Changing Conditional Beliefs Unconditionally*

Abhaya C. Nayak[1], Norman Y. Foo[1], Maurice Pagnucco[1] and Abdul Sattar[2]

[1]Knowledge Systems Group
Basser Dept. of Computer Science (F09)
University of Sydney
NSW 2006, Australia
email: {abhaya|norman|morri}@cs.su.oz.au

[2]School of Computing and Info. Technology
Griffith University, Brisbane
QLD 4111, Australia
email: A.Sattar@cit.gu.edu.au

Abstract

Although the AGM account of belief change tells us how to change *un*conditional beliefs, it fails to guide us in changing our conditional beliefs. That explains why, in the AGM account of belief change proper, a decent account of iterated belief change is not forthcoming. Darwiche and Pearl provide an account of changing conditional beliefs. We argue that the Darwiche-Pearl postulates are, on the one hand, limited in scope and, on the other, excessively strong and suggest how they should be supplemented. We show, contrary to the generally held view, that the revision operation changes to a new (though in some cases the same) revision operation after each instance of belief change. Finally, we prove what may be viewed as representation results connecting the acount offered in this paper with entrenchment kinematics.

Keywords: Conditionals, iterated belief change, entrenchment kinematics.

Suppose an agent currently believes that Skippy is an elephant. If the agent were to learn from a credible source that Skippy is actually not an elephant, but a kangaroo, then she would infer that Skippy jumps. This fact may be explained by ascribing to the agent a conditional belief `kangaroo_skippy > jumps_skippy`. (Note the subjunctive nature of the conditional $>$.) This way of determining which conditionals are accepted by an agent in a given belief state is originally due to Frank Ramsey [11]. In some cases of belief change, belief in such conditionals will be carried over to the next belief state, whereas, in others, the evidence will nullify such conditional beliefs.

It is only by learning what conditional beliefs are carried over to the new belief state from the old, and what conditional beliefs are acquired anew that we can compute the result of further revising the current (already revised) belief state. Such questions are currently attracting considerable interest from the area of Artificial Intelligence and have important implications for updating (logical) databases.

In §1 we discuss how the AGM account of belief change [1, 4], through its lack of concern for change in conditional beliefs, fails to provide an account of iterated belief change. We also show in this section how the Darwiche-Pearl (D-P) account [3] attempts, but fails, to do so. In §2 we provide an extension of a modified D-P account, based on a dynamic belief revision operation, which is successful in this task. In §3 we characterize this new framework via entrenchment kinematics [10].

* To appear in *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge*, De Zeeuwse Stromen, The Netherlands, March 1996.

1 Background

In the AGM approach, the object language (the language in which the beliefs of an agent are represented) is a propositional language $\mathcal{L}$ closed under the usual connectives $\neg$, $\rightarrow$, $\leftrightarrow$, $\wedge$ and $\vee$. Two eminent members of $\mathcal{L}$ are $\top$ (Truth) and $\bot$ (Falsity). The underlying logic (the logic of the agent) is represented by its consequence operation Cn which satisfies the following conditions: For all sets of sentences Γ, Γ' and all sentences x and y,

Inclusion:	$\Gamma \subseteq Cn(\Gamma)$
Iteration:	$Cn(Cn(\Gamma)) = Cn(\Gamma)$
Monotonicity:	$Cn(\Gamma) \subseteq Cn(\Gamma')$ whenever $\Gamma \subseteq \Gamma'$
Supraclassicality:	$x \in Cn(\Gamma)$ if Γ classically implies x
Deduction:	$y \in Cn(\Gamma \cup \{x\})$ iff $(x \rightarrow y) \in Cn(\Gamma)$
Compactness:	If $x \in Cn(\Gamma)$ then $x \in Cn(\Gamma')$ for some finite $\Gamma' \subseteq \Gamma$.

We often write $\Gamma \vdash x$ for $x \in Cn(\Gamma)$.

Any set of sentences $K \subseteq \mathcal{L}$ closed under Cn is called a belief set (or theory) and represents a possible belief state. We may interpret the members of K to be the sentences that the agent holds as unconditional beliefs. By **K** we denote the set of all possible belief sets (in $\mathcal{L}$). We use $K_\bot$ to denote the absurd belief set $\mathcal{L}$.

At least three forms of epistemic change are recognized in the AGM framework. These changes are represented by three operations $+, -, * :$ $\mathbf{K} \times \mathcal{L} \rightarrow \mathbf{K}$ that, for every belief set and a sentential input, returns a belief set. These are, respectively, called the belief expansion, belief contraction and belief revision operations. Since members of **K** are interpreted as belief states, these operations may be viewed as state transition functions. In this paper we confine our discussion only to belief revision ($*$).

Note that the symbol $>$ is not a connective in $\mathcal{L}$. So the conditional beliefs that are of interest are not members of any AGM belief set K. However, whether an agent in state K holds a conditional belief $x > y$ can be determined by finding out whether or not $y \in K^*_x$. The same procedure can be used to find out whether old conditional beliefs are lost or new conditional beliefs are acquired through acquiring new (unconditional) information.

We state below the constraints that the revision operation is required to satisfy in the AGM framework. Motivation and interpretation of these constraints can be found in [4].

(1^*)	K^*_x is a theory	(Closure)
(2^*)	$x \in K^*_x$	(Success)
(3^*)	$K^*_x \subseteq Cn(K \cup \{x\})$	(Inclusion)
(4^*)	If $K \not\vdash \neg x$ then $Cn(K \cup \{x\}) \subseteq K^*_x$	(Preservation)
(5^*)	$K^*_x = K_\bot$ iff $\vdash \neg x$	(Consistency)
(6^*)	If $\vdash x \leftrightarrow y$, then $K^*_x = K^*_y$	(Extensionality)
(7^*)	$K^*_{(x \wedge y)} \subseteq Cn(K^*_x \cup \{y\})$	(Super-expansion)
(8^*)	If $\neg y \notin K^*_x$ then $Cn(K^*_x \cup \{y\}) \subseteq K^*_{(x \wedge y)}$	(Sub-expansion)

It is easily noticed that the AGM system tells us precious little about iterated belief change. The only interesting inference about iterated belief change that we can draw from the AGM postulates is

- **AGM-It:** If $\neg y \notin K^*_x$ then $(K^*_x)^*_y = K^*_{x \wedge y}$.

(This follows primarily from 4*, 7* and 8*.) But the AGM system does not constrain iterated belief change in any manner when $\neg y \in K^*_x$. It has occasionally been interpreted as the AGM system permitting, when $\neg y \in K^*_x$, all possible iterated belief changes consistent with 1*-6* (see, for example, [3]). In other words, in the envisioned situation, $(K^*_x)^*_y$ may be equated with *any* belief set K' such that $y \in K'$ without violating the AGM constraints.

In order to alleviate this situation, Darwiche and Pearl [3] impose four further constraints on the AGM revision operation $*$. Their constraints are couched in the K-M formalism [7] which assumes that each belief set can be expressed as a single sentence. We will present these constraints in the AGM terminology instead. The extra constraints they propose are the following (DP1-DP4):

DP1: If $y \vdash x$ then $(K^*_x)^*_y = K^*_y$
DP2: If $y \vdash \neg x$ then $(K^*_x)^*_y = K^*_y$
DP3: If $x \in K^*_y$ then $x \in (K^*_x)^*_y$
DP4: If $\neg x \notin K^*_y$ then $\neg x \notin (K^*_x)^*_y$

There are two major problems with the D-P postulates. On the one hand, they are too limited in scope to assist in iterated belief change in some very common situations. On the other hand, they are excessively strong; in particular, DP2 is inconsistent with the AGM postulates of belief revision.

In order to see that the DP postulates are too limited in scope, consider the following scenario. Our agent believes that Tweety is a singing bird. However, since there is no strong correlation between singing and birdhood, the agent is prepared to retain the belief that Tweety sings even after accepting the information that Tweety is not a bird, and conversely, if the agent were to be informed that Tweety does not sing, she would still retain the belief that Tweety is a bird. Imagine that the agent first receives the information that Tweety is not a bird, then the information that Tweety does not sing. On such an occasion, it is reasonable to assume that the agent should believe that Tweety is a non-singing non-bird. However, this is not guaranteed by the Darwiche-Pearl account. This is so, because, as easily seen, AGM-It (if $\neg y \notin K^*_x$ then $(K^*_x)^*_y = K^*_{x \wedge y}$) and DP1-DP4 are inapplicable in the above case.

Freund and Lehmann [5] have shown that DP2 conflicts with the AGM postulates. Even weakening DP2 in the following manner does not help:[1]

- DP2$'$. If both $\not\vdash \neg x$ and $y \vdash \neg x$ then $(K^*_x)^*_y = K^*_y$.

For instance, consider $K = Cn(\emptyset)$ and $K' = Cn(\{x\})$, and let x and y be atomic sentences. It follows from 1*-6* that $K'^*_{\neg y} = K^*_{x \wedge \neg y}$. Hence $(K'^*_{\neg y})^*_y = (K^*_{x \wedge \neg y})^*_y$. Applying DP2$'$ to this equality, we get $K'^*_y = K^*_y$, i.e., $Cn(\{x, y\}) = Cn(\{y\})$. Contradiction! ■

Such considerations have led Lehmann [8] to further weaken DP2. In fact, it is possible to show that Lehmann's proposal is a special case of William's [13] adjustment operation.[2] Our

[1]We are thankful to Daniel Lehmann for pointing this out in personal communication.

[2]Lehmann's [8] *Widening Ranked Model* can be captured via William's [13] $\langle x, \alpha \rangle$-adjustment. Readers acquainted with the semantics of adjustment (see the principle "OCF-Adjs" in [9], p. 364) will easily verify that $\langle x, \alpha \rangle$-adjustment, where α is the next non-empty rank greater than $rank(x)$, corresponds to revision by x in a widening ranked model.

aim, on the other hand, is to adopt the DP framework in spirit and to extend it, without courting inconsistency, so as to handle iterated belief change in all situations. We do that in the next section.

2 An Extension of the D-P Framework

There are primarily two reasons why a complete specification of iterated belief change is lacking in the D-P framework. (1) Although inconsistent belief sets are not beyond the reach of an agent, no account is given of how one may revise one's beliefs once she is in $K_\perp$. (2) None of the postulates are applicable when we compute $(K^*_x)^*_y$, if x and y are mutually consistent, $\neg x \in K^*_y$ and $\neg y \in K^*_x$.

In order to amend the situation with respect to problem (1), we propose the following postulate:

- **Absurdity:** $(K_\perp)^*_x = Cn(\{x\})$ for any sentence x.

This postulate tells us that once the agent has plunged into logical inconsistency, she loses all the genuine information she had acquired so far and starts her epistemic life anew.

As to problem (2), consider the following example. Let us suppose that all that an agent knows about Kim is that it is a living thing. If the agent were to learn that x : `mammal_kim`, she would infer that $\neg y$: ¬`fly_kim` since very few mammals fly. On the other hand, if the agent were to learn that y : `fly_kim`, the agent would infer that $\neg x$: ¬`mammal_kim`, since most flyers are non-mammals. Suppose that the agent first learns that Kim is a mammal, and then that Kim flies. In this case, because of the evident underlying conflict, the agent might be inclined to check the veracity of these two pieces of evidence; but given that the evidence in question is as good as the AGM system requires (cf. postulate of Success), we submit that the agent should infer that Kim is a flying mammal (such as a bat). So the subsequent evidence (y) does not override the previous evidence x.

This example suggests that when $x \wedge y \not\vdash \perp$, the sentence x is in $(K^*_x)^*_y$ even if $\neg x \in K^*_y$ and $\neg y \in K^*_x$. If we allow this principle, then it is reasonable to allow the seemingly weaker principle that the sentence x is in $(K^*_x)^*_y$ if $x \wedge y \not\vdash \neg x$, $\neg x \notin K^*_y$ and $\neg y \in K^*_x$ (for in this case, y is less "in conflict" with x given the background knowledge K, than it is in the earlier case). From these two principles, together with **AGM-It** and DP3, follows the following principle:

- **#Recalcitrance:** If $x \wedge y \not\vdash \perp$ then $x \in (K^*_x)^*_y$.

However, the following result shows that there is no revision operation $*$ satisfying the AGM postulates and #Recalcitrance.

Theorem 1 *No belief revision operation $*$ that satisfies the AGM postulates satisfies #Recalcitrance.*

This result has two alternative explanations: (1) #Recalcitrance is too strong. We should reject/weaken it in a judicious manner. (2) When we revise our current belief set, it is not only the belief set that changes; the accompanying revision operation also undergoes modification.

The revision operation used to revise prior belief set K and the revision operation used to revise the posterior belief set K^*_x are, in general, different.

Of the above two explanations, we adopt the second one. We will present three different arguments in support of our position, namely: (a) If this impossibility result is reason enough to discard the account of belief change we are proposing, then there are more or less equally compelling reasons to discard competing accounts of belief change such as Boutilier's [2]; (b) the notion of a fixed belief revision operation is unintuitive and overly restrictive; and (c) literature in the area supports the view that a revision mechanism changes when there is an act of belief change.

First we demonstrate that if our account is to be rejected on the basis of this impossibility result, then so should the account of iterated belief change proposed by Boutilier which has drawn significant attention. It is not necessary to go into the detailed structure of Boutilier's account; all we need is the following consequence of his assumptions (p.524, Theorem 7, [2]):

Theorem 2 (Boutilier 1993) *Let $A_1, \ldots, A_n$ be a revision sequence with one incompatible update A_n. Then $((K^*_{A_1})^*_{A_2} \ldots)^*_{A_n} = K^*_{A_1 \wedge \ldots A_k \wedge A_n}$ where A_k is the most recent compatible update for A_n. If there is no such A_k then $((K^*_{A_1})^*_{A_2} \ldots)^*_{A_n} = K^*_{A_n}$.*

By saying that A_n is the only incompatible update, he means that each A_i $(1 < i < n)$ is consistent with $K^*_{A_{i-1}}$ but A_n is inconsistent with $K^*_{A_{n-1}}$. The following observation, is a consequence of the theorem (2).

Observation 1 *Let $*$ be a fixed revision operation satisfying Boutilier constraints; let $x \in K$ such that $x \vdash y$. Then, $\neg y \in K^*\neg x$.*

The behavior of a revision operation $*$, as depicted by this observation, is very disturbing. Consider an agent who believes that Cleopatra had a son (and hence a child). According to this result, if the agent were to learn that Cleopatra had no son, then the agent must conclude that Cleopatra had no children. But that is counterintuitive: having no son is not sufficient evidence to infer barrenness. We conclude that in Boutilier's system, iterated revision makes no sense unless it is assumed that the revision operation $*$ changes along with belief change.

Next we give a commonsensical account of why the notion of a fixed revision operation is counterintuitive. It is a general AGM assumption that a revision operation $*$ determines a unique belief set $K' = K^*_x$ given any belief set K and any evidential statement x. If this is so, then an agent would have no need to change her revision operation $*$. It, however, has the counterintuitive consequence that *whenever* the agent has the belief set K and accepts the evidence x, she will invariably end up in the same revised belief set K'. Intuitively the firmness of beliefs plays a crucial role in determining the revised belief set. It is conceivable that at two different times, t_1 and t_2, an agent has the same set of beliefs but the relative firmness of the beliefs are different. If the agent accepts the same evidence at t_1 and t_2, the resultant belief sets would be different. But this cannot be accommodated if we require the belief revision operation to be fixed once and for all. Hence, we have an a priori case against a fixed belief revision operation $*$.

Finally we mention the work in the literature that supports our view. Spohn [12], Nayak [9], Nayak *et al.* [10] and Williams[13], among others, study how the beliefs and their firmness change during belief change. Since the relative firmness of various beliefs plays a crucial role in determining the behavior of belief revision operations, all these works are indirect evidence that belief revision operations change over time.

In light of the above discussion, we propose that the *prior* revision operation used to revise K and the *posterior* revision operation used to revise K^*_x are not necessarily identical. We denote the former by $*$ as before; the latter is denoted by $*|x$. Accordingly, #Recalcitrance is recast as:

- **Recalcitrance:** If $x \wedge y \nvdash \bot$ then $x \in (K^*_x)^{*|x}_y$.

The other principles about iterated revision, namely AGM-It, DP1-DP4 are similarly modified. It is easily noticed that DP2′, when thus modified, is no longer inconsistent with the AGM postulates. Hence, once we adopt a dynamic belief revision operation, there is no need to weaken DP2′. This, we consider, is further evidence that the belief revision operation is dynamic, and is not fixed once and for all.

Recalcitrance mandates that if x and y are mutually consistent, x is in $(K^*_x)^{*|x}_y$. According to the postulate of Success, y is also in $(K^*_x)^{*|x}_y$. Hence $x \wedge y$ is in $(K^*_x)^{*|x}_y$. It would appear reasonable to assume then, that when x and y are mutually consistent, $(K^*_x)^{*|x}_y$ is actually the result of revising K by $x \wedge y$. Accordingly, we replace Recalcitrance by Conjunction:

- **Conjunction:** If $x \wedge y \nvdash \bot$ then $(K^*_x)^{*|x}_y = K^*_{x\wedge y}$.

This postulate tells us that if two sequentially received pieces of information are consistent with each other, then they may be conjoined together into a single piece of information.

Obviously Conjunction implies Recalcitrance. What is more, if we accept Conjunction, we get AGM-It, DP1, DP3 and DP4 for free:

Observation 2 *In the presence of 1*-6*, Conjunction implies AGM-It, DP1, DP3 and DP4, provided that the second occurrence of $*$ in them are replaced by an occurrence of $*|x$.*

Accordingly, the final list of postulates for the extended D-P framework, presented below for convenience, is quite short:

(0^*)	$(K_\bot)^*_x = Cn(\{x\})$ for any sentence x	(Absurdity)
(1^*-6^*)	As in the AGM	
(7^*new)	If $x \wedge y \nvdash \bot$ then $(K^*_x)^{*|x}_y = K^*_{x\wedge y}$	(Conjunction)
(8^*new)	If $x \wedge y \vdash \bot$ but $\nvdash \neg x$ then $(K^*_x)^{*|x}_y = K^*_y$	(DP2′)

The last two postulates may be viewed as constrains on $|$ rather than on $*$.

3 Character of Entrenchment Kinematics

In this section we connect the account of belief change described in this paper with the account of entrenchment kinematics offered by Nayak *et al.* [10]. They represent a belief state as an epistemic entrenchment (EE) relation which is slightly different from the standard epistemic entrenchment (SEE) relation introduced in AGM [6]. An EE relation, $\preceq$, is defined as any relation over the language $\mathcal{L}$ that satisfies the following four constraints:

$(EE1)$	If $x \preceq y$ and $y \preceq z$ then $x \preceq z$	(transitivity)
$(EE2)$	If $x \vdash y$ then $x \preceq y$	(dominance)
$(EE3)$	For any x and y, $x \preceq x \wedge y$ or $y \preceq x \wedge y$	(conjunctiveness)
$(EE4)$	Given that $\bot \prec z$ for some z, if $y \preceq x$ for all y, then $\vdash x$	(maximality).

The belief set accompanying the EE relation is the epistemic content of this relation, defined as follows:

Definition 1 (Epistemic Content)

$$EC(\preceq) = \begin{cases} \{x \mid \bot \prec x\} & \textit{if } \bot \textit{ is not } \preceq\textit{-maximal} \\ K_\bot & \textit{otherwise} \end{cases}$$

In this account the revision of the EE relation $\preceq$ in light of evidence e is denoted by $\preceq_e^\odot$. Here we equivalently reformulate that definition of $\preceq_e^\odot$ as:

Definition 2 *(***Entrenchment Revision)**
$x \preceq_e^\odot y$ iff any of the following three conditions is satisfied:
(a) $e \vdash \bot$, (b) $e \vdash x,\ e \vdash y,$ AND either $\vdash y$ or both $\not\vdash x$ and $x \preceq y$, (c) $e \not\vdash x$ and $(e \rightarrow x) \preceq (e \rightarrow y)$.

They show that if the epistemic contents of $\preceq$ and $\preceq_e^\odot$ are, respectively, identified with K and K_e^*, then the revision operation $*$ so defined satisfies the AGM postulates of belief revision. We show below that the postulates 0*-8*(new) actually describe the account of iterated belief *revision* provided in [10]. First we present two definitions which are based on analogous definitions in the AGM literature. These definitions provide recipes for constructing an entrenchment relation from a given revision operation and, conversely, for constructing a revision operation from a given entrenchment relation.

Definition 3 (**EE to** $*$)
Let $\preceq$ be an EE relation. Define $*_\preceq$ as: $y \in K_x^{*\preceq}$ iff either $(x \rightarrow \neg y) \prec (x \rightarrow y)$ or $\vdash \neg x$ or $x \vdash y$.

Definition 4 ($*$ **to EE**)
Let $*$ be a belief revision operation and K a belief set. Define $\preceq_{*,K}$ as: $x \preceq_{*,K} y$ iff either $x \notin K^*_{\neg x \vee \neg y}$ or $\vdash x \wedge y$ or $K = K_\bot$.

Given that $*$ satisfies the eight AGM postulates of revision, $\preceq_{*,K}$ is demonstrably an EE relation with epistemic content K. On the other hand, given that $\preceq$ is an EE relation, the operation $*_\preceq$ is also demonstrably a revision operation satisfying the AGM postulates 1*-8*. Now we are in a position to state some formal results.

Theorem 3 *Let e be any arbitrary sentence. Let $*$ and its revision $*|e$ by e be two revision operations such that the postulates 0*-8*(new) are satisfied. Let K be an arbitrary belief set and $K' = K_e^*$. Then $(\preceq_{*,K})_e^\odot = \preceq_{*|e,K'}$*

This result shows that given two revision operations $*$ and $*|e$ satisfying 0*-8*(new), Definition 4 generates from them two entrenchment relations, the latter of which is the revision of the former by the evidence e in accordance with Definition 2.

Theorem 4 *Let $\preceq$ be an EE relation with epistemic content $EC(\preceq) = K$, and $\preceq_e^{\odot}$ the result of revising it by a sentence e. Then the revision operations $* = *_{\preceq}$ and $*|e = *_{\preceq_e^{\odot}}$ satisfy the postulates 0*-8*(new) with respect to the belief set $K = EC(\preceq)$.*

This theorem tells us that given two entrenchment relations $\preceq$ and $\preceq'$, the latter being the revision of the former by evidence e in accordance with Definition 2, we can generate from them revision operations $*$ and $*|e$ with the help of Definition 3 which satisfy postulates 0*-8*(new) *for a fixed belief set* K, namely the epistemic content of $\preceq$. The reason why $*$ and $*|e$ do not necessarily satisfy the postulates for every arbitrary belief set is that the entrenchment relation, being defined for a given belief set (its epistemic content) contains less information than an AGM revision operation which operates on the set of all belief sets.[3]

The two theorems, 3 and 4, together, provide what may be seen as a partial representation result connecting the account we offer in this paper and the account of entrenchment kinematics offered in [10]. Our next two theorems examine what happens when we reconstruct a pair of revision operations (entrenchment relations) from another pair by completing the circle (using both the definitions (3) and (4) consecutively).

Theorem 5 *Let $*$ and $*|e$ be two revision operations satisfying the postulates 0*-8*(new). Let K be an arbitrary belief set and $K' = K_e^*$. Then the revision operations $*_{\preceq_{*,K}}$ and $*_{\preceq_{*|e,K'}}$ are belief revision operations that satisfy 0*-8*(new) with respect to the belief set K.*

As expected, in light of Theorem 4, we see that there is a loss of information when we reconstruct a pair of revision operations from another via an entrenchment relation, since in constructing the intermediary entrenchment relations $\preceq_{*,K}$ and $\preceq_{*|e,K'}$ we lose information. But the interesting fact revealed by this theorem is that the appropriate relation between the reconstituted revision operations hold, albeit only with respect to the belief set K used to construct $\preceq_{*,K}$.

Theorem 6 *Let $\preceq$ be an EE relation with epistemic content $EC(\preceq) = K$, $\preceq_e^{\odot}$ the result of revising $\preceq$ by a sentence e and $K' = EC(\preceq_e^{\odot})$ the epistemic content of $\preceq_e^{\odot}$. Then $\preceq_{*_{\preceq},K} = \preceq$ and $\preceq_{*_{\preceq_e^{\odot}},K'} = \preceq_e^{\odot}$.*

This result is not surprising, since it easily follows from results in the AGM literature, but we list it here only for the sake of completeness in exposition. This completes the formal results obtained in this paper. On the basis of these results, we conclude that the postulates 0*-8*(new) characterize the account of entrenchment kinematics discussed in [10].

[3]Construction of a posterior revision operation that is more general than $*|x$, though easy, is rather tedious, and the study of its behavior is beyond the scope of this paper. One possibility is to construct a suitable function that behaves like $*|x$ when revising K_x^*, and behaves like $*$ when revising any other belief set.

4 Summary And Conclusion

In this paper we have argued that a reasoned account of changing conditional beliefs (or iterated belief change), should supplement the AGM account of belief revision. We have shown that an extension of a slight variation of the Darwiche-Pearl framework, based on a dynamic belief revision operation, is a suitable candidate for this task. Finally we explored the connection between the approach taken in this paper and the account of entrenchment kinematics offered by Nayak *et al.* [10] and proved formal results, including two that may be jointly taken as a partial representation theorem connecting our account with entrenchment kinematics.

Acknowledgements

The authors would like to thank Daniel Lehman, Pavlos Peppas, Rex Kwok, Mary-Anne Williams and Wayne Wobcke whose comments have greatly helped to improve this paper.

References

[1] Alchourrón, C. E., P. Gärdenfors, and D. Makinson, "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions", *Journal of Symbolic Logic,* vol. 50, pp. 510-530, 1985.

[2] Boutilier, C., "Revision Sequences and Nested Conditionals", in *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence,*, R. Bajcsy (Ed), pp. 519-25, 1993.

[3] Darwiche, A. and Pearl, J., "On the Logic of Iterated Belief Revision", in R. Fagin (Ed), *Proceedings of the 1994 Conference on the Theoretical Aspects of Reasoning About Knowledge*, pp. 5-23, 1994.

[4] Gärdenfors, P. *Knowledge in Flux. Modeling the Dynamics of Epistemic States,* Bradford/MIT, Cambridge (Massachusetts), 1988.

[5] Freund, M. and Lehmann, D., "Belief Revision and Rational Inference", Technical Report 94-16, Leibniz Center for Research in Computer Science, Inst. of Computer Science, Hebrew Univ. of Jerusalem.

[6] Gärdenfors, P. and D.Makinson, "Revision of Knowledge Systems using Epistemic Entrenchment," in *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge,* M. Vardi (ed.), Los Altos, CA: Morgan Kaufmann. pp. 83-95, 1988.

[7] Katsuno, H. and A. O. Mendelzon, "On the Difference Between Updating a Knowledge Base and Revising it" in Peter Gärdenfors (ed.) *Belief Revision*, Cambridge University Press, pp. 183-203, 1992.

[8] Lehmann, D. "Belief Revision, Revised", in *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, Montreal, pp. 1534 – 1540, August 1995.

[9] Nayak, A. C. "Iterated Belief Change Based on Epistemic Entrenchment", *Erkenntnis* **41**: pp 353-390, 1994.

[10] Nayak, A., Foo, N., Pagnucco, M. and Sattar, A., "Entrenchment Kinematics 101", *Proceedings of the 7th Australian Joint Conference on Artificial Intelligence,* C. Zhang, J. Debenham and D. Lukose (eds.), World Scientific, pp. 157-164, 1994.

[11] Ramsey, F. P. "General Propositions and Causality", in his *Foundations of Mathematics and Other Logical Essays*, pp. 237-257, 1931.

[12] Spohn, W. "Ordinal Conditional Functions: A Dynamic Theory of Epistemic States", in W. L. Harper and B. Skyrms (eds), *Causation in Decision, Belief Change and Statistics,* vol.2, pp. 105-134, 1988.

[13] Williams, M. A. "Transmutations of Knowledge Systems", in *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning*, J. Doyle and E. Sandewall eds., pp. 619-629, 1994.

A Proofs

NOTATION: $\prec$ (with possible decoration) is used as the strict part of the relation $\preceq$. $x \equiv y$ abbreviates $(x \preceq y) \wedge (y \preceq x)$.

In some of the proofs, we will abuse notation for the sake of readability. We will use $\prec_\alpha^\odot$ to denote the strict part of $\preceq_\alpha^\odot$. This should not be confused with the revision of $\prec$ by e! Similarly, for instance, $(\prec_{*,K})_\alpha^\odot$ is used to denote the strict part of the relation $(\preceq_{*,K})_\alpha^\odot$.

More importantly, $x\ (\npreceq_{*,K})_\alpha^\odot\ y$ will, for instance, abbreviate $\neg(x\ (\preceq_{*,K})_\alpha^\odot\ y)$. The notation $x\ \npreceq_\alpha^\odot\ y$ should be similarly interpreted.

Theorem 1 *No belief revision operation $*$ that satisfies the AGM postulates satisfies #Recalcitrance.*

Proof
Assume, for *reductio*, that there is a belief revision operation $*$ that satisfies the AGM postulates and #Recalcitrance, a belief set K and two sentential inputs x, y such that $x \wedge y \not\vdash \bot$, $x \in K$, but $x \notin K_y^*$.[4] Now consider the resultant belief set $(K_x^*)_y^*$. By #Recalcitrance $x \in (K_x^*)_y^*$. On the other hand, since $x \in K$, by 3* and 4*, we get $K_x^* = K$ whereby $(K_x^*)_y^* = K_y^*$. Thus, since $x \notin K_y^*$, it follows that $x \notin (K_x^*)_y^*$. Contradiction! ■

[4]Although we do not exhibit such a revision operation here, it is not difficult to construct one satisfying these conditions.

Observation 1 *Let $*$ be a fixed revision operation satisfying Boutilier constraints; let $x \in K$ such that $x \vdash y$. Then, $\neg y \in K^*\neg x$.*

Proof
Assume that $x \in K$ and $x \vdash y$; assume also that $*$ is a fixed Boutilier revision operation. Since $x \vdash y$ it follows that $\neg x \in K^*_{\neg y}$ whence, by Theorem 2, $(K^*_{\neg y})^*_x = K^*_x = K$. Hence $((K^*_{\neg y})^*_x)^*_{\neg x} = K^*_{\neg x}$. On the other hand, denote by H the belief set $K^*_{\neg y}$. Then $((K^*_{\neg y})^*_x)^*_{\neg x} = (H^*_x)^*_{\neg x}$. Now, applying Theorem 2 on the RHS, we get $(H^*_x)^*_{\neg x} = H^*_{\neg x}$. However, as noted earlier, already $\neg x \in K^*_{\neg y} = H$ whereby $H^*_{\neg x} = H$. Thus $(K^*_{\neg y})^*_x)^*_{\neg x} = H = K^*_{\neg y}$. Thus we get the identity $K^*_{\neg x} = K^*_{\neg y}$. Obviously, $\neg y \in K^*_{\neg y}$. Hence, $\neg y \in K^*_{\neg x}$. ■

Observation 2 *In the presence of 1*-6*, Conjunction implies AGM-It, DP1, DP3 and DP4, provided that the second occurrence of $*$ in them are replaced by an occurrence of $*|x$.*

Proof
For easy reference, we reproduce Conjunction, AGM-It, DP1, DP3 and DP4 (with the required modification) below:

AGM-It:	If $\neg y \notin K^*_x$, then $(K^*_x)^{*\|x}_y = K^*_{x\wedge y}$.
DP1:	If $y \vdash x$ then $(K^*_x)^{*\|x}_y = K^*_y$
DP3:	If $x \in K^*_y$ then $x \in (K^*_x)^{*\|x}_y$
DP4:	If $\neg x \notin K^*_y$ then $\neg x \notin (K^*_x)^{*\|x}_y$

AGM-It. Assume that $\neg y \notin K^*_x$. By Closure, it follows that K^*_x is consistent, whereby, from Success it follows that $x \wedge y \nvdash \perp$. It follows from Conjunction then that $(K^*_x)^{*|x}_y = K^*_{x\wedge y}$, as desired.

DP1. Assume that $y \vdash x$. In case $\vdash \neg y$ then the proof is trivial (due to Success). Otherwise, given $y \vdash x$, it follows that $\vdash (x \wedge y) \leftrightarrow y$. Hence, by Extensionality, $K^*_{x\wedge y} = K^*_y$. Furthermore, by Conjunction, $(K^*_x)^{*|x}_y = K^*_{x\wedge y}$. Hence the desired result follows.

DP3. Assume that $x \in K^*_y$. If x and y are mutually inconsistent, it must be because $\vdash \neg y$ (use Consistency) whereby $x \in (K^*_x)^{*|x}_y$ (by Success and Closure). On the other hand, assume that x and y are mutually consistent. Hence, by Conjunction $K^*_{x\wedge y} = (K^*_x)^{*|x}_y$. Now, by Success and Closure $x \in K^*_{x\wedge y}$ wherefrom the desired result follows.

DP4. Assume that $\neg x \notin K^*_y$. Due to Success and Closure, surely then x and y are mutually consistent. Then by Conjunction, $(K^*_x)^{*|x}_y = K^*_{x\wedge y}$ which, by Consistency, is consistent, and to which, by Success and Closure, x belongs. Hence $\neg x \notin (K^*_x)^{*|x}_y$. ■

Theorem 3 *Let e be any arbitrary sentence. Let $*$ and its revision $*|e$ by e be two revision operations such that the postulates 0*-8*(new) are satisfied. Let K be an arbitrary belief set and $K' = K^*_e$. Then $(\preceq_{*,K})^{\odot}_e = \preceq_{*|e,K'}$*

Proof
($\subseteq$)

First we show that if $x\ (\preceq_{*,K})^{\odot}_e\ y$ then $x \preceq_{*|e,K'} y$. Assume that $x\ (\preceq_{*,K})^{\odot}_e\ y$. Then, by Definition 2, either (a) $e \vdash \perp$ or (b) $e \vdash x,\ e \vdash y$ AND either $\vdash y$ or both $\nvdash x$ and $x \preceq_{*,K} y$ or (c) $e \nvdash x$ and $e \rightarrow x\ \preceq_{*,K}\ e \rightarrow y$.

CASE (a). $e \vdash \bot$. Hence $K' = K^*_e = K_\bot$. Then, trivially, $x \preceq_{*|e,K'} y$ by Definition 4.

CASE (b). First consider the situation $\vdash y$. Now, either $\vdash x$ or $\nvdash x$. If $\vdash x$, then we get $\vdash x \wedge y$, in which case, by definition (4), we get $x \preceq_{*|e,K'} y$. On the other hand, consider $\nvdash x$. Note that since $\vdash y$, it follows that $\neg x \vee \neg y$ is logically equivalent to $\neg x$. Furthermore, since $\nvdash x$, surely $\neg x \nvdash \bot$ whereby, $K'^{*|e}_{\neg x \vee \neg y} = K'^{*|e}_{\neg x} \neq K_\bot$ (Extensionality and Consistency). Since, by Success $\neg x \in K'^{*|e}_{\neg x}$ it follows that $x \notin K'^{*|e}_{\neg x} = K'^{*|e}_{\neg x \vee \neg y}$. Hence, by Definition 4 it follows that $x \preceq_{*|e,K'} y$.

Now consider the other relevant situation, i.e., both $\nvdash x$ and $x \preceq_{*,K} y$. (Also note the active assumptions that $e \vdash x,\ e \vdash y$). By Definition 4 then, either (i) $x \notin K^*_{\neg x \vee \neg y}$ or (ii) $\vdash x \wedge y$ or (iii) $K = K_\bot$. In case (ii), when $\vdash x \wedge y$, we get the desired result trivially (using Definition 4 again). So we consider only cases (i) and (iii).

Note that since both x and y are consequences of e, both x and y are in K'. Now we claim that $x \notin K'^{*|e}_{\neg x \vee \neg y}$ whereby, from Definition 4 we get the desired result that $x \preceq_{*|e,K'} y$. The demonstration is simple. Note that since in (a) we considered the case $e \vdash \bot$, we can assume, without any loss of generality, that $e \nvdash \bot$. Furthermore, since both $e \vdash x$ and $e \vdash y$, it follows that $e \wedge (\neg x \vee \neg y) \vdash \bot$. Hence, by (8*new, i.e., DP2′) it follows that $K'^{*|e}_{\neg x \vee \neg y} = (K^*_e)^{*|e}_{\neg x \vee \neg y} = K^*_{\neg x \vee \neg y}$.

In case (iii), $K = K_\bot$. Then $K' = (K_\bot)^*_e = Cn(\{e\})$ (by Absurdity). Furthermore, $K'^{*|e}_{\neg x \vee \neg y} = (K_\bot)^*_{\neg x \vee \neg y} =$ (by Absurdity) $Cn(\{\neg x \vee \neg y\})$. By assumption $\nvdash x$ whereby $\neg x \nvdash x$; therefore $x \notin Cn(\{\neg x \vee \neg y\}) = K'^{*|e}_{\neg x \vee \neg y}$.

In order to complete case (b), we now consider case (i) when $x \notin K^*_{\neg x \vee \neg y}$. Furthermore, $K'^{*|e}_{\neg x \vee \neg y} = (K^*_e)^{*|e}_{\neg x \vee \neg y} = K^*_{\neg x \vee \neg y}$ (using (8*new)). Hence $x \notin K'^{*|e}_{\neg x \vee \neg y}$ from which, with the help of Definition 4 the desired result follows.

CASE (c). Assume that $e \nvdash x$ and $e \rightarrow x \preceq_{*,K} e \rightarrow y$. Now, since $e \nvdash x$ it follows that $e \nvdash x \wedge y$. Hence $e \wedge (\neg x \vee \neg y) \nvdash \bot$. Hence, by (7*new), we get $K'^{*|e}_{\neg x \vee \neg y} = (K^*_e)^{*|e}_{\neg x \vee \neg y} = K^*_{e \wedge (\neg x \vee \neg y)}$. Now, applying Definition 4 to the assumption that $e \rightarrow x \preceq_{*,K} e \rightarrow y$, we get, either (i) $e \rightarrow x \notin K^*_{\neg(e \rightarrow x) \vee \neg(e \rightarrow y)}$ or (ii) $\vdash (e \rightarrow x) \wedge (e \rightarrow y)$ or (iii) $K = K_\bot$. Consider case (i). Note that $\neg(e \rightarrow x) \vee \neg(e \rightarrow y)$ is logically equivalent to $e \wedge (\neg x \vee \neg y)$ whereby $K^*_{\neg(e \rightarrow x) \vee \neg(e \rightarrow y)} = K'^{*|e}_{\neg x \vee \neg y}$. Furthermore, since $x \vdash e \rightarrow x$ and $e \rightarrow x \notin K^*_{\neg(e \rightarrow x) \vee \neg(e \rightarrow x)}$ it follows that $x \notin K'^{*|e}_{\neg x \vee \neg y}$, as desired.

Next consider case (ii). This is an impossible case since, given $e \nvdash x$, it follows that $e \nvdash x \wedge y$ contradicting the initial assumption that $\vdash (e \rightarrow x) \wedge (e \rightarrow y)$.

Finally we consider case (iii). Since $K'^{*|e}_{\neg x \vee \neg y} = K^*_{e \wedge (\neg x \vee \neg y)}$ and $K = K_\bot$, it follows from (0*) that $K'^{*|e}_{\neg x \vee \neg y} = Cn(\{e \wedge (\neg x \vee \neg y)\})$. By (Deduction), it follows that x is in this set if and only if $e \vdash (\neg x \vee \neg y) \rightarrow x$ only if $e \vdash \neg x \rightarrow x$ only if $e \vdash x$. But by assumption $e \nvdash x$ whereby, $x \notin K'^{*|e}_{\neg x \vee \neg y}$, as desired.

($\supseteq$)

Now we show that if $x \preceq_{*|e,K'} y$ then $x (\preceq_{*,K})^{\odot}_{e} y$. Assume that $x \preceq_{*|e,K'} y$. Then, by Definition 4, either (1) $\vdash x \wedge y$ or (2) $K' = K_\bot$ or (3) $x \notin K'^{*|e}_{\neg x \vee \neg y}$.

CASE 1. Since $\vdash x \wedge y$ it follows trivially from part (b) of Definition 2 that $x\ (\preceq_{*,K})^{\odot}_{e}\ y$.

CASE 2. $K^*_e = K' = K_\bot$. By (Consistency) it follows that $e \vdash \bot$. So from part (a) of

Definition 2 it follows that $x\ (\preceq_{*,K})^{\odot}_{e}\ y$.

CASE 3. $x \notin K'^{*|e}_{\neg x \vee \neg y}$. Now, either (i) $e \vdash \bot$ or (ii) $e \nvdash \bot$ and $e \vdash x \wedge y$ or (iii) $e \nvdash \bot$ and $e \nvdash x \wedge y$.

In case (i) we get the desired result trivially (part (a) of Definition 2).

In case (ii) Since $e \nvdash \bot$ and $e \vdash x \wedge y$, it follows from (8*new) that $K'^{*|e}_{\neg x \vee \neg y} = (K^*_e)^{*|e}_{\neg x \vee \neg y} = K^*_{\neg x \vee \neg y}$. Hence $x \notin K^*_{\neg x \vee \neg y}$ from which by Definition 4 it follows that $x \ \preceq_{*,K}\ y$. However, since $e \vdash x \wedge y$, the desired result easily follows from it with the help of Definition 2b.

In case (iii) since $e \nvdash x \wedge y$, we get by (7*new) that $K'^{*|e}_{\neg x \vee \neg y} = K^*_{e \wedge (\neg x \vee \neg y)}$. So $x \notin K^*_{e \wedge (\neg x \vee \neg y)}$. Noting that $e \wedge (\neg x \vee \neg y)$ is logically equivalent to $e \wedge \neg(x \wedge y)$, from Definition 4 it then follows, among other things, that $e \wedge \neg(x \wedge y) \rightarrow x \preceq_{K,*} e \wedge \neg(x \wedge y) \rightarrow \neg x$. By substituting logical equivalents we get $(e \rightarrow x \wedge y) \vee x \preceq_{K,*} (e \rightarrow x \wedge y) \vee \neg x$. Since the conjunction of $(e \rightarrow x \wedge y) \vee x$ and $(e \rightarrow x \wedge y) \vee \neg x$ is logically equivalent to $e \rightarrow x \wedge y$, from the EE postulates we get $e \rightarrow x \wedge y \equiv_{K,*} (e \rightarrow x \wedge y) \vee x$. However, since $(e \rightarrow x \wedge y) \vee x$ is logically equivalent to $e \rightarrow x$, we get $e \rightarrow x \wedge y \equiv_{K,*} e \rightarrow x$. Using the EE postulates again (since $e \rightarrow x \wedge y$ is logically equivalent $(e \rightarrow x) \wedge (e \rightarrow y)$) we get $e \rightarrow x \preceq_{K,*} e \rightarrow y$.

To get the desired result with the help of Definition 2c all we need to do now is to show that $e \nvdash x$. We do that below. Suppose to the contrary that $e \vdash x$. Then $\vdash e \rightarrow x$. Furthermore, since $e \nvdash x \wedge y$ it follows that $\nvdash e \rightarrow y$. Now, $e \rightarrow x$, being a tautology, $p \preceq_{*,K} e \rightarrow x$ for all sentences p (by Dominance). Since $e \rightarrow x \preceq_{*,K} e \rightarrow y$, it follows by Transitivity and (EE4) that if $\bot \prec_{*,K} z$ for some z, then $\vdash e \rightarrow y$. But since (as we saw above) $\nvdash e \rightarrow y$, it follows that $\bot \not\prec_{*,K} z$ for every sentence z. Since $\preceq_{*,K}$ is connected (as follows from EE1-3), it follows that $z \preceq_{*,K} \bot$ for all z. Thus $\bot$ is $\preceq_{*,K}$-maximal. Hence, by Definition 1, it follows that $K = K_\bot$. Hence, by Absurdity it follows that $K^*_{e \wedge (\neg x \vee \neg y)} = Cn(\{e \wedge (\neg x \vee \neg y\})$. Since by assumption $e \vdash x$, it follows that $x \in K^*_{e \wedge (\neg x \vee \neg y)}$. But this conflicts with the starting assumption that $x \notin K^*_{e \wedge (\neg x \vee \neg y)}$. Contradiction! ■

Theorem 4 *Let $\preceq$ be an EE relation with epistemic content $EC(\preceq) = K$, and $\preceq^{\odot}_{e}$ the result of revising it by a sentence e. Then the revision operations $* = *_{\preceq}$ and $*|e = *_{\preceq^{\odot}_{e}}$ satisfy the postulates 0*-8*(new) with respect to the belief set $K = EC(\preceq)$.*

Proof

Let $\preceq$ be an EE relation with epistemic content $EC(\preceq) = K$, and $\preceq^{\odot}_{e}$ the result of revising it by a sentence e. Define the revision operations $* = *_{\preceq}$ and $*|e = *_{\preceq^{\odot}_{e}}$. We need to show that $*$ and $*|e$ satisfy the postulates 0*-8*(new) with respect to the belief set $K = EC(\preceq)$.

From Definition 3 we get that (i) $y \in K^*_x$ iff either $x \rightarrow \neg y \prec x \rightarrow y$ or $\vdash \neg x$ or $x \vdash y$, and (ii) $y \in K^{*|e}_x$ iff either $x \rightarrow \neg y \prec^{\odot}_{e} x \rightarrow y$ or $\vdash \neg x$ or $x \vdash y$. From Theorem 1a of [10] we know that $\preceq^{\odot}_{e}$ is an EE relation. By slightly modifying the AGM proof to the same effect, we learn that, since $\preceq$ and $\preceq^{\odot}_{e}$ are EE relations, the operations $*$ and $*|e$ satisfy (1* - 6*). We need to show only 0*, 7*(new) and 8*(new).

0*. First we show that $*$ satisfies (0*). Assume that $K = K_\bot$. From Definition 1 of Epistemic Content, it follows then that $\bot$ is $\preceq$-maximal. I.e., $z \preceq \bot$ for all z. On the other hand, it follows from Dominance that $\bot \preceq z$ for all z. From the properties of an EE relation, it follows then that $a \equiv b$ for all a and b. Hence, it follows from (i) that $y \in K^*_x$ iff either $\vdash \neg x$ or $x \vdash y$. I.e., if $\vdash \neg x$, then $K^*_x = Cn(\bot)$ and otherwise, $K^*_x = Cn(x)$. In other words,

$K_x^* = Cn(x)$. Thus, $*$ satisfies (0*).

Next we show that $*|e$ satisfies (0*). Assume then that $EC(\preceq_e^{\odot}) = K_\perp$. By similar argument as above, we then obtain that $K_x^{*|e} = Cn(x)$.

(7* new). Assume that $e \wedge x \nvdash \perp$. We need to show that $(K_e^*)_x^{*|e} = K_{e\wedge x}^*$. Now there are two cases: either $K = K_\perp$ or $K \neq K_\perp$.

CASE 1.

Assume that $K = K_\perp$. It follows from Definition 1 and EE1-4 that $a \preceq b$ for all a and b. Hence, it follows from Definition 3 that RHS = $Cn(\{e \wedge x\})$. On the other hand, from Definition 2 it follows that

- $x \preceq_e^{\odot} y$ iff either $e \vdash \perp$ or $e \nvdash x$ or both $e \vdash x \wedge y$ and either $\nvdash x$ or $\vdash y$.

It would be sufficient to show that $y \in (K_e^*)_x^{*|e}$ iff $y \in Cn(e \wedge x)$. Note that from Definition 3 it follows that $y \in (K_e^*)_x^{*|e}$ iff either $x \to \neg y \prec_e^{\odot} x \to y$ or $\vdash \neg x$ or $x \vdash y$.

(Only if.) Assume that $y \in (K_e^*)_x^{*|e}$. Hence, either $x \to \neg y \prec_e^{\odot} x \to y$ or $\vdash \neg x$ or $x \vdash y$. Obviously, if either $\vdash \neg x$ or $x \vdash y$, then we get the desired result, that $y \in Cn(e \wedge x)$ trivially. We need only consider the remaining case. Assume then that $x \to \neg y \prec_e^{\odot} x \to y$. It follows then that $x \to y \npreceq_e^{\odot} x \to \neg y$. From the Definition 2 [note the displayed item above], it then follows, among other things, that $e \vdash x \to y$ whereby $y \in Cn(e \wedge x)$.

(If.) Assume that $y \in Cn(e \wedge x)$. Assume further that $\nvdash \neg x$ and $x \nvdash y$. We need to show that $x \to \neg y \prec_e^{\odot} x \to y$. By our initial assumption $e \wedge x \nvdash \perp$, whereby $e \nvdash \perp$. Since $y \in Cn(e \wedge x)$, it follows that $e \vdash x \to y$. Furthermore, since $e \wedge x \nvdash \perp$, we get $e \nvdash \neg x$ whereby, $e \nvdash x \to (y \wedge \neg y)$. Hence, vacuously, we get that if both $e \vdash x \to y$ and $e \vdash x \to \neg y$, then both $\vdash x \to y$ and $\nvdash x \to \neg y$. (As we noted above, $x \preceq_e^{\odot} y$ iff either $e \vdash \perp$ or $e \nvdash x$ or both $e \vdash x \wedge y$ and either $\nvdash x$ or $\vdash y$.) Hence it follows that $x \to y \npreceq_e^{\odot} x \to \neg y$ whereby we get the desired result that $x \to \neg y \prec_e^{\odot} x \to y$.

CASE 2.

Assume that $K \neq K_\perp$. Note that since $K \neq K_\perp$, from Definition 1 it follows that $\perp$ is not $\preceq$-maximal. Then by EE4, every $\preceq$-maximal sentence is a theorem ($\top$). Hence, in this case, part (b) of Definition 2 can be equivalently replaced by:

- $e \vdash x$, $e \vdash y$ and $x \preceq y$.

(Only if.) Assume that $y \in (K_e^*)_x^{*|e}$. We need to show that $y \in K_{e\wedge x}^*$.

By Definition 3, we get either $x \to \neg y \prec_e^{\odot} x \to y$ or $\vdash \neg x$ or $x \vdash y$. The last two cases are trivial; so consider the principal case: $x \to \neg y \prec_e^{\odot} x \to y$. I.e., $x \to y \npreceq_e^{\odot} x \to \neg y$. It follows from Definition 2 then [note the simplification to clause (b) mentioned above] that (i) $e \nvdash \perp$, (ii) If $e \vdash x \to y$ and $e \vdash x \to \neg y$, then $x \to y \npreceq x \to \neg y$ (iii) If $e \nvdash x \to y$ then $e \to (x \to y) \npreceq e \to (x \to \neg y)$.

Now, either $e \vdash x \to y$ or $e \nvdash x \to y$. Consider the first case. $e \vdash x \to y$. Then the desired result, that $y \in K_{e\wedge x}^*$, follows trivially from Definition 3. Now consider the second case, $e \nvdash x \to y$. By (iii) we get $e \to (x \to \neg y) \prec e \to (x \to y)$ whereby, $(e \wedge x) \to \neg y \prec (e \wedge x) \to y$. Hence, by Definition 3, we get $y \in K_{e\wedge x}^*$, as desired.

(If.) Assume that $y \in K_{e\wedge x}^*$. We need to show that $y \in (K_e^*)_x^{*|e}$. It follows from Definition 3 that $y \in (K_e^*)_x^{*|e}$ if and only if either $\vdash \neg x$ or $x \vdash y$ or $x \to \neg y \prec_e^{\odot} x \to y$. So we assume that $\nvdash \neg x$ and $x \nvdash y$. It will be sufficient to show that $x \rightarrow y \npreceq^{\odot}_{e} x \rightarrow \neg y$. It follows from Definition 2 that we need to demonstrate the following three claims:

(A) $e \nvdash \bot$, (B) If $e \vdash \neg x$ then $x \rightarrow y \npreceq x \rightarrow \neg y$ and (C) If $e \nvdash x \rightarrow y$ then $e \rightarrow (x \rightarrow \neg y) \prec e \rightarrow (x \rightarrow y)$. Claims (A) and (B) follow from the assumption that $e \wedge x \nvdash \bot$. So we need to demonstrate only the last claim (C).

Now, since $y \in K^{*}_{e \wedge x}$, it follows by Definition 3 that either $\vdash \neg(e \wedge x)$ or $e \wedge x \vdash y$ or $e \wedge x \rightarrow \neg y \prec e \wedge x \rightarrow y$. The first case is impossible since by assumption $e \wedge x \nvdash \bot$. In the second case, the claim (C) follows trivially, since, if $e \wedge x \vdash y$, then by Deduction, $e \vdash x \rightarrow y$ (which contradicts the antecedent of (C)). Similarly, in the third case (C) follows trivially, since, $e \wedge x \rightarrow \neg y \prec e \wedge x \rightarrow y$ is logically equivalent to the consequent, $e \rightarrow (x \rightarrow \neg y) \prec e \rightarrow (x \rightarrow y)$, of (C).

(8*new)

We need to show that if both $\nvdash \neg e$ and $e \wedge x \vdash \bot$ then $(K^{*}_{e})^{*|e}_{x} = K^{*}_{x}$.

Assume that $\nvdash \neg e$ and $e \wedge x \vdash \bot$. Now, either $K = K_{\bot}$ or $K \neq K_{\bot}$.

CASE 1. Assume that $K = K_{\bot}$. Now, by Definition 3, $y \in LHS$ iff either $x \rightarrow \neg y \prec^{\odot}_{e} x \rightarrow y$ or $\vdash \neg x$ or $x \vdash y$. On the other hand, $y \in RHS$ iff either $x \rightarrow \neg y \prec x \rightarrow y$ or $\vdash \neg x$ or $x \vdash y$. Now, right to left is easy. Assume that $y \in RHS$, i.e., either $x \rightarrow \neg y \prec x \rightarrow y$ or $\vdash \neg x$ or $x \vdash y$. Since $K = K_{\bot}$, it follows from Definition 1 that $a \preceq b$ for all $a,\ b$ whereby $x \rightarrow \neg y \nprec x \rightarrow y$. Hence, either $\vdash \neg x$ or $x \vdash y$. In either of these cases, it trivially follows that $y \in LHS$. Hence consider the proof for left to right. Assume that $y \in LHS$. Hence either $x \rightarrow \neg y \prec^{\odot}_{e} x \rightarrow y$ or $\vdash \neg x$ or $x \vdash y$. The last two cases are trivial. So consider the principal case: $x \rightarrow \neg y \prec^{\odot}_{e} x \rightarrow y$. I.e., $x \rightarrow y \npreceq^{\odot}_{e} x \rightarrow \neg y$. Applying Definition 2 we get, among other things, that either $e \nvdash x \rightarrow y$ or $e \nvdash x \rightarrow \neg y$ or both $\vdash x \rightarrow \neg y$ and either $\vdash x \rightarrow y$ or $x \rightarrow y \npreceq x \rightarrow \neg y$. (This particularly follows from Definition 2b.) Now, by assumption $e \wedge x \vdash \bot$ (i.e. $e \vdash \neg x$) from which it follows both that $e \vdash x \rightarrow y$ and $e \vdash x \rightarrow \neg y$. Hence, both $\vdash x \rightarrow \neg y$ and either $\vdash x \rightarrow y$ or $x \rightarrow y \npreceq x \rightarrow \neg y$. However, since $K = K_{\bot}$ it follows that $x \rightarrow y \preceq x \rightarrow \neg y$. Hence both $\vdash x \rightarrow \neg y$ and $\vdash x \rightarrow y$ from which it follows that $\vdash \neg x$. From this it trivially follows that $y \in RHS$.

CASE 2. Assume that $K \neq K_{\bot}$. Again, $y \in LHS$ iff either $x \rightarrow \neg y \prec^{\odot}_{e} x \rightarrow y$ or $\vdash \neg x$ or $x \vdash y$. From Definition 2, the connectedness of $\preceq$, and the fact that $K \neq K_{\bot}$, it follows that $x \rightarrow \neg y \prec^{\odot}_{e} x \rightarrow y$ iff (a) $e \nvdash \bot$, (b) If $e \vdash \neg x$ then $x \rightarrow \neg y \prec x \rightarrow y$ and (c) If $e \nvdash x \rightarrow y$ then $e \rightarrow (x \rightarrow \neg y) \prec e \rightarrow (x \rightarrow y)$. Now, (a) is satisfied by assumption. Furthermore, (c) is vacuously satisfied, since by assumption $e \vdash \neg x$. Besides, the antecedent of (b) is also satisfied by assumption. Hence, $x \rightarrow \neg y \prec^{\odot}_{e} x \rightarrow y$ iff $x \rightarrow \neg y \prec x \rightarrow y$. Thus, $y \in LHS$ iff either $x \rightarrow \neg y \prec x \rightarrow y$ or $\vdash \neg x$ or $x \vdash y$ iff, by Definition 3, $y \in RHS$. ■

Theorem 5 *Let $*$ and $*|e$ be two revision operations satisfying the postulates 0*-8*(new). Let K be an arbitrary belief set and $K' = K^{*}_{e}$. Then the revision operations $*_{\preceq_{*,K}}$ and $*_{\preceq_{*|e,K'}}$ are belief revision operations that satisfy 0*-8*(new) with respect to the belief set K.*

Proof

Let $*$ and $*|e$ be two revision operations satisfying the postulates 0*-8*(new). Let K be an arbitrary belief set and $K' = K^{*}_{e}$. It follows from Theorem 3 that $\preceq_{*|e,K'} = (\preceq_{*,K})^{\odot}_{e}$. Hence it would be sufficient to show that the operations $*_{\preceq_{*,K}}$ and $*_{(\preceq_{*,K})^{\odot}_{e}}$ satisfy 0*-8*(new) with respect to the belief set K.

It is easily seen ([6]; note the extra clause, $K = K_\perp$ in Definition 4 to compensate for the difference between SEE and EE relations) that $\preceq_{*,K}$ is an EE relation. Furthermore, we know from Theorem 1a of [10] that, since $\preceq_{*,K}$ is an EE relation, the relation $(\preceq_{*,K})^{\odot}_{e}$ is also an EE relation. Hence, surely $*_{\preceq_{*,K}}$ and $*_{(\preceq_{*,K})^{\odot}_{e}}$ satisfy the basic AGM revision postulates 1*-6* (see [6]; note the extra clause, $x \vdash y$, in Definition 3 to compensate for the difference between SEE and EE relations). So we need only show that $*_{\preceq_{*,K}}$ and $*_{(\preceq_{*,K})^{\odot}_{e}}$ satisfy 0*, 7*new and 8*new.

0*.

(i) Suppose $K = K_\perp$. Then using definitions (4) and (3) sequentially, and noting that $x \preceq_{*,K} y$ for all x and y (by Definition 4), it is easily seen that $y \in K_x^{*\preceq_{*,K}}$ iff either $\vdash \neg x$ or $x \vdash y$. Hence, irrespective of whether $\vdash \neg x$ or not, we get $K_x^{*\preceq_{*,K}} = Cn(x)$. Thus, the operation $*_{\preceq_{*,K}}$ satisfies 0*. (ii) By similar argument as above, it is easily seen that the operation $*_{(\preceq_{*,K})^{\odot}_{e}}$ satisfied 0*.

7*new.

Assume that $e \wedge x \nvdash \perp$. We need to show that $y \in (K_e^{*\preceq_{*,K}})_x^{*(\preceq_{*,K})^{\odot}_{e}}$ iff $y \in K_{e\wedge x}^{*\preceq_{*,K}}$.

(Only if.) Assume that $y \in (K_e^{*\preceq_{*,K}})_x^{*(\preceq_{*,K})^{\odot}_{e}}$. It follows from Definition 3 then that either $x \rightarrow \neg y \; (\prec_{*,K})^{\odot}_{e} \; x \rightarrow y$ or $\vdash \neg x$ or $x \vdash y$. The last two cases are easy: we know that $*_{\preceq_{*,K}}$ satisfies 1*-6*, in particular, Closure and Success; hence, irrespective of whether $\vdash \neg x$ or $x \vdash y$, surely $y \in K_{e\wedge x}^{*\preceq_{*,K}}$. So, consider the principal case. Assume that $x \rightarrow y \; (\not\preceq_{*,K})^{\odot}_{e} \; x \rightarrow \neg y$. Hence, by Definition 2, (a) $e \nvdash \perp$, (b) If $e \vdash x \rightarrow y$ and $e \vdash x \rightarrow \neg y$, then $x \rightarrow y \not\preceq_{*,K} x \rightarrow \neg y$, and (c) If $e \nvdash x \rightarrow y$ then $e \rightarrow (x \rightarrow y) \not\preceq_{*,K} e \rightarrow (x \rightarrow \neg y)$. Now, by Definition 3, we need to show that either $e \wedge x \vdash \perp$ or $e \wedge x \vdash y$ or $(e \wedge x) \rightarrow \neg y \prec_{*,K} (e \wedge x) \rightarrow y$. We assume $e \wedge x \nvdash y$. It will be sufficient to show that $(e \wedge x) \rightarrow y \not\preceq_{*,K} (e \wedge x) \rightarrow \neg y$.

Since $e \wedge x \nvdash y$, it follows that $e \nvdash x \rightarrow y$. Hence, the desired result follows from (c).

(If). Assume that $y \in K_{e\wedge x}^{*\preceq_{*,K}}$. Hence, by Definition 3, either (i) $e \wedge x \vdash \perp$ or (ii) $e \wedge x \vdash y$ or (iii) $(e \wedge x) \rightarrow \neg y \prec_{*,K} (e \wedge x) \rightarrow y$. We need to show that $y \in (K_e^{*\preceq_{*,K}})_x^{*(\preceq_{*,K})^{\odot}_{e}}$. In light of Definition 3 we assume that $\nvdash \neg x$ and $x \nvdash y$. It will be sufficient to show that $x \rightarrow y \; (\not\preceq_{*,K})^{\odot}_{e} \; x \rightarrow \neg y$. By Definition 2, it follows then that it will be sufficient to show that (a) $e \nvdash \perp$, (b) If both $e \vdash x \rightarrow y$ and $e \vdash x \rightarrow \neg y$, then $x \rightarrow y \not\preceq_{*,K} x \rightarrow \neg y$, and (c) If $e \nvdash x \rightarrow y$ then $e \rightarrow (x \rightarrow y) \not\preceq_{*,K} e \rightarrow (x \rightarrow \neg y)$. Note that $e \wedge x \nvdash \perp$ by assumption. Hence $e \nvdash \perp$ whereby (a) is trivially satisfied. Furthermore, it also follows from the same assumption that not both $e \vdash x \rightarrow y$ and $e \vdash x \rightarrow \neg y$, whereby (b) is vacuously satisfied. So we need to show only (c).

Case (i) is impossible since, by assumption, $e \wedge x \nvdash \perp$.

Case (ii) it follows by Deduction that $e \vdash x \rightarrow y$. Hence (c) is trivially satisfied.

Case (iii). Trivially (c) is satisfied.

8* (new).

Assume that $\nvdash \neg e$ and $\vdash \neg(e \wedge x)$. We need to show that $y \in (K_e^{*\preceq_{*,K}})_x^{*(\preceq_{*,K})^{\odot}_{e}}$ iff $y \in K_x^{*\preceq_{*,K}}$.

(Only if.)

Assume that $y \in (K_e^{*\preceq_{*,K}})_x^{*(\preceq_{*,K})^{\odot}_{e}}$. Hence, by Definition 3, either (i) $x \rightarrow \neg y \; (\prec_{*,K})^{\odot}_{e} \; x \rightarrow$

y or (ii) $\vdash \neg x$ or (iii) $x \vdash y$. Cases (ii) and (iii) are trivial; so we consider only the principal case (i). So assume that $x \rightarrow \neg y \ (\prec_{*,K})^{\odot}_{e} \ x \rightarrow y$. It follows from definition (2, particularly 2b) then, among other things, that if $e \vdash \neg x$ (i.e., if both $e \vdash x \rightarrow y$ and $e \vdash x \rightarrow \neg y$), then $x \rightarrow \neg y \ \prec_{*,K} \ x \rightarrow y$. It follows from Definition 3 then that $y \in K_x^{*\preceq_{*,K}}$.

(If.)

Conversely, assume that $y \in K_x^{*\preceq_{*,K}}$. Hence, by Definition 3, either (i) $x \rightarrow \neg y \ \prec_{*,K} \ x \rightarrow y$ or (ii) $\vdash \neg x$ or (iii) $x \vdash y$. Cases (ii) and (iii) are trivial. So we consider only case (i). Assume then that $x \rightarrow \neg y \ \prec_{*,K} \ x \rightarrow y$. Hence, vacuously, we get that if both $e \vdash x \rightarrow y$ and $e \vdash x \rightarrow \neg y$ then $x \rightarrow y \ \not\preceq_{*,K} \ x \rightarrow \neg y$ (use part b of Definition 2). Furthermore, by assumption, $\not\vdash \neg e$. Besides, since by assumption $e \vdash \neg x$, we vacuously get that if $e \not\vdash x \rightarrow y$ then $e \rightarrow (x \rightarrow y) \ \not\preceq_{*,K} \ e \rightarrow (x \rightarrow \neg y)$ (use part c of Definition 2). From Definition 2 we then get that $x \rightarrow \neg y \ (\prec_{*,K})^{\odot}_{e} \ x \rightarrow y$. It follows from definition (3) then that $y \in (K_e^{*\preceq_{*,K}})_x^{*(\preceq_{*,K})^{\odot}_{e}}$.
■

Theorem 6 *Let $\preceq$ be an EE relation with epistemic content $EC(\preceq) = K$, $\preceq^{\odot}_{e}$ the result of revising $\preceq$ by a sentence e and $K' = EC(\preceq^{\odot}_{e})$ the epistemic content of $\preceq^{\odot}_{e}$. Then $\preceq_{*_{\preceq},K} = \preceq$ and $\preceq_{*_{\preceq^{\odot}_{e}},K'} = \preceq^{\odot}_{e}$.*

Proof

[We omit the proof of this theorem. The proof can be easily constructed by slightly modifying the analogous proof in [6].] ■