{E}) \subseteq \triangle(\mathbb{E})$ representing the agent's *core* beliefs in $\mathbb{E}$. This is all captured by the definition of an *epistemic frame*:

Definition 1. *An* epistemic frame *is a triple $\langle Ep, \triangle, \blacktriangle \rangle$, where Ep is a set, whose elements will be called* epistemic states*, and $\triangle : Ep \rightarrow \mathcal{K}$ and $\blacktriangle : Ep \rightarrow \mathcal{K}$ are functions such that, for all $\mathbb{E} \in Ep$, $\blacktriangle(\mathbb{E}) \subseteq \triangle(\mathbb{E})$.*

For most of this paper we will assume that we are working with some arbitrary, but fixed, epistemic frame $\langle Ep, \triangle, \blacktriangle \rangle$ in the background. Not until Sect. 6 will we get more specific and employ a more concrete representation of an epistemic frame. An obvious fact which is worth keeping in mind is that, since $\blacktriangle(\mathbb{E}) \subseteq \triangle(\mathbb{E})$, we always have $[\triangle(\mathbb{E})] \subseteq [\blacktriangle(\mathbb{E})]$. The set $\blacktriangle(\mathbb{E})$ can in general be *any* sub(belief)set of $\triangle(\mathbb{E})$. As two special cases, at opposite extremes, we have $\blacktriangle(\mathbb{E}) = Cn(\emptyset)$, i.e., the only core beliefs are the tautologies, and $\blacktriangle(\mathbb{E}) = \triangle(\mathbb{E})$, i.e., *all* regular beliefs are also core beliefs. One of our main aims in this paper will be to try and formulate rational constraints on the behaviour of both the regular beliefs $\triangle(\mathbb{E})$ **and** the core beliefs $\blacktriangle(\mathbb{E})$ under operations of change **to the underlying epistemic state** $\mathbb{E}$. We begin with the case when the operation of change is triggered by a regular belief input.

3 Regular Revision Inputs

In this section we consider the usual case where the revision input is a (consistent) sentence to be included in the regular belief set $\triangle(\mathbb{E})$. Given an epistemic state $\mathbb{E} \in Ep$ and a regular input $\phi \in L^+$, we shall let $\mathbb{E} \circ \phi$ denote the resulting epistemic state. We consider the single-step case and the iterated case in turn.

3.1 Single-Step Regular Revision

As indicated in the introduction, we follow the spirit of screened revision and assume that the new regular belief set $\triangle(\mathbb{E} \circ \phi)$ is given by

$$\triangle(\mathbb{E} \circ \phi) = \begin{cases} \triangle(\mathbb{E}) *^{\mathbb{E}}_{\triangle} \phi & \text{if } \neg\phi \notin \blacktriangle(\mathbb{E}) \\ \triangle(\mathbb{E}) \text{ otherwise.} \end{cases}$$

where, for each epistemic state $\mathbb{E}$, $*^{\mathbb{E}}_{\triangle}$ is a simple AGM revision function for $\triangle(\mathbb{E})$. This is also very similar to the definition of *endorsed core beliefs revision* in [13]. The difference is that in that paper the function $*^{\mathbb{E}}_{\triangle}$ is not assumed to satisfy the postulate **(K*6)** from Sect. 2. By Theorem 1, the above method is equivalent to assuming that for each $\mathbb{E}$ there exists some total pre-order $\leq^{\mathbb{E}}_{\triangle}$ on $\mathcal{W}$, anchored on $[\triangle(\mathbb{E})]$, such that

$$\triangle(\mathbb{E} \circ \phi) = \begin{cases} Th(\min([\phi], \leq^{\mathbb{E}}_{\triangle})) \text{ if } \neg\phi \notin \blacktriangle(\mathbb{E}) \\ \triangle(\mathbb{E}) \text{ otherwise.} \end{cases} \tag{1}$$

We remark that the subscript on $\leq^{\mathbb{E}}_{\triangle}$ does not actually denote the $\triangle$-function itself, but is merely a decoration to remind us that this order is being used to revise the *regular* beliefs in $\mathbb{E}$. We now make the following definition:

Definition 2. *Let $\circ : Ep \times L^+ \rightarrow Ep$ be a function. Then $\circ$ is a* regular revision operator (on the epistemic frame $\langle Ep, \triangle, \blacktriangle \rangle$) *if, for each $\mathbb{E} \in Ep$, there exists a total pre-order $\leq^{\mathbb{E}}_{\triangle}$ on $\mathcal{W}$, anchored on $[\triangle(\mathbb{E})]$, such that $\triangle(\mathbb{E} \circ \phi)$ may be determined as in (1) above. We call $\leq^{\mathbb{E}}_{\triangle}$ the* regular pre-order associated to $\mathbb{E}$ (according to $\circ$).

For some properties satisfied by this general type of construction the reader is referred to [13, 19]. One intuitive property which is not guaranteed to hold under the above definition as it stands[2] is the following, which essentially corresponds to the rule (Strong Regularity) from [13]:

(SR) $\blacktriangle(\mathbb{E}) \subseteq \triangle(\mathbb{E} \circ \phi)$

[2] For a counter-example suppose $\mathbb{E}$ is such that $\triangle(\mathbb{E}) = Cn(p)$ and $\blacktriangle(\mathbb{E}) = Cn(p \vee q)$ where p, q are distinct propositional variables, and suppose $*^{\mathbb{E}}_{\triangle}$ is the "trivial" simple AGM revision function for $\triangle(\mathbb{E})$ given by $\triangle(\mathbb{E}) *^{\mathbb{E}}_{\triangle} \phi = \triangle(\mathbb{E}) + \phi$ if $\neg\phi \notin \triangle(\mathbb{E})$, $\triangle(\mathbb{E}) *^{\mathbb{E}}_{\triangle} \phi = Cn(\phi)$ otherwise. Then $\blacktriangle(\mathbb{E}) \nsubseteq \triangle(\mathbb{E} \circ \neg p) = Cn(\neg p)$.

This postulate states that the set of core beliefs are retained as regular beliefs after revision, while leaving open the question of whether they are again retained as *core* beliefs. For this property to hold of a regular revision operator $\circ$ we require that the pre-orders $\leq^{\mathbb{E}}_{\triangle}$ associated with each $\mathbb{E}$ satisfy an extra condition, namely that $\leq^{\mathbb{E}}_{\triangle}$ considers all the worlds in $[\blacktriangle(\mathbb{E})]$ as strictly more plausible than all the worlds *not* in $[\blacktriangle(\mathbb{E})]$.[3]

Proposition 1. *Let $\circ$ be a regular revision operator. Then $\circ$ satisfies* **(SR)** *iff, for each $\mathbb{E} \in Ep$, $\leq^{\mathbb{E}}_{\triangle}$ satisfies $w_1 <^{\mathbb{E}}_{\triangle} w_2$ whenever $w_1 \in [\blacktriangle(\mathbb{E})]$ and $w_2 \notin [\blacktriangle(\mathbb{E})]$.*

The reader may have noticed that, since inputs which contradict $\blacktriangle(\mathbb{E})$ are simply rejected, the $\leq^{\mathbb{E}}_{\triangle}$-ordering of the worlds outside of $[\blacktriangle(\mathbb{E})]$ never plays any role in determining the new regular belief set. It *will*, however, play a role later on when we come to look at core revision.

What effect should performing a regular revision $\circ$ have on the core belief set? In this paper we take the position that $\circ$ is concerned exclusively with changes to $\triangle(\mathbb{E})$, and so the core belief set does not change at all.

(X1) $\blacktriangle(\mathbb{E} \circ \phi) \subseteq \blacktriangle(\mathbb{E})$ (Core Non-expansion)
(X2) $\blacktriangle(\mathbb{E}) \subseteq \blacktriangle(\mathbb{E} \circ \phi)$ (Core Preservation)

Thus **(X1)**, respectively **(X2)**, says that no core beliefs are added, respectively lost, during an operation of regular revision.

Definition 3. *Let $\circ$ be a regular revision operator. Then $\circ$ is* core-invariant *iff $\circ$ satisfies both* **(X1)** *and* **(X2)**.

Since clearly **(X2)** implies **(SR)**, we have that every core-invariant regular revision operator satisfies **(SR)**. The reasonableness of core-invariance may be questioned. For example a consequence of **(X2)** is that we automatically get that if $\neg\phi \in \blacktriangle(\mathbb{E})$ then $\phi \notin \triangle(\mathbb{E} \circ \phi \circ \phi \circ \cdots \circ \phi)$, and this holds regardless of how many times we revise by ϕ, be it one or one billion. It might be expected here that repeatedly receiving ϕ might have the effect of gradually "loosening" $\neg\phi$ from the core beliefs until eventually at some point it "falls out", leading ϕ to become acceptable. Similarly, rule **(X1)** precludes the situation in which repeated input of a non-core belief eventually leads to the admittance of that belief into the core. On the other hand there exist situations in which core-invariance does seem reasonable in these cases. An example is when the regular belief inputs are assumed to be coming from a single source throughout, i.e., the source is just repeating itself. Weaker alternatives to the rules **(X1)** and **(X2)** which come to mind are:

(wX1) $\blacktriangle(\mathbb{E} \circ \phi) \subseteq \blacktriangle(\mathbb{E}) + \phi$ (Weak Core Non-expansion)
(wX2) $\blacktriangle(\mathbb{E}) \subseteq \blacktriangle(\mathbb{E} \circ \phi) + \neg\phi$ (Weak Core Preservation)

In terms of propositional worlds, **(wX1)** is equivalent to requiring $[\blacktriangle(\mathbb{E})] \cap [\phi] \subseteq [\blacktriangle(\mathbb{E} \circ \phi)]$, while **(wX2)**, which is reminiscent of the "recovery" postulate

[3] Due to space limitations, proofs are omitted from this version of the paper.

from belief contraction [10], is equivalent to requiring $[\blacktriangle(\mathbb{E} \circ \phi)] \subseteq [\blacktriangle(\mathbb{E})] \cup [\phi]$. Thus **(wX1)** says that, in the transformation of $[\blacktriangle(\mathbb{E})]$ into $[\blacktriangle(\mathbb{E} \circ \phi)]$, the only worlds which can possibly be *removed* from $[\blacktriangle(\mathbb{E})]$ are those in $[\neg\phi]$, while **(wX2)** says that the only worlds which can possibly be *added* are those in $[\phi]$. For this paper, however, we will assume that both **(X1)** and **(X2)** hold throughout, and so we will make no further reference to the above weaker versions.

3.2 Iterating Regular Revision

Now we consider iteration of $\circ$. How should $\blacktriangle(\mathbb{E})$ and $\triangle(\mathbb{E})$ behave under sequences of regular inputs? Clearly since we are accepting both **(X1)** and **(X2)** this question is already answered in the case of $\blacktriangle(\mathbb{E})$ — the core beliefs remain constant throughout. What about the regular beliefs $\triangle(\mathbb{E})$? Here we take our lead from the work on iterated AGM ("prioritised") revision by Darwiche and Pearl [8]. They suggest a list of four postulates to rationally constrain the beliefs under iterated AGM revision (we will write "$\mathbb{E} \circ \theta \circ \phi$" rather than "$(\mathbb{E} \circ \theta) \circ \phi$" etc.):

$\mathbf{(C1)}_\triangle$ If $\phi \to \theta \in Cn(\emptyset)$ then $\triangle(\mathbb{E} \circ \theta \circ \phi) = \triangle(\mathbb{E} \circ \phi)$
$\mathbf{(C2)}_\triangle$ If $\phi \to \neg\theta \in Cn(\emptyset)$ then $\triangle(\mathbb{E} \circ \theta \circ \phi) = \triangle(\mathbb{E} \circ \phi)$
$\mathbf{(C3)}_\triangle$ If $\theta \in \triangle(\mathbb{E} \circ \phi)$ then $\theta \in \triangle(\mathbb{E} \circ \theta \circ \phi)$
$\mathbf{(C4)}_\triangle$ If $\neg\theta \notin \triangle(\mathbb{E} \circ \phi)$ then $\neg\theta \notin \triangle(\mathbb{E} \circ \theta \circ \phi)$

Briefly, these postulates can be explained as follows: The rule $\mathbf{(C1)}_\triangle$ says that if two inputs are received, the second being more specific than the first, then the first is rendered redundant (at least regarding its effects on the regular belief set). Rule $\mathbf{(C2)}_\triangle$ says that if two contradictory inputs are received, then the most recent one prevails. Rule $\mathbf{(C3)}_\triangle$ says that an input θ should be in the regular beliefs after receiving the subsequent input ϕ if θ would have been believed given input ϕ to begin with. Finally $\mathbf{(C4)}_\triangle$ says that if θ is not contradicted after receipt of input ϕ, then it should still be uncontradicted if input ϕ is preceded by input θ itself.[4] Which of these postulates are suitable for core-invariant regular revision? While $\mathbf{(C3)}_\triangle$ and $\mathbf{(C4)}_\triangle$ seem to retain their validity in our setting, there is a slight problem with $\mathbf{(C1)}_\triangle$ and $\mathbf{(C2)}_\triangle$ concerning the case when ϕ is taken to be a core-contravening sentence, i.e., when $\neg\phi \in \blacktriangle(\mathbb{E}) = \blacktriangle(\mathbb{E} \circ \theta)$. Consider momentarily the following two properties:

$\mathbf{(wC1)}_\triangle$ If $\neg\phi \in \blacktriangle(\mathbb{E})$ and $\phi \to \theta \in Cn(\emptyset)$ then $\triangle(\mathbb{E} \circ \theta) = \triangle(\mathbb{E})$
$\mathbf{(wC2)}_\triangle$ If $\neg\phi \in \blacktriangle(\mathbb{E})$ and $\phi \to \neg\theta \in Cn(\emptyset)$ then $\triangle(\mathbb{E} \circ \theta) = \triangle(\mathbb{E})$

Then it can easily be shown that any core-invariant regular revision operator $\circ$ which satisfies $\mathbf{(C1)}_\triangle$, respectively $\mathbf{(C2)}_\triangle$, also satisfies $\mathbf{(wC1)}_\triangle$, respectively $\mathbf{(wC2)}_\triangle$. However one can easily find examples of core-invariant regular revision operators which fail to satisfy the latter two properties. For example let p and q be propositional variables and suppose $\mathbb{E}$ is such that $\blacktriangle(\mathbb{E}) = \triangle(\mathbb{E}) = Cn(\neg p)$.

[4] We remark that these postulates have not been *totally* immune to criticism in the literature. In particular $\mathbf{(C2)}_\triangle$ is viewed by some as problematic (see [5, 7, 17]).

Then clearly we have $\neg p \in \blacktriangle(\mathbb{E})$ and $p \to (p \vee q), p \to \neg(\neg p \wedge q) \in Cn(\emptyset)$. But $\triangle(\mathbb{E} \circ (p \vee q)) = \triangle(\mathbb{E}) + (p \vee q) = Cn(\neg p \wedge q)$ (contradicting **(wC1)**$_\triangle$), while $\triangle(\mathbb{E} \circ (\neg p \wedge q)) = \triangle(\mathbb{E}) + (\neg p \wedge q) = Cn(\neg p \wedge q)$ (contradicting **(wC2)**$_\triangle$). Thus we conclude that **(C1)**$_\triangle$ and **(C2)**$_\triangle$ are not suitable as they stand. Instead we propose to modify them so that they apply only when $\neg\phi \notin \blacktriangle(\mathbb{E})$.

(C1′)$_\triangle$ If $\neg\phi \notin \blacktriangle(\mathbb{E})$ and $\phi \to \theta \in Cn(\emptyset)$ then $\triangle(\mathbb{E} \circ \theta \circ \phi) = \triangle(\mathbb{E} \circ \phi)$
(C2′)$_\triangle$ If $\neg\phi \notin \blacktriangle(\mathbb{E})$ and $\phi \to \neg\theta \in Cn(\emptyset)$ then $\triangle(\mathbb{E} \circ \theta \circ \phi) = \triangle(\mathbb{E} \circ \phi)$

We tend to view **(C1′)**$_\triangle$, **(C2′)**$_\triangle$, **(C3)**$_\triangle$ and **(C4)**$_\triangle$ as being minimal conditions on iterated regular revision. An interesting consequence of **(C2′)**$_\triangle$ is revealed by the following proposition.

Proposition 2. *Let $\circ$ be a core-invariant regular revision operator which satisfies* **(C2′)**$_\triangle$*. Then, for all $\mathbb{E} \in Ep$ and $\theta, \phi \in L^+$, we have that $\neg\theta \in \blacktriangle(\mathbb{E})$ implies $\triangle(\mathbb{E} \circ \theta \circ \phi) = \triangle(\mathbb{E} \circ \phi)$.*

The above proposition says that not only does revising by a core-contravening sentence θ have no effect on the regular belief set, it also has no impact on the result of revising by any subsequent *regular* inputs. (As we will see later, this does not necessarily mean that core-contravening regular inputs are *totally* devoid of impact.)

As in our current set-up, Darwiche and Pearl assume that the new belief set $\triangle(\mathbb{E} \circ \theta)$ resulting from the single revision by θ is determined AGM-style by a total pre-order $\leq^{\mathbb{E}}_{\triangle}$ anchored on $[\triangle(\mathbb{E})]$. Likewise the new belief set $\triangle(\mathbb{E} \circ \theta \circ \phi)$ following a *subsequent* revision by ϕ is then determined by the total pre-order $\leq^{\mathbb{E}\circ\theta}_{\triangle}$ anchored on $[\triangle(\mathbb{E} \circ \theta)]$. Thus the question of which properties of iterated revision are satisfied is essentially the same as asking what the new pre-order $\leq^{\mathbb{E}\circ\theta}_{\triangle}$ looks like. In a result in [8], Darwiche and Pearl show how each of their postulates **(C1)**$_\triangle$–**(C4)**$_\triangle$ regulates a different aspect of the relationship between $\leq^{\mathbb{E}}_{\triangle}$ and the new regular pre-order $\leq^{\mathbb{E}\circ\theta}_{\triangle}$. The following proposition, which may be viewed as a *generalisation* of Darwiche and Pearl's result (roughly speaking, Darwiche and Pearl are looking at the special case when $\blacktriangle(\mathbb{E}) = Cn(\emptyset)$), does the same for **(C1′)**$_\triangle$, **(C2′)**$_\triangle$, **(C3)**$_\triangle$ and **(C4)**$_\triangle$ in our non-prioritised setting.

Proposition 3. *Let $\circ$ be a core-invariant regular revision operator. Then $\circ$ satisfies* **(C1′)**$_\triangle$*,* **(C2′)**$_\triangle$*,* **(C3)**$_\triangle$ *and* **(C4)**$_\triangle$ *iff each of the following conditions holds for all $\mathbb{E} \in Ep$ and $\theta \in L^+$:*

(1) For all $w_1, w_2 \in [\blacktriangle(\mathbb{E})] \cap [\theta]$, $w_1 \leq^{\mathbb{E}\circ\theta}_{\triangle} w_2$ iff $w_1 \leq^{\mathbb{E}}_{\triangle} w_2$
(2) For all $w_1, w_2 \in [\blacktriangle(\mathbb{E})] \cap [\neg\theta]$, $w_1 \leq^{\mathbb{E}\circ\theta}_{\triangle} w_2$ iff $w_1 \leq^{\mathbb{E}}_{\triangle} w_2$
(3) For all $w_1, w_2 \in [\blacktriangle(\mathbb{E})]$, if $w_1 \in [\theta], w_2 \in [\neg\theta]$ and $w_1 <^{\mathbb{E}}_{\triangle} w_2$, then $w_1 <^{\mathbb{E}\circ\theta}_{\triangle} w_2$
(4) For all $w_1, w_2 \in [\blacktriangle(\mathbb{E})]$, if $w_1 \in [\theta], w_2 \in [\neg\theta]$ and $w_1 \leq^{\mathbb{E}}_{\triangle} w_2$, then $w_1 \leq^{\mathbb{E}\circ\theta}_{\triangle} w_2$

Thus, according to the above proposition, **(C1′)**$_\triangle$ corresponds to the requirement that, in the transformation from $\leq^{\mathbb{E}}_{\triangle}$ to $\leq^{\mathbb{E}\circ\theta}_{\triangle}$, the relative ordering between the $[\theta]$-worlds in $[\blacktriangle(\mathbb{E})]$ remains unchanged. $\mathbf{(C2')}_\triangle$ corresponds to the same requirement but with regard to the $[\neg\theta]$-worlds in $[\blacktriangle(\mathbb{E})]$. $\mathbf{(C3)}_\triangle$ corresponds to the requirement that if a given $[\theta]$-world in $[\blacktriangle(\mathbb{E})]$ was regarded as strictly more plausible than a given $[\neg\theta]$-world in $[\blacktriangle(\mathbb{E})]$ *before* receipt of the input θ, then this relation should be preserved *after* receipt of θ. Finally $\mathbf{(C4)}_\triangle$ matches the same requirement as $\mathbf{(C3)}_\triangle$, but with "at least as plausible as" substituted for "strictly more plausible than". Note how each property only constrains the transformation from $\leq^{\mathbb{E}}_\triangle$ to $\leq^{\mathbb{E}\circ\theta}_\triangle$ **within** $[\blacktriangle(\mathbb{E})]$. We will later see some conditions which constrain the movement of the other worlds.

The Darwiche and Pearl postulates form our starting point in the study of iterated revision. However, other postulates have been suggested. In particular, another postulate of interest which may be found in the literature on iterated AGM revision (cf. the rule (Recalcitrance) in [21]) is:

$\mathbf{(C5)}_\triangle$ If $\phi \to \neg\theta \notin Cn(\emptyset)$ then $\theta \in \triangle(\mathbb{E} \circ \theta \circ \phi)$

Note that this postulate is in fact a strengthening of $\mathbf{(C3)}_\triangle$ and $\mathbf{(C4)}_\triangle$. (This will also soon follow from Proposition 5.) In fact $\mathbf{(C5)}_\triangle$ might just have well have been called "strong success" since it also implies that $\theta \in \triangle(\mathbb{E} \circ \theta)$ for all $\theta \in L^+$. (Hint: substitute $\top$ for ϕ.) For this reason the postulate, as it stands, is obviously *not* suitable in our non-prioritised setting. However the following weaker version will be of interest to us:

$\mathbf{(C5')}_\triangle$ If $\phi \to \neg\theta \notin \blacktriangle(\mathbb{E})$ then $\theta \in \triangle(\mathbb{E} \circ \theta \circ \phi)$.

$\mathbf{(C5')}_\triangle$ entails that if, having received a regular input θ, we *do* decide to accept it, then we do so wholeheartedly (or as wholeheartedly as we can without actually elevating it to the status of a core belief!) in that the only way it can be dislodged from the belief set by a succeeding regular input is if that input contradicts it given the core beliefs $\blacktriangle(\mathbb{E})$. This postulate too can be translated into a somewhat plausible constraint on the new regular pre-order $\leq^{\mathbb{E}\circ\theta}_\triangle$.

Proposition 4. *Let $\circ$ be a core-invariant regular revision operator. Then $\circ$ satisfies* $\mathbf{(C5')}_\triangle$ *iff, for each $\mathbb{E} \in Ep$ and $\theta \in L^+$, and for all $w_1, w_2 \in [\blacktriangle(\mathbb{E})]$, if $w_1 \in [\theta]$ and $w_2 \in [\neg\theta]$ then $w_1 <^{\mathbb{E}\circ\theta}_\triangle w_2$.*

Thus $\mathbf{(C5')}_\triangle$ corresponds to the property that all the $[\theta]$-worlds in $[\blacktriangle(\mathbb{E})]$ are deemed strictly more plausible by $\leq^{\mathbb{E}\circ\theta}_\triangle$ than all the $[\neg\theta]$-worlds in $[\blacktriangle(\mathbb{E})]$. $\mathbf{(C5')}_\triangle$ is related to our previous postulates in the following way:

Proposition 5. *Let $\circ$ be a core-invariant regular revision operator which satisfies* $\mathbf{(C5')}_\triangle$. *Then $\circ$ satisfies* $\mathbf{(C3)}_\triangle$ *and* $\mathbf{(C4)}_\triangle$.

4 Core Belief Inputs

So far we have assumed that the set of core beliefs in an epistemic state $\mathbb{E}$ remains constant under regular revision inputs. In this section we want to consider the case when the core beliefs are themselves subject to revision by external inputs. To do this we shall now assume that we are given a second type of revision operator on epistemic states which we denote by $\bullet$. Given $\mathbb{E} \in Ep$ and $\phi \in L^+$, $\mathbb{E} \bullet \phi$ will denote the result of revising $\mathbb{E}$ so that ϕ is included as a core belief.[5] The operator $\bullet$ is distinct from $\circ$, though intuitively we should expect some interaction between the two. Once again we consider single-step revision and iterated revision in turn.

4.1 Single-Step Core Revision

What constraints should we put on $\blacktriangle(\mathbb{E} \bullet \phi)$? Well first of all, in order to simplify matters and unlike for $\circ$, we shall assume that **every revision using $\bullet$ is successful**, i.e., $\phi \in \blacktriangle(\mathbb{E} \bullet \phi)$.[6] For example core belief inputs might correspond to information from a source which the agent deems to be highly reliable or trustworthy, such as first-hand observations. A reasonable possibility is then to treat the core beliefs as we would any other belief set in this case and assume that the new core can be obtained by applying some simple AGM revision function for $\blacktriangle(\mathbb{E})$. Equivalently, by Theorem 1, we assume, for each $\mathbb{E} \in Ep$, the existence of a total pre-order $\leq^{\mathbb{E}}_{\blacktriangle}$ on $\mathcal{W}$, anchored on $[\blacktriangle(\mathbb{E})]$, such that, for all $\phi \in L^+$,

$$\blacktriangle(\mathbb{E} \bullet \phi) = Th(\min([\phi], \leq^{\mathbb{E}}_{\blacktriangle})). \tag{2}$$

Definition 4. *Let $\bullet : Ep \times L^+ \to Ep$ be a function. Then $\bullet$ is a* core revision operator (on the epistemic frame $\langle Ep, \triangle, \blacktriangle\rangle$) *if, for each $\mathbb{E} \in Ep$, there exists a total pre-order $\leq^{\mathbb{E}}_{\blacktriangle}$ on $\mathcal{W}$, anchored on $[\blacktriangle(\mathbb{E})]$, such that $\blacktriangle(\mathbb{E} \bullet \phi)$ may be determined as in (2) above. We call $\leq^{\mathbb{E}}_{\blacktriangle}$ the* core pre-order associated to $\mathbb{E}$ (according to $\bullet$)*.*

So, to have both a regular revision operator *and* a core revision operator on an epistemic frame $\langle Ep, \triangle, \blacktriangle\rangle$ means to assume that each epistemic state $\mathbb{E} \in Ep$ comes equipped with **two** total pre-orders $\leq^{\mathbb{E}}_{\triangle}$ and $\leq^{\mathbb{E}}_{\blacktriangle}$, anchored on $[\triangle(\mathbb{E})]$ and $[\blacktriangle(\mathbb{E})]$ respectively. The interplay between these two orders will be of concern throughout the rest of the paper.

What constraints should we be putting on $\triangle(\mathbb{E} \bullet \phi)$? This question isn't so easy to answer. Here we need to keep in mind that we must have $\blacktriangle(\mathbb{E} \bullet \phi) \subseteq \triangle(\mathbb{E}\bullet\phi)$ and so, since we are assuming we always have $\phi \in \blacktriangle(\mathbb{E}\bullet\phi)$, we necessarily require $\phi \in \triangle(\mathbb{E} \bullet \phi)$. Hence if $\phi \notin \triangle(\mathbb{E})$ then some changes to the regular beliefs will certainly be necessary. In the case when $\neg\phi \notin \blacktriangle(\mathbb{E})$ it seems reasonable to expect that $\triangle(\mathbb{E})$ should be revised just as if ϕ was a regular belief input, i.e,:

(Y1) If $\neg\phi \notin \blacktriangle(\mathbb{E})$ then $\triangle(\mathbb{E} \bullet \phi) = \triangle(\mathbb{E} \circ \phi)$ (Cross-Vacuity)

[5] To put it another way in terms of revising epistemic states by *conditional* beliefs [4, 16]: whereas a regular revision by ϕ may be equated with a revision by the conditional $\top \Rightarrow \phi$, a *core* revision by ϕ may be equated with a revision by the conditional $\neg\phi \Rightarrow \bot$.

[6] An interesting alternative could be to reject ϕ from the core belief set, but include it instead merely as a regular belief.

This postulate gives us our basic point of contact between core revision and regular revision.

What should we do if $\neg\phi \in \blacktriangle(\mathbb{E})$? In this case we can't set $\triangle(\mathbb{E}\bullet\phi) = \triangle(\mathbb{E}\circ\phi)$ since ϕ is not contained in the right-hand side. One possibility could be to just throw away the distinctions between core belief and regular belief in this case by setting

(sY2) If $\neg\phi \in \blacktriangle(\mathbb{E})$ then $\triangle(\mathbb{E} \bullet \phi) = \blacktriangle(\mathbb{E} \bullet \phi)$. (Regular Collapse)

However this seems a bit drastic. A more interesting possibility which we intend to explore in future work could be to adopt a Levi-style approach (cf. the Levi Identity [10]) and decompose the operation into two steps: first remove $\neg\phi$ from $\blacktriangle(\mathbb{E})$ using some sort of "core contraction" operation, and *then* revise by ϕ using $\circ$. For now, though, we take a different approach. Note that **(Y1)** says that, in the case when $\neg\phi \notin \blacktriangle(\mathbb{E})$, we should just use the pre-order $\leq^{\mathbb{E}}_{\triangle}$ to determine the new regular belief set. Why not just use $\leq^{\mathbb{E}}_{\triangle}$ also in the case when $\neg\phi \in \blacktriangle(\mathbb{E})$? That is we just set, in *all* cases

$$\triangle(\mathbb{E} \bullet \phi) = Th(\min([\phi], \leq^{\mathbb{E}}_{\triangle})). \tag{3}$$

However we need to be careful here, for remember we must have $\blacktriangle(\mathbb{E} \bullet \phi) \subseteq \triangle(\mathbb{E} \bullet \phi)$. This will be ensured if we require the two pre-orders $\leq^{\mathbb{E}}_{\triangle}$ and $\leq^{\mathbb{E}}_{\blacktriangle}$ to *cohere* with one another in a certain respect. Namely if we require

$$\leq^{\mathbb{E}}_{\triangle} \subseteq \leq^{\mathbb{E}}_{\blacktriangle},$$

i.e., that $\leq^{\mathbb{E}}_{\triangle}$ is a *refinement* of $\leq^{\mathbb{E}}_{\blacktriangle}$. This is confirmed by the following result.

Proposition 6. *Let $\mathbb{E} \in Ep$ and let $\leq^{\mathbb{E}}_{\triangle}, \leq^{\mathbb{E}}_{\blacktriangle}$ be two total pre-orders on $\mathcal{W}$ anchored on $[\triangle(\mathbb{E})]$ and $[\blacktriangle(\mathbb{E})]$ respectively. If $\leq^{\mathbb{E}}_{\triangle} \subseteq \leq^{\mathbb{E}}_{\blacktriangle}$ then, for all $\phi \in L^+$, we have $Th(\min([\phi], \leq^{\mathbb{E}}_{\blacktriangle})) \subseteq Th(\min([\phi], \leq^{\mathbb{E}}_{\triangle}))$.*

As we will shortly see in Theorem 2, defining $\triangle(\mathbb{E} \bullet \phi)$ as in (3) above has the consequence that, in addition to **(Y1)**, the following two properties are satisfied.

(Y2) If $\neg\phi \notin \triangle(\mathbb{E} \bullet \theta)$ then $\triangle(\mathbb{E} \bullet (\theta \wedge \phi)) = \triangle(\mathbb{E} \bullet \theta) + \phi$ (Cross-Conjunction 1)

(Y3) If $\phi_1 \leftrightarrow \phi_2 \in Cn(\emptyset)$ then $\triangle(\mathbb{E} \bullet \phi_1) = \triangle(\mathbb{E} \bullet \phi_2)$ (Cross-Extensionality)

The first property above is similar to the AGM postulate **(K*6)** from Sect. 2. It says that the new regular belief set after core-revising by $\theta \wedge \phi$ should be obtainable by first core-revising by θ and then simply expanding the resultant regular belief set $\triangle(\mathbb{E} \bullet \theta)$ by ϕ, *provided* ϕ is consistent with $\triangle(\mathbb{E} \bullet \theta)$. It is easy to see that, for core-revision operators, **(sY2)** implies **(Y2)**. The second property above expresses the reasonable requirement that core-revising by logically equivalent sentences should yield the same regular belief set. We make the following definition:

Definition 5. *Let $\circ$ and $\bullet$ be a core-invariant regular revision operator and a core revision operator on the epistemic frame $\langle Ep, \triangle, \blacktriangle\rangle$ respectively. If $\circ$ and $\bullet$ together satisfy* **(Y1)** *and $\bullet$ satisfies* **(Y2)** *and* **(Y3)** *then we call the pair $\langle \circ, \bullet\rangle$ a* revision system *(on $\langle Ep, \triangle, \blacktriangle\rangle$).*

The next theorem is one of the main results of this paper. It gives a characterisation for revision systems.

Theorem 2. *Let $\langle Ep, \triangle, \blacktriangle\rangle$ be an epistemic frame and let $\circ, \bullet : Ep \times L^+ \to Ep$ be two functions. Then the following are equivalent:*
(i). *$\langle \circ, \bullet\rangle$ is a revision system on $\langle Ep, \triangle, \blacktriangle\rangle$.*
(ii). *For each $\mathbb{E} \in Ep$ there exist a total pre-order $\leq^{\mathbb{E}}_{\triangle}$ on $\mathcal{W}$ anchored on $[\triangle(\mathbb{E})]$, and a total pre-order $\leq^{\mathbb{E}}_{\blacktriangle}$ on $\mathcal{W}$ anchored on $[\blacktriangle(\mathbb{E})]$ such that $\leq^{\mathbb{E}}_{\triangle} \subseteq \leq^{\mathbb{E}}_{\blacktriangle}$ and, for all $\phi \in L^+$,*

$$\triangle(\mathbb{E} \bullet \phi) = Th(\min([\phi], \leq^{\mathbb{E}}_{\triangle})) \qquad \blacktriangle(\mathbb{E} \bullet \phi) = Th(\min([\phi], \leq^{\mathbb{E}}_{\blacktriangle}))$$

$$\triangle(\mathbb{E} \circ \phi) = \begin{cases} Th(\min([\phi], \leq^{\mathbb{E}}_{\triangle})) & \textit{if } \neg\phi \notin \blacktriangle(\mathbb{E}) \\ \triangle(\mathbb{E}) & \textit{otherwise} \end{cases} \qquad \blacktriangle(\mathbb{E} \circ \phi) = \blacktriangle(\mathbb{E})$$

Proof (Sketch). To show that (i) implies (ii), let $\langle \circ, \bullet\rangle$ be a revision system on $\langle Ep, \triangle, \blacktriangle\rangle$. Then, by definition, $\circ$ is a core-invariant regular revision operator. Hence there exists, for each $\mathbb{E} \in Ep$, a total pre-order $\leq^{\mathbb{E}}_{r}$ on $\mathcal{W}$ anchored on $[\triangle(\mathbb{E})]$ such that, for all $\phi \in L^+$,

$$\triangle(\mathbb{E} \circ \phi) = \begin{cases} Th(\min([\phi], \leq^{E}_{r})) & \text{if } \neg\phi \notin \blacktriangle(\mathbb{E}) \\ \triangle(\mathbb{E}) & \text{otherwise} \end{cases}$$

and $\blacktriangle(\mathbb{E} \circ \phi) = \blacktriangle(\mathbb{E})$. We also know that $\bullet$ is a core revision operator. Hence, for each $\mathbb{E} \in Ep$ there also exists a total pre-order $\leq^{\mathbb{E}}_{\blacktriangle}$ on $\mathcal{W}$ anchored on $[\blacktriangle(\mathbb{E})]$ such that, for all $\phi \in L^+$ we have $\blacktriangle(\mathbb{E} \bullet \phi) = Th(\min([\phi], \leq^{\mathbb{E}}_{\blacktriangle}))$. It might be hoped now that $\leq^{\mathbb{E}}_{r}$ and $\leq^{\mathbb{E}}_{\blacktriangle}$ then give us our required pair of pre-orders, however we first need to make some modification to $\leq^{\mathbb{E}}_{r}$. We define a new ordering $\leq^{\mathbb{E}}_{\triangle}$ which agrees with $\leq^{\mathbb{E}}_{r}$ within $[\blacktriangle(\mathbb{E})]$ and likewise makes all $[\blacktriangle(\mathbb{E})]$-worlds more plausible than all non-$[\blacktriangle(\mathbb{E})]$-worlds. However, $\leq^{\mathbb{E}}_{\triangle}$ orders the non-$[\blacktriangle(\mathbb{E})]$-worlds differently from $\leq^{\mathbb{E}}_{r}$. Precisely we set, for $w_1, w_2 \in \mathcal{W}$,

$$\begin{aligned} w_1 \leq^{\mathbb{E}}_{\triangle} w_2 \text{ iff } & w_1, w_2 \in [\blacktriangle(\mathbb{E})] \text{ and } w_1 \leq^{\mathbb{E}}_{r} w_2 \\ & \textbf{or } w_1 \in [\blacktriangle(\mathbb{E})] \text{ and } w_2 \notin [\blacktriangle(\mathbb{E})] \\ & \textbf{or } w_1, w_2 \notin [\blacktriangle(\mathbb{E})] \text{ and } \neg\alpha_1 \notin \triangle(\mathbb{E} \bullet (\alpha_1 \vee \alpha_2)) \end{aligned}$$

In the last line here, α_i is any sentence such that $[\alpha_i] = \{w_i\}$ $(i = 1, 2)$. (By **(Y3)** the precise choice of α_i is irrelevant.) It can then be shown that $\leq^{\mathbb{E}}_{\triangle}$ is a total pre-order anchored on $[\triangle(\mathbb{E})]$ and that $\leq^{\mathbb{E}}_{\triangle}$ and $\leq^{\mathbb{E}}_{\blacktriangle}$ then give the required pair of pre-orders.

The proof that (ii) implies (i) is straightforward. □

For the rest of this section we assume $\langle \circ, \bullet\rangle$ to be an arbitrary but fixed revision system.

4.2 Iterating Core Beliefs Revision

What should be the effect on $\blacktriangle(\mathbb{E})$ and $\triangle(\mathbb{E})$ of iterated applications of $\bullet$? For the former, since we are assuming $\bullet$ behaves just like an AGM revision operator with regard to $\blacktriangle(\mathbb{E})$, the iterated AGM revision postulates mentioned in Sect. 3.2 are again relevant. Rephrased in terms of core revision, they are:

(C1)$_\blacktriangle$ If $\phi \to \theta \in Cn(\emptyset)$ then $\blacktriangle(\mathbb{E} \bullet \theta \bullet \phi) = \blacktriangle(\mathbb{E} \bullet \phi)$
(C2)$_\blacktriangle$ If $\phi \to \neg\theta \in Cn(\emptyset)$ then $\blacktriangle(\mathbb{E} \bullet \theta \bullet \phi) = \blacktriangle(\mathbb{E} \bullet \phi)$
(C3)$_\blacktriangle$ If $\theta \in \blacktriangle(\mathbb{E} \bullet \phi)$ then $\theta \in \blacktriangle(\mathbb{E} \bullet \theta \bullet \phi)$
(C4)$_\blacktriangle$ If $\neg\theta \notin \blacktriangle(\mathbb{E} \bullet \phi)$ then $\neg\theta \notin \blacktriangle(\mathbb{E} \bullet \theta \bullet \phi)$
(C5)$_\blacktriangle$ If $\phi \to \theta \notin Cn(\emptyset)$ then $\theta \in \blacktriangle(\mathbb{E} \bullet \theta \bullet \phi)$

We take **(C1)$_\blacktriangle$**–**(C4)$_\blacktriangle$** to be minimal requirements. We remind the reader that **(C5)$_\blacktriangle$** implies both **(C3)$_\blacktriangle$** and **(C4)$_\blacktriangle$**. The characterisation result of Darwiche and Pearl already tells us how each of **(C1)$_\blacktriangle$**–**(C4)$_\blacktriangle$** regulates a certain aspect of the relationship between $\leq^{\mathbb{E}}_\blacktriangle$ and the new *core* pre-order $\leq^{\mathbb{E}\bullet\theta}_\blacktriangle$. **(C5)$_\blacktriangle$** also corresponds to a constraint on $\leq^{\mathbb{E}\bullet\theta}_\blacktriangle$. The proof of this correspondence is implicit in [21].

Proposition 7 ([8, 21]). *Let $\bullet$ be a core revision operator. Then $\bullet$ satisfies* **(C1)$_\blacktriangle$**–**(C5)$_\blacktriangle$** *iff each of the following conditions hold for all $\mathbb{E} \in Ep$ and $\theta \in L^+$:*

(1) For all $w_1, w_2 \in [\theta]$, $w_1 \leq^{\mathbb{E}\bullet\theta}_\blacktriangle w_2$ iff $w_1 \leq^{\mathbb{E}}_\blacktriangle w_2$
(2) For all $w_1, w_2 \in [\neg\theta]$, $w_1 \leq^{\mathbb{E}\bullet\theta}_\blacktriangle w_2$ iff $w_1 \leq^{\mathbb{E}}_\blacktriangle w_2$
(3) For all $w_1, w_2 \in \mathcal{W}$, if $w_1 \in [\theta], w_2 \in [\neg\theta]$ and $w_1 <^{\mathbb{E}}_\blacktriangle w_2$, then $w_1 <^{\mathbb{E}\bullet\theta}_\blacktriangle w_2$
(4) For all $w_1, w_2 \in \mathcal{W}$, if $w_1 \in [\theta], w_2 \in [\neg\theta]$ and $w_1 \leq^{\mathbb{E}}_\blacktriangle w_2$, then $w_1 \leq^{\mathbb{E}\bullet\theta}_\blacktriangle w_2$
(5) For all $w_1, w_2 \in \mathcal{W}$, if $w_1 \in [\theta]$ and $w_2 \in [\neg\theta]$ then $w_1 <^{\mathbb{E}\bullet\theta}_\blacktriangle w_2$

For the case of $\triangle(\mathbb{E})$ we expect that the behaviour of the regular belief set under a sequence of core inputs should be connected in some way with the behaviour of the core itself. But how? Here we present one idea, which is perhaps best motivated directly in terms of the two pre-orders $\leq^{\mathbb{E}}_\triangle$ and $\leq^{\mathbb{E}}_\blacktriangle$ which we take to underlie a given epistemic state $\mathbb{E}$. First note that the question of how $\triangle(\mathbb{E})$ should behave under sequences of core inputs essentially reduces to the question of what the new regular pre-order $\leq^{\mathbb{E}\bullet\theta}_\triangle$ following the core input θ should look like. One constraint on $\leq^{\mathbb{E}\bullet\theta}_\triangle$ is already in place, namely that $\leq^{\mathbb{E}\bullet\theta}_\triangle \subseteq \leq^{\mathbb{E}\bullet\theta}_\blacktriangle$. Our idea is to carry over as much of the structure of $\leq^{\mathbb{E}}_\triangle$ to $\leq^{\mathbb{E}\bullet\theta}_\triangle$ as possible, while obeying this constraint. This can be achieved by defining $\leq^{\mathbb{E}\bullet\theta}_\triangle$ simply to be the lexicographic *refinement* of $\leq^{\mathbb{E}\bullet\theta}_\blacktriangle$ *by* $\leq^{\mathbb{E}}_\triangle$, i.e., for all $w_1, w_2 \in \mathcal{W}$,

$$w_1 \leq^{\mathbb{E}\bullet\theta}_\triangle w_2 \text{ iff either } w_1 <^{\mathbb{E}\bullet\theta}_\blacktriangle w_2 \text{ or } w_1 \sim^{\mathbb{E}\bullet\theta}_\blacktriangle w_2 \text{ and } w_1 \leq^{\mathbb{E}}_\triangle w_2.$$

We remark that the idea of combining pre-orders using lexicographic refinement crops up several times in the literature on belief revision, e.g. [20]. It turns out that this behaviour may be characterised here by the following property:

(Z1) If $\blacktriangle(\mathbb{E} \bullet \theta \bullet \phi) \subseteq \blacktriangle(\mathbb{E} \bullet \phi)$ then $\triangle(\mathbb{E} \bullet \theta \bullet \phi) = \triangle(\mathbb{E} \bullet \phi)$ (Coupling)

This property says that if a core input ϕ can be preceded by another core input θ *without increasing* the set of core beliefs, then preceding ϕ by θ should lead to precisely the same set of regular beliefs.

Proposition 8. *$\bullet$ satisfies* **(Z1)** *iff, for each $\mathbb{E} \in Ep$ and $\theta \in L^+$, $\leq_{\triangle}^{\mathbb{E}\bullet\theta}$ is equal to the lexicographic refinement of $\leq_{\blacktriangle}^{\mathbb{E}\bullet\theta}$ by $\leq_{\triangle}^{\mathbb{E}}$.*

The appearance of "$\subseteq$" rather than "$=$" in the antecedent of **(Z1)** may seem surprising. However, as the next proposition shows, if $\bullet$ satisfies certain other properties then there is no difference.

Proposition 9. *Let $\bullet$ be a core revision operator which satisfies* **(C1)**$_\blacktriangle$, **(C2)**$_\blacktriangle$ *and* **(C5)**$_\blacktriangle$. *Then, for all $\mathbb{E} \in Ep$ and $\theta, \phi \in L^+$, we have $\blacktriangle(\mathbb{E} \bullet \theta \bullet \phi) \subseteq \blacktriangle(\mathbb{E} \bullet \phi)$ iff $\blacktriangle(\mathbb{E} \bullet \theta \bullet \phi) = \blacktriangle(\mathbb{E} \bullet \phi)$.*

Finally in this section we have the following result, which attests to the strength of **(Z1)**.

Proposition 10. *If $\bullet$ satisfies* **(Z1)** *and* **(C1)**$_\blacktriangle$, **(C2)**$_\blacktriangle$ *and* **(C5)**$_\blacktriangle$, *then $\bullet$ also satisfies* **(C1)**$_\triangle$, **(C2)**$_\triangle$ *and* **(C5)**$_\triangle$ *(with $\bullet$ in place of $\circ$).*

5 Mixing Regular and Core Revision

In this section we explore the possibility of performing mixed sequences of revisions, containing both regular and core inputs. We give a number of possible properties, relating each one to a condition in terms of the underlying regular and core pre-orders, and also showing how they relate to the postulates of the previous sections. In the following section we will give a concrete pair of revision operators which in fact satisfies all the postulates from this section. The intuitive simplicity of these constructed operators will lend some support to the reasonableness of these postulates. Throughout this section we again assume $\langle \circ, \bullet \rangle$ is an arbitrary but fixed revision system.

Recall that we have assumed regular revision to be core-invariant, i.e., a regular belief input leaves the core belief set unchanged. Our first property seeks to lessen the impact of regular inputs on the core even further. It says that a regular input also has no influence on what the core beliefs should look like after the *next core input*.

(M0) $\blacktriangle(\mathbb{E} \circ \theta \bullet \phi) = \blacktriangle(\mathbb{E} \bullet \phi)$ (Core-Conditional-Invariance)

Since $\circ$ is core-invariant, the right-hand-side here is equal to $\blacktriangle(\mathbb{E} \bullet \phi \circ \theta)$. Hence, for a revision system, **(M0)** is equivalent to the following property which says that when receiving two consecutive inputs, one of which is a regular input and the other a core input, the resulting core belief set is actually independent of the order in which these two inputs are received.

(Comm)$_\blacktriangle$ $\blacktriangle(\mathbb{E} \circ \theta \bullet \phi) = \blacktriangle(\mathbb{E} \bullet \phi \circ \theta)$

In terms of the underlying pre-orders, the property **(M0)** corresponds to the simple requirement that the entire core pre-order remains unchanged under regular revision. Thus **(M0)** can be seen as a kind of "scaling-up" of the property of core-invariance to apply to the iterated case.

Proposition 11. *$\langle \circ, \bullet \rangle$ satisfies* **(M0)** *iff, for all $\mathbb{E} \in Ep$ and $\theta \in L^+$, we have $\leq_\blacktriangle^{\mathbb{E}\circ\theta} = \leq_\blacktriangle^{\mathbb{E}}$.*

For the next batch of postulates, we consider again Darwiche and Pearl's original postulates **(C1)$_\triangle$**–**(C4)$_\triangle$** for iterated regular revision listed at the beginning of Sect. 3.2. The postulates **(C1)$_\triangle$** and **(C2)$_\triangle$** provide conditions under which the effects of a regular input on the regular belief set should be "overruled" by a succeeding regular input. Namely when the second is more specific than the first (**(C1)$_\triangle$**), or when it contradicts it (**(C2)$_\triangle$**). Since core inputs can be viewed as carrying more "weight" than regular inputs, it seems reasonable to suggest that in both cases this overruling should also occur when the second input is upgraded from being a regular input to being a *core* input:

(M1) If $\phi \to \theta \in Cn(\emptyset)$ then $\triangle(\mathbb{E} \circ \theta \bullet \phi) = \triangle(\mathbb{E} \bullet \phi)$
(M2) If $\phi \to \neg\theta \in Cn(\emptyset)$ then $\triangle(\mathbb{E} \circ \theta \bullet \phi) = \triangle(\mathbb{E} \bullet \phi)$

We propose similar modifications to **(C3)$_\triangle$** and **(C4)$_\triangle$**. If θ would be a regular belief after receiving the core input ϕ then preceding this core input with the regular input θ should not change this fact. Similarly if θ is not discounted as a regular belief after receiving the core input ϕ then preceding this core input with the regular input θ should not change this fact.

(M3) If $\theta \in \triangle(\mathbb{E} \bullet \phi)$ then $\theta \in \triangle(\mathbb{E} \circ \theta \bullet \phi)$
(M4) If $\neg\theta \notin \triangle(\mathbb{E} \bullet \phi)$ then $\neg\theta \notin \triangle(\mathbb{E} \circ \theta \bullet \phi)$

Recall Proposition 3 which showed how each of the postulates **(C1′)$_\triangle$**, **(C2′)$_\triangle$**, **(C3)$_\triangle$** and **(C4)$_\triangle$** corresponded to a constraint on the new regular pre-order $\leq_\triangle^{\mathbb{E}\circ\theta}$ **within** $[\blacktriangle(\mathbb{E})]$. The following result shows how each of **(M1)**–**(M4)** corresponds to the same constraints on $\leq_\triangle^{\mathbb{E}\circ\theta}$, but extended to apply to the *whole* of $\mathcal{W}$.

Proposition 12. *$\langle \circ, \bullet \rangle$ satisfies* **(M1)**–**(M4)** *iff the following hold, for each $\mathbb{E} \in Ep$ and $\theta \in L^+$:*

(1) For all $w_1, w_2 \in [\theta]$, $w_1 \leq_\triangle^{\mathbb{E}\circ\theta} w_2$ iff $w_1 \leq_\triangle^{\mathbb{E}} w_2$
(2) For all $w_1, w_2 \in [\neg\theta]$, $w_1 \leq_\triangle^{\mathbb{E}\circ\theta} w_2$ iff $w_1 \leq_\triangle^{\mathbb{E}} w_2$
(3) For all $w_1, w_2 \in \mathcal{W}$, if $w_1 \in [\theta], w_2 \in [\neg\theta]$ and $w_1 <_\triangle^{\mathbb{E}} w_2$, then $w_1 <_\triangle^{\mathbb{E}\circ\theta} w_2$
(4) For all $w_1, w_2 \in \mathcal{W}$, if $w_1 \in [\theta], w_2 \in [\neg\theta]$ and $w_1 \leq_\triangle^{\mathbb{E}} w_2$, then $w_1 \leq_\triangle^{\mathbb{E}\circ\theta} w_2$

It is then easy to see:

Proposition 13. *The postulates* **(M1)**–**(M4)** *imply* **(C1′)$_\triangle$**, **(C2′)$_\triangle$**, **(C3)$_\triangle$** *and* **(C4)$_\triangle$** *respectively.*

Now suppose the agent receives a regular input θ followed by a core input ϕ. Then our next property says θ should be in the regular beliefs as long as it is consistent with the resultant core beliefs.

(M5) If $\neg\theta \notin \blacktriangle(\mathbb{E} \circ \theta \bullet \phi)$ then $\theta \in \triangle(\mathbb{E} \circ \theta \bullet \phi)$ (Retro-success)

Moreover this should hold even if θ is *not* believed after the first revision. Thus this property allows a regular input which may initially have been rejected to be accepted *retrospectively* as long as it does not contradict the new core beliefs. To roughly illustrate, let us return to Example 1 from the introduction. Suppose your initial epistemic state is $\mathbb{E}$ and let θ and ϕ stand for the information "your son ate lunch at school with King Gustav" and "King Gustav was at a local school" respectively. Assume that your expectations of the King's behaviour are such that you very strongly believe that, of all the places where the King might be, a local school is not one of them, i.e., $\neg\phi \in \blacktriangle(\mathbb{E})$. You also strongly believe that no-one can be in two places at the same time, i.e., $\neg\phi \rightarrow \neg\theta \in \blacktriangle(\mathbb{E})$, and so you deduce $\neg\theta \in \blacktriangle(\mathbb{E})$. Assuming you always treat information provided to you by your son as a regular input, your epistemic state after your son provides you with θ is $\mathbb{E} \circ \theta$. However, since $\neg\theta \in \blacktriangle(\mathbb{E})$, we have $\theta \notin \triangle(\mathbb{E} \circ \theta) = \triangle(\mathbb{E})$. Now suppose you receive the information ϕ from the TV news. Since you take the TV news to be highly reliable, you treat this as a core input, and so your epistemic state after this input is $\mathbb{E} \circ \theta \bullet \phi$. Now, because $\neg\phi$ is necessarily removed from the core at this point, your grounds for deducing that $\neg\theta$ is a core belief have been taken away. If this was really your *only* argument for deducing $\neg\theta \in \blacktriangle(\mathbb{E})$ then you should now have $\neg\theta \notin \blacktriangle(\mathbb{E} \circ \theta \bullet \phi)$. **(M5)** says that, in this case, $\theta \in \triangle(\mathbb{E} \circ \theta \bullet \phi)$, i.e., you should now believe your son's information.

In terms of pre-orders **(M5)** corresponds to the following property:

Proposition 14. *$\langle \circ, \bullet \rangle$ satisfies* **(M5)** *iff, for all $\mathbb{E} \in Ep$ and $\theta \in L^+$ and for all $w_1, w_2 \in \mathcal{W}$ such that $w_1 \in [\neg\theta]$ and $w_2 \in [\theta]$, we have $w_1 \leq_{\triangle}^{\mathbb{E}\circ\theta} w_2$ implies $w_1 <_{\blacktriangle}^{\mathbb{E}\circ\theta} w_2$.*

The above property says that after a regular revision by θ, and for each $[\theta]$-world w_2, the only $[\neg\theta]$-worlds considered at least as plausible as w_2 by the new regular pre-order are those considered *strictly* more plausible than w_2 by the new core pre-order. (Recall that we always have $\leq_{\triangle}^{\mathbb{E}\circ\theta} \subseteq \leq_{\blacktriangle}^{\mathbb{E}\circ\theta}$ and so we necessarily have that $w_1 <_{\blacktriangle}^{\mathbb{E}\circ\theta} w_2$ implies $w_1 \leq_{\triangle}^{\mathbb{E}\circ\theta} w_2$.) The following proposition reveals **(M5)** to be quite a strong property.

Proposition 15. *Let $\langle \circ, \bullet \rangle$ be a revision system which satisfies* **(M5)***. Then $\circ$ satisfies* $\mathbf{(C5')}_{\triangle}$*. If, in addition, $\langle \circ, \bullet \rangle$ satisfies* **(M0)** *then $\langle \circ, \bullet \rangle$ also satisfies* **(M3)** *and* **(M4)***.*

Our next postulate is inspired by similar considerations to those behind the AGM revision postulate **(K*6)** (see [10]). Suppose an agent receives a core input θ followed by a regular input ϕ. Then if ϕ is consistent with the core beliefs *after* the first revision, then the regular belief set after the second revision should be just the same as if the agent had received θ and ϕ *together* as a core input.

(M6) If $\neg\phi \notin \blacktriangle(\mathbb{E} \bullet \theta)$ then $\triangle(\mathbb{E} \bullet \theta \circ \phi) = \triangle(\mathbb{E} \bullet (\theta \wedge \phi))$ (Cross-Conjunction 2)

It is easy to see that, for a revision system, this postulate implies **(Y2)** (i.e., (Cross-Conjunction 1)). In terms of pre-orders **(M6)** can be understood as follows:

Proposition 16. *$\langle \circ, \bullet \rangle$ satisfies* **(M6)** *iff, for all $\mathbb{E} \in Ep$ and $\theta \in L^+$ and for all $w_1, w_2 \in [\blacktriangle(\mathbb{E} \bullet \theta)]$, we have $w_1 \leq_{\triangle}^{\mathbb{E}\bullet\theta} w_2$ iff $w_1 \leq_{\triangle}^{\mathbb{E}} w_2$.*

The above condition places a constraint on the new regular pre-order $\leq_{\triangle}^{\mathbb{E}\bullet\theta}$ which follows a core input θ. It says that within $[\blacktriangle(\mathbb{E} \bullet \theta)]$ the old regular ordering $\leq_{\triangle}^{\mathbb{E}}$ is preserved. It is quite easy to see that if $\leq_{\triangle}^{\mathbb{E}\bullet\theta}$ is taken to be the lexicographic refinement of $\leq_{\blacktriangle}^{\mathbb{E}\bullet\theta}$ by $\leq_{\triangle}^{\mathbb{E}}$ as in Proposition 8 then this constraint is satisfied. Thus, using propositions 8 and 16, we see that

Proposition 17. *If $\bullet$ satisfies* **(Z1)** *then $\langle \circ, \bullet \rangle$ satisfies* **(M6)**.

Our last property is the regular-belief analogue of **(Comm)**$_\blacktriangle$:

(Comm)$_\triangle$ $\triangle(\mathbb{E} \circ \theta \bullet \phi) = \triangle(\mathbb{E} \bullet \phi \circ \theta)$

Rather surprisingly it turns out that, in the presence of **(M0)**, this rule is equivalent to the conjunction of some of our previous postulates:

Proposition 18. *Let $\langle \circ, \bullet \rangle$ be a revision system which satisfies* **(M0)**. *Then $\langle \circ, \bullet \rangle$ satisfies* **(Comm)**$_\triangle$ *iff $\langle \circ, \bullet \rangle$ satisfies* **(M1)**, **(M2)**, **(M5)** *and* **(M6)**.

We remark that the derivations of **(M1)** and **(M2)** from **(Comm)**$_\triangle$ do not require **(M0)**. Figure 1 summarises the postulates from this section, together with their inferential interrelations and their relations with some of our previous postulates. A dashed line indicates a derivation which requires the presence of the postulate **(M0)**.

6 A Construction

We now give an explicit construction of a pair of operators $\circ$ and $\bullet$ which display much of the behaviour described in the previous sections. For this construction we will use a specific representation of an epistemic frame. For the set of epistemic states the basic idea is that the agent keeps two separate lists of sentences, one which records all the regular inputs he receives, and one which records all the core inputs he receives. We take an agent's epistemic state simply *to be* this record. More precisely we take for the set of epistemic states the set $seq(L^+)^2$ of all pairs of finite sequences of consistent sentences $\mathcal{E} = \langle (\beta_1, \ldots, \beta_m), (\alpha_1, \ldots, \alpha_n) \rangle$, where the β_i are the regular inputs which the agent has so far received, and the α_i are the core inputs.[7] Revision of such an epistemic state is then a trivial affair

[7] In the literature on iterated AGM revision, the idea of taking an epistemic state to be a *single* sequence of sentences reflecting the revision history has already been suggested in, e.g., [6, 17, 18].

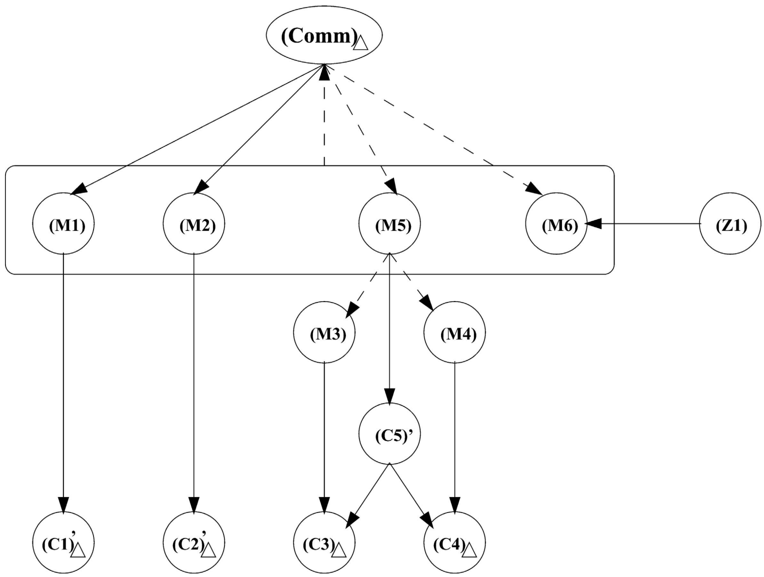

Fig. 1. The inferential interrelations of the postulates for a revision system

– we obtain our new epistemic state by simply adding the sentence received as input to the end of the appropriate list, depending on whether the sentence is received as a regular input or a core input. More precisely, given $\mathcal{E} = \langle \boldsymbol{b}, \boldsymbol{a} \rangle$ where $\boldsymbol{b} = (\beta_1, \ldots, \beta_m)$ and $\boldsymbol{a} = (\alpha_1, \ldots, \alpha_n)$, we define

$$\mathcal{E} \circ \phi = \langle \boldsymbol{b} \cdot \phi,\ \boldsymbol{a} \rangle \quad \text{and} \quad \mathcal{E} \bullet \phi = \langle \boldsymbol{b},\ \boldsymbol{a} \cdot \phi \rangle$$

where $\cdot$ denotes sequence concatenation. The properties of the revision operators are then determined entirely by the particular ways we choose to extract the belief set $\triangle(\mathcal{E})$ and the core belief set $\blacktriangle(\mathcal{E})$ from any given $\mathcal{E} \in seq(L^+)^2$.[8] However there are some things we can say immediately about $\bullet$ and $\circ$ without referring to $\blacktriangle(\mathcal{E})$ or $\triangle(\mathcal{E})$, for clearly we have $\mathcal{E} \circ \theta \bullet \phi = \langle \boldsymbol{b} \cdot \theta,\ \boldsymbol{a} \cdot \phi \rangle = \mathcal{E} \bullet \phi \circ \theta$. Hence we already have:

Proposition 19. $\bullet$ *and* $\circ$ *defined above together satisfy the rules* **(Comm)**$_\blacktriangle$ *and* **(Comm)**$_\triangle$.

Hence, once we have shown that $\langle \circ, \bullet \rangle$ is a revision system on the epistemic frame $\langle seq(L^+)^2, \triangle, \blacktriangle \rangle$, it will immediately follow that $\langle \circ, \bullet \rangle$ satisfies many of the postulates we have considered, including *all* those of the previous section (cf. Fig. 1 and recall that, for a revision system, **(Comm)**$_\blacktriangle$ is equivalent to **(M0)**).

[8] The general approach of leaving all the work in performing revision to some operation of *retrieval* on epistemic states is known as the *vertical* approach to belief revision [22].

Now let's determine $\triangle(\mathcal{E})$ and $\blacktriangle(\mathcal{E})$. Here we again assume $\mathcal{E}$ is of the form $\langle(\beta_1,\ldots,\beta_m),(\alpha_1,\ldots,\alpha_n)\rangle$. Turning first to $\blacktriangle(\mathcal{E})$, we first form an increasing sequence of sets of sentences Γ_i starting with $\Gamma_0=\emptyset$ by setting, for each $i=1,\ldots,n$,

$$\Gamma_i=\begin{cases}\Gamma_{i-1}\cup\{\alpha_{n+1-i}\} & \text{if this is consistent}\\ \Gamma_{i-1} & \text{otherwise}\end{cases}$$

In other words we form the set of sentences Γ_n by taking the sequence $(\alpha_1,\ldots,\alpha_n)$ and, starting at the end with α_n, working our way backwards through the sequence and adding each sentence as we go, provided it is consistent with the sentences already collected up to that point. We then define

$$\blacktriangle(\mathcal{E})=Cn(\Gamma_n).$$

Note that $\blacktriangle(\mathcal{E})$ so defined is bound to be consistent. Also note that in the construction of $\blacktriangle(\mathcal{E})$ no mention is made of the first sequence $(\beta_1,\ldots,\beta_m)$ of $\mathcal{E}$, i.e., $\blacktriangle(\mathcal{E})$ depends only the second sequence $(\alpha_1,\ldots,\alpha_n)$. This is all that is needed to show:

Proposition 20. $\circ$ *satisfies* **(X1)** *and* **(X2)***.*

We also have

Proposition 21. $\bullet$ *is a core revision operator (on the epistemic frame* $\langle seq(L^+)^2,\triangle,\blacktriangle\rangle$*) according to Definition 4. Furthermore* $\bullet$ *satisfies* **(C1)**$_\blacktriangle$*,* **(C2)**$_\blacktriangle$ *and* **(C5)**$_\blacktriangle$*.*

We now turn to our definition of $\triangle(\mathcal{E})$. This is a matter of just picking up from where we left off in the construction of $\blacktriangle(\mathcal{E})$. We define an increasing sequence of sets of sentences Θ_i, starting with $\Theta_0=\Gamma_n$, by setting, for $i=1,\ldots,m$,

$$\Theta_i=\begin{cases}\Theta_{i-1}\cup\{\beta_{m+1-i}\} & \text{if this is consistent}\\ \Theta_{i-1} & \text{otherwise}\end{cases}$$

We then set

$$\triangle(\mathcal{E})=Cn(\Theta_m).$$

Clearly, since $\Gamma_n\subseteq\Theta_m$, we get $\blacktriangle(\mathcal{E})\subseteq\triangle(\mathcal{E})$. Also since Γ_n is consistent we also have $\triangle(\mathcal{E})$ is consistent. We would now like to show that $\circ$ forms a core-invariant regular revision operator on the epistemic frame $\langle seq(L^+)^2,\triangle,\blacktriangle\rangle$. As a first step we show the following:

Proposition 22. *For all* $\mathcal{E}\in seq(L^+)^2$ *and all* $\phi\in L^+$ *we have*

$$\triangle(\mathcal{E}\circ\phi)=\begin{cases}\triangle(\mathcal{E}\bullet\phi) & \textit{if } \neg\phi\notin\blacktriangle(\mathcal{E})\\ \triangle(\mathcal{E}) & \textit{otherwise}\end{cases}$$

(In particular $\circ$ *and* $\bullet$ *together satisfy* **(Y1)***.)*

Next we have the following result:

Proposition 23. *For each $\mathcal{E} \in seq(L^+)^2$ there exists some simple AGM revision function $*$ for $\triangle(\mathcal{E})$ such that $\triangle(\mathcal{E} \bullet \phi) = \triangle(\mathcal{E}) * \phi$ for all $\phi \in L^+$.*

Combining propositions 20, 22 and 23 then allows us to prove:

Proposition 24. *$\circ$ is a core-invariant regular revision operator (on the epistemic frame $\langle seq(L^+)^2, \triangle, \blacktriangle\rangle$) according to Definitions 2 and 3.*

Meanwhile, as a corollary to Proposition 23 we have:

Corollary 1. *$\bullet$ satisfies* **(Y2)** *and* **(Y3)**.

We remark, however, that $\bullet$ does *not* satisfy **(sY3)**, i.e., core-revising $\mathcal{E}$ by a core-contravening sentence will not necessarily erase the distinctions between core belief set and regular belief set. To see this let $\mathcal{E} = \langle (p \wedge q), (\neg q)\rangle$. Then $\blacktriangle(\mathcal{E}) = \triangle(\mathcal{E}) = Cn(\neg q)$. Consider core-revising by q. Then obviously $\neg q \in \blacktriangle(\mathcal{E})$ but $\mathcal{E} \bullet q = \langle (p \wedge q), (\neg q, q)\rangle$, leading to $\blacktriangle(\mathcal{E} \bullet q) = Cn(q) \neq Cn(p \wedge q) = \triangle(\mathcal{E} \bullet q)$.

We are now finally in a position to show:

Theorem 3. *The pair $\langle \circ, \bullet\rangle$ forms a revision system (on the epistemic frame $\langle seq(L^+)^2, \triangle, \blacktriangle\rangle$).*

As a final piece in our jigsaw we have the following result regarding the behaviour of $\triangle(\mathcal{E})$ under iterated core revision.

Proposition 25. *$\bullet$ satisfies* **(Z1)**. *Hence $\bullet$ satisfies* $\mathbf{(C1')}_\triangle$, $\mathbf{(C2')}_\triangle$ *and* $\mathbf{(C5')}_\triangle$ *(with $\bullet$ in place of $\circ$).*

The second part of this proposition follows from propositions 10 and 21.

7 Conclusion

We have taken a close look at iterated non-prioritised revision, using as a starting point a basic model of non-prioritised revision which makes use of a set of core beliefs amongst the set of regular beliefs. We considered two types of revision operator on epistemic states: a normal, regular revision corresponding to a direction to include a given sentence in the regular beliefs, and a core revision operator. We presented some postulates for the iteration of both these operators, including some for the particularly interesting case of mixed iterated sequences of revisions consisting of both types of revision operations. In many cases we have shown how these postulates correspond to conditions on the dynamics of the plausibility orderings on worlds which underlie an agent's epistemic state. Finally we provided a construction which illustrated some of the ideas.

As further work we would like to examine also operations of contraction in this context. Of particular interest would be an operation of "core contraction" in which a sentence is removed from the core beliefs. What effect should such an operation have on the *regular* belief set? For instance, should the sentence removed from the core be retained as a regular belief? Also, in this paper we considered revision systems consisting of just *two* revision operators $\langle \circ, \bullet\rangle$ with $\circ$

being, in a sense, *dominated* by $\bullet$. A possible extension would be to consider an extended *family* of revision operators $\langle \circ_1, \circ_2, \ldots, \circ_n \rangle$ of increasing "strength", perhaps with each $\circ_i$ corresponding to a different source of information. Finally we would like to investigate other natural ways of constructing a revision system $\langle \circ, \bullet \rangle$ which perhaps do not satisfy some of the stronger postulates considered here (such as $\mathbf{(Comm)}_\triangle$).

Acknowledgements

This work is supported by the DFG research project "Computationale Dialektik". The motivation to begin the work came from some interesting discussions with Emil Weydert. Thanks are also due to Gerhard Brewka for his general support and to the anonymous referees for their detailed comments.

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. R. Booth, A negotiation-style framework for non-prioritised revision, in: *Proceedings of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2001)* (2001) 137–150.
3. C. Boutilier, Iterated revision and minimal change of conditional beliefs, *Journal of Philosophical Logic* **25** (1996) 262–305.
4. C. Boutilier and M. Goldszmidt, Revision by conditional beliefs, in: *Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI'93)* (1993) 649–654.
5. J. Cantwell, Some logics of iterated belief change, *Studia Logica* **63** (1999) 49–84.
6. S. Chopra, K. Georgatos and R. Parikh, Relevance sensitive non-monotonic inference on belief sequences, *Journal of Applied Non-Classical Logics* **11**(1-2) (2001) 131–150.
7. S. Chopra, A. K. Ghose and T. A. Meyer, Iterated revision and recovery: a unified treatment via epistemic states, in: *Proceedings of the Fifteenth European Conference on Artificial Intelligence (ECAI 2002)* (2002) 541–545.
8. A. Darwiche and J. Pearl, On the logic of iterated belief revision, *Artificial Intelligence* **89** (1997) 1–29.
9. E. Fermé, *Revising the AGM Postulates*, PhD Thesis, Universidad de Buenos Aires (1999).
10. P. Gärdenfors, *Knowledge in Flux*, MIT Press (1988).
11. A. Grove, Two modellings for theory change, *Journal of Philosophical Logic*, **17** (1988) 157–170.
12. S. O. Hansson, A survey of non-prioritized belief revision, *Erkenntnis* **50** (1999) 413–427.
13. S. O. Hansson, E. Fermé, J. Cantwell and M. Falappa, Credibility-limited revision, *Journal of Symbolic Logic* **66**(4) (2001) 1581–1596.
14. A. Herzig, S. Konieczny and L. Perrussel, On iterated revision in the AGM framework, in: *Proceedings of the Seventh European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'03)* (2003).

15. H. Katsuno and A. O. Mendelzon, Propositional knowledge base revision and minimal change, *Artificial Intelligence* **52** (1991) 263–294.
16. G. Kern-Isberner, Postulates for conditional belief revision, in: *Proceedings of the Sixteenth International Conference on Artificial Intelligence (IJCAI'99)* (1999) 186–191.
17. S. Konieczny and R. Pino Pérez, A framework for iterated revision, *Journal of Applied Non-Classical Logics* **10**(3-4) (2000) 339–367.
18. D. Lehmann, Belief revision, revised, in: *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI'95)* (1995) 1534–1540.
19. D. Makinson, Screened revision, *Theoria* **63** (1997) 14–23.
20. A. C. Nayak, Iterated belief change based on epistemic entrenchment, *Erkenntnis* **41** (1994) 353–390.
21. A. C. Nayak, M. Pagnucco and P. Peppas, Dynamic belief revision operators, *Artificial Intelligence* **146** (2003) 193–228.
22. H. Rott, *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*, Oxford University Press (2001).