ARTICLE IN PRESS

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.1 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 1

ELSEVIER

Available at www.ComputerScienceWeb.com POWERED BY SCIENCE @ DIRECT®

Artificial Intelligence ●●● (●●●●) ●●●–●●●

Artificial Intelligence

www.elsevier.com/locate/artint

Dynamic belief revision operators ☆

Abhaya C. Nayak [a], Maurice Pagnucco [b,∗], Pavlos Peppas [c]

[a] *Department of Computing, Division of Information and Communication Sciences, Macquarie University, NSW 2109, Australia*
[b] *Artificial Intelligence Laboratory, School of Computer Science and Engineering, The University of New South Wales, Sydney, NSW 2052, Australia*
[c] *AIT—Athens Information Technology, 19.5km Markopoulo Avenue, Peania 19002, Athens, Greece*

Received 30 April 2002

Abstract

The AGM approach to belief change is not geared to provide a decent account of iterated belief change. Darwiche and Pearl have sought to extend the AGM proposal in an interesting way to deal with this problem. We show that the original Darwiche–Pearl approach is, on the one hand excessively strong and, on the other rather limited in scope. The later Darwiche–Pearl approach, we argue, although addresses the first problem, still remains rather permissive. We address both these issues by (1) assuming a dynamic revision operator that changes to a new revision operator after each instance of belief change, and (2) strengthening the Darwiche–Pearl proposal. Moreover, we provide constructions of this dynamic revision operator via entrenchment kinematics as well as a simple form of lexicographic revision, and prove representation results connecting these accounts.

Keywords: Iterated belief change; Entrenchment kinematics; Lexicographic revision; Unary revision

The study of belief change focuses on the way in which a reasoning entity should modify its stock of beliefs in light of new information. The issues that are important here are also

☆ Part of this paper appeared in: Y. Shoham (Ed.), Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge, Morgan Kaufmann Publishers, 1996 as "Changing Conditional Beliefs Unconditionally" (pp. 119–135).

∗ Corresponding author.
E-mail addresses: abhaya@ics.mq.edu.au (A.C. Nayak), morri@cse.unsw.edu.au (M. Pagnucco), pavlosp@cse.unsw.edu.au (P. Peppas).
URLs: http://www.comp.mq.edu.au/~abhaya/ (A.C. Nayak), http://www.cse.unsw.edu.au/~morri/ (M. Pagnucco), http://www.cse.unsw.edu.au/~pavlosp/ (P. Peppas).

0004-3702/03/$ – see front matter
doi:10.1016/S0004-3702(03)00017-1

explored in other guises, for example, logics for conditionals [4–6] and non-monotonic logics [11,15].

Classical accounts of belief change (such as AGM [8]) are geared to deal with "one-shot" belief change. While providing a cogent account of how a rational agent should change her beliefs in light of a piece of evidence, they fail to give a systematic account of belief change as an iterative process—of how an epistemic agent should deal with a (possibly temporal) sequence of evidential pieces of information. This shortcoming is addressed in the literature on *iterated belief change* [2,4,7,13,18,20,26]. One account of iterated belief change that has attracted a lot of scrutiny is that of Darwiche and Pearl [6,7]. However, this approach has been undermined on two fronts. Firstly, it has been shown by Freund and Lehmann [9] (and later by Nayak et al. [17]) that this account is inconsistent with the classical AGM belief change framework [8] which it seeks to extend and hence is overly strong. Secondly, it allowed iterated changes of belief that are intuitively unappealing. In the later development of their work Darwiche and Pearl [7] attempt to address the former problem but the latter persists. Even so, the solution to the first problem is not without concern as it requires a modification to the way in which the AGM framework is viewed.

In this paper, we address the deficiencies of the Darwiche–Pearl approach to iterated belief revision. It is based on Nayak et al. [17] but substantially extends that account. In particular, while that work [17] lacked a semantics, the current work supplies the semantics of the approach taken. The semantics is provided in a fashion similar to one used in the Darwiche–Pearl account for ease of comparison.

In Section 1 we discuss how the AGM account of belief change [1,8] fails to provide an account of iterated belief change. We also show in this section how the Darwiche–Pearl (D–P) account [6] attempts to deal with this matter, but fails to do so. We point out that Darwiche and Pearl's later work does not address the over-permissiveness of their framework. In Section 2 we provide an extension of a modified D–P account, based on a dynamic belief revision operation, which is successful in this task. We also show that this operation must be viewed as a unary operation for the sake of consistency. In Section 3 we characterise this new framework via entrenchment kinematics. In Section 4 we provide a semantic construction of such dynamic belief revision operators. In Section 5 we provide a brief discussion of how the approach taken in this paper blocks the counterexamples that Darwiche and Pearl use against Natural Revision [4], but their own account does not. We also briefly sketch an account of how our proposal can be modified to deal with inconsistent belief states in a more interesting way. Finally, we conclude with a brief summary.

1. Background

Let us start with the classical account of belief change propounded by Alchourrón, Gärdenfors and Makinson, popularly known as the AGM account of belief change [1,8]. 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). For technical reasons we take this language $\mathcal{L}$ to be a finitary language, i.e., a language generated from a finite set of atomic propositions. 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 $(x \to y) \in Cn(\Gamma)$.
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 beliefs. By $\mathbf{K}$ we denote the set of all possible belief sets (in $\mathcal{L}$). We use $K_\perp$ to denote the absurd belief set $\mathcal{L}$.

At least three forms of epistemic change are recognised in the AGM framework. These changes are represented by three operations $+$ (expansion), $-$ (contraction) and $*$ (revision). These functions are intended to output, given a sentence x, the result of expanding, contracting or revising one's body of beliefs by x. It is however unclear whether these operations, specifically the contraction operation $-$ and the revision operation $*$, are binary operations, mapping an arbitrary pair $\langle K, x\rangle$ of belief set and sentence to a resultant belief set K', or unary operations that map a sentence x to a resultant belief set K' with the understanding that a fixed belief set K is taken to be the background knowledge. In the literature there is support for both these readings, and each of these choices substantially limits what can be achieved in the ensuing framework (see Rott's eminently readable paper [24] for the implications of making these choices). For the present purpose, we take these operations to be binary operations, i.e., $+, -, * : \mathbf{K} \times \mathcal{L} \to \mathbf{K}$ that, for every belief set and a sentential input, returns a belief set.[1]

Since members of $\mathbf{K}$ are often interpreted as belief states, these operations may be viewed as state transition functions.[2] In this paper we confine our discussion only to belief revision ($*$).

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 [8].

(1^*)	K^*_x is a theory	(Closure)
(2^*)	$x \in K^*_x$	(Success)
(3^*)	$K^*_x \subseteq Cn(K \cup \{x\})$	(Inclusion)

[1] We will be forced to discard this assumption in favour of a unary operation in Section 2.5.

[2] Of late [7,18] however, belief states are taken to be more than simply belief sets—belief states are now taken to mean an entity that has information as to what beliefs are currently held, as well as their relative firmness/entrenchment.

(4*) $If\ K \nvdash \neg x$ then $Cn(K \cup \{x\}) \subseteq K^*_x$ (Preservation)

(5*) $K^*_x = K_\perp$ 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).

1.1. The problem of iterated belief change

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 (3*), (4*), (7*) and (8*).) But the AGM system does not constrain iterated belief change in any manner when $\neg y \in K^*_x$. This has occasionally been interpreted as the AGM system permitting, when $\neg y \in K^*_x$, all possible iterated belief changes consistent with 1*–6*. 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. The following example by Darwiche and Pearl [7] illustrates this point:

Example 1 (*Darwiche and Pearl*, 1997). We are introduced to a lady X who sounds smart and looks rich, so we believe that X is smart and X is rich. Moreover, since we profess to no prejudice, we also maintain that X is smart even if found to be poor and, conversely, X is rich even if found not to be smart. Now, we obtain some evidence that X is in fact not smart and we remain of course convinced that X is rich. Still, it would be strange for us to say, "If the evidence turns out false, and X turns out smart after all, we would no longer believe that X is rich". If we currently believe X is smart and rich, then evidence first refuting then supporting that X is smart should not in any way change our opinion about X being rich. Strangely, the AGM postulates do permit such a change of opinion. . . .

Let us formalise and examine this example. We know that

$$K = Cn(\{smart, rich\}), \quad \text{and} \quad \neg smart \in K^*_{\neg smart}.$$

We are to find out whether or not $rich \in (K^*_{\neg smart})^*_{smart}$. Since $\neg smart \in K^*_{\neg smart}$, clearly AGM-It has no bearing on this example, so as far as the classical account of belief revision is concerned, any belief set K' can be this revised belief set so long as $smart \in K'$. For instance, $Cn(\{smart, \neg rich\})$ is as good a candidate as any other!

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 aij1970 P.5 (1-36) by:violeta p. 5

1.2. Darwiche–Pearl proposal

In order to alleviate this situation, Darwiche and Pearl [6] impose four further constraints on the AGM revision operation $*$.[3] Their constraints are couched in the K-M formalism [14] which assumes that each belief set can be expressed as a single sentence.[4] 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 $\neg x \notin (K^*_x)^*_y$.

These four constraints appear to be very plausible indeed. The first postulate is justified on the grounds of specificity: since the subsequent evidence y is more specific than the initial evidence x, the later evidence washes away the earlier evidence. The second constraint says that in the case of two contradictory pieces of evidence, it is the later piece of evidence that prevails. The third postulate is based on the intuition that if learning y is grounds for believing x, then learning x followed by learning y cannot constitute grounds for suspending belief in x. Finally, the last postulate says that if learning y is not sufficient grounds for *dis*-believing x, then learning x and then learning y cannot constitute such grounds either—“no evidence can contribute to its own demise” [7].

Despite the intuitive appeal of this approach, it was realised very early that these postulates are too strong. Freund and Lehmann [9] have shown that DP2 conflicts with the AGM postulates. Even weakening DP2 in the following manner does not help:[5]

- DP2′. If both $\nvdash \neg x$ 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! $\square$

Such considerations have led Lehmann [16] to further weaken DP2. Darwiche and Pearl themselves, on the other hand, have later sought to alter the AGM framework itself to suit their needs [7]. This they do by assuming analogues of the AGM postulates in a system that seeks to revise belief states as opposed to belief sets. This re-interpretation of the AGM system is largely innocuous, except for the modification it requires in the Extensionality Postulate (6^*) which states that logically equivalent pieces of evidence

---

[3] In a later paper [7] Darwiche and Pearl slightly altered their framework to the effect that their work is about the revision of belief *states*, not about belief sets. We will discuss this further later in this section.

[4] Once a finitary language is assumed, a belief set K can be represented by a sentence k such that $Cn(\{k\}) = K$. The constraints on a belief revision operation can then be represented as relations among sentences; e.g., the AGM postulate (7^*) can be rephrased as: $(k * x) \wedge y \models k * (x \wedge y)$. See Katsuno and Mendelzon [14] for details.

[5] We are thankful to Daniel Lehmann for pointing this out in private communication.

have the same effect on a body of beliefs. The postulate (6*) licences, for instance, that $Cn(\{a,b\})^*_{\neg a} = Cn(\{a, a \to b\})^*_{a\to\neg a}$ since $a \to \neg a$ is logically equivalent to $\neg a$ and the two belief sets $Cn(\{a,b\})$ and $Cn(\{a, a \to b\})$ are identical. According to the modification sought by Darwiche and Pearl, it is not sufficient for this purpose that the pieces of evidence be logically equivalent and the body of beliefs identical; the belief states in question must be identical as well. To continue this example, we cannot infer that $Cn(\{a,b\})^*_{\neg a} = Cn(\{a, a \to b\})^*_{a\to\neg a}$ unless we also know that the (identical) belief sets $Cn(\{a,b\})$ and $Cn(\{a, a \to b\})$ are associated with an identical belief state. It is easily verified that given such a weakened version of Extensionality, the proof above that DP2′ (and indeed DP2) is inconsistent with the AGM postulates does not go through. In particular, in the above proof, from the fact that $K'^*_{\neg y} = K^*_{x\wedge\neg y}$, the derivation of $(K'^*_{\neg y})^*_y = (K^*_{x\wedge\neg y})^*_y$ will be blocked.

We do not have any reservation about this particular way of resolving the issue except for the impression it gives that the problem is being altered instead of being solved. Furthermore, there are many accounts of belief *state* revision already available in the literature [18,25,26]. We instead look at the problem from a different perspective. We maintain that the AGM framework can still be retained as is, however the revision operation $*$ adopted by an agent itself evolves in light of new evidence. We recognise that the derivation of $(K'^*_{\neg y})^*_y = (K^*_{x\wedge\neg y})^*_y$ from $K'^*_{\neg y} = K^*_{x\wedge\neg y}$ should be blocked. But the underlying reason, in our view, is different—the two sequential revision operations used are different. We will give more reasons for taking this decision in the next section.

1.3. A stubborn problem

There is however a more serious problem with the Darwiche–Pearl approach that was first reported by Nayak et al. [17]: it is too limited in scope to assist in iterated belief change in some very common situations. This problem is rather resilient in that although the weakening of Extensionality buys Darwiche and Pearl consistency with the underlying belief revision framework, it still fails to address this problem. In order to see that the DP postulates are too limited in scope, consider the following scenario.

Example 2 (*Nayak et al.*, 1996). 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. Note that this problem is not affected by whether the D–P postulates and AGM postulates are taken to be about belief sets or belief states.

As we mentioned earlier, we will develop an account of a dynamic belief revision operation to accommodate Darwiche–Pearl postulates in an AGM framework. Towards this

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.7 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 7

end, we will adopt the original Darwiche–Pearl framework [6] in spirit, without courting inconsistency. In order to handle iterated belief change in all situations, including scenarios depicted by Example 2, we will extend the account of iterated belief change proposed by Darwiche and Pearl. The next section is devoted to this end.

2. Dynamic belief revision

We noted in the last section that although Darwiche and Pearl's later work [7] addresses the issue of their framework being rather too strong, it still leaves the under-specification problem as is. There are primarily two reasons why a complete specification of iterated belief change is lacking in their framework.

(1) *Escaping from inconsistency.* Inconsistent belief sets are not beyond the reach of an agent. However, the D–P approach does not provide an account of how one may revise one's beliefs once she is in $K_\perp$.

(2) *Problem of counteracting evidence.* We will say that x is counter-evidence for y if $\neg y \in K^*_x$. If both $\neg x \in K^*_y$ and $\neg y \in K^*_x$, that is both x and y are counter-evidence for each other, we will say that they are pieces of evidence that counteract. It is possible for x and y to be mutually consistent, and yet counteract: e.g., ability to fly and being a mammal are a counteracting pair, yet they jointly occur. However, none of the D–P postulates are applicable to compute $(K^*_x)^*_y$, if x and y are mutually consistent, but are a counteracting pair.

Both of these problems need to be dealt with for completely specifying a strategy for iterated belief change. Let us look at them in turn.

2.1. Escaping from inconsistency

We propose a very straightforward solution to this problem. We will assume that once an agent has plunged into logical inconsistency, she loses all the genuine information she had acquired so far and so starts her epistemic life anew. This is a rather drastic measure, but our purpose is to provide an escape route from inconsistency so that a strategy for iterated belief change can be fully specified. Accordingly we propose:

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

This is perhaps not the best way to deal with this situation. There are other means available to recover from such inconsistency. For instance, Rott [23] provides a mechanism for extracting useful information even after plunging into an inconsistent belief state. This framework, however, involves explicit representation of the "basic beliefs"; our framework does not have such richness in structure. It is certainly possible to obtain a more interesting means of escaping from inconsistency without assuming such rich representation of beliefs by recourse to certain rather "artificial" means. For instance, as we briefly discuss in Section 5, we can escape from inconsistency by assuming an "impossible world" apart

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.8 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 8

from all the possible worlds. For the purpose of this paper, however, we merely recognise that it is a special case, and treat it as such.

Darwiche and Pearl themselves [7] claim that they have dealt with this problem. They have slightly altered the definition of Spohn's Ordinal Conditionalisation Functions (OCF's) [25] for this purpose. Spohn assumes that at least one world is assigned the κ-value of 0; that is, the knowledge of the agent in question is always satisfiable. Darwiche and Pearl dispense with this assumption, thus allowing belief states with inconsistent belief sets. This they do in the context of showing that their framework can be modeled via OCF's. However, their framework is completely relational, and uses total preorders instead of Spohn's κ functions. Since a total preorder does not allow the set of minimal elements to be empty, it is not clear how their framework can deal with inconsistent belief sets, or for that matter with self-contradictory evidence.

2.2. Problem of counteracting evidence

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. Thus, x and y here are a counteracting pair. They are moreover mutually consistent. 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 $x \wedge y \nvdash \bot$, 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 $x \wedge y \nvdash \neg x$, $\neg x \notin K^*_y$ 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 $x \wedge y \nvdash \bot$ then $x \in (K^*_x)^*_y$.

2.3. Need for a dynamic revision operation

In Section 1 we noticed that use of a dynamic revision operation can block unwanted conflict between DP2 and the AGM postulates. In this section we provide further reason for adopting such a dynamic revision operation. In particular, if we want to retain #Recalcitrance, then we have to discard the notion of a fixed belief revision operator.[6]

[6] Alternatively we can, like Darwiche and Pearl, discard the basic AGM assumption that the revision operation is applied to a belief set.

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) P.9 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 aij1970 by:violeta p. 9

We have seen that #Recalcitrance appears to be a well justified principle. However, it turns out to be rather strong. The following impossibility result shows that there is no revision operation $*$ satisfying the AGM postulates and #Recalcitrance.

Theorem 3. *No belief revision operation $*$ that satisfies the AGM postulates satisfies #Recalcitrance.*

Proof.[7] Assume, for *reductio*, that $*$ is a revision operation that satisfies the AGM postulates and #Recalcitrance. Let the belief set K contain two atomic sentences a and b. Then, by the AGM postulates, $K^*_a = K^*_b = K$. Hence, $(K^*_a)^*_{\neg a \vee \neg b} = (K^*_b)^*_{\neg a \vee \neg b} = K^*_{\neg a \vee \neg b}$. Since each of a and b are individually consistent with $\neg a \vee \neg b$, it follows from #Recalcitrance that either of them is in $K^*_{\neg a \vee \neg b}$. On the other hand, by Success, $\neg a \vee \neg b$ is in $K^*_{\neg a \vee \neg b}$. Hence $K^*_{\neg a \vee \neg b} \vdash \bot$. But, by Consistency, since $\neg a \vee \neg b \nvdash \bot$, it is required that $K^*_{\neg a \vee \neg b} \nvdash \bot$. Contradiction.

This result has two[8] alternative explanations:

(1) #Recalcitrance is too strong. We should reject/weaken it in a judicious manner.
(2) When we revise our current belief set K to a new belief set K^*_x, it is not only the belief set K that changes; the revision operation involved, namely $*$, also undergoes modification. The revision operation used to revise the "prior" belief set K to K' and the revision operation that would be used to revise the "posterior" belief set K' to, say, K'' are, in general, different. Hence the derivation of $(K^*_a)^*_{\neg a \vee \neg b} = (K^*_b)^*_{\neg a \vee \neg b} = K^*_{\neg a \vee \neg b}$ from $K^*_a = K^*_b = K$ is illegal.

Of the above two explanations, we adopt the second one. We have noted earlier that use of a dynamic revision operation can prevent the conflict between DP2 and the AGM framework. Our position receives further support from two different considerations:

(a) competing accounts of belief change such as Boutilier's [4] must be interpreted as implicitly using dynamic belief revision operations in order to make sense;
(b) the notion of a fixed belief revision operation is unintuitive and overly restrictive.

First we demonstrate that the account of iterated belief change proposed by Boutilier which has drawn significant attention must be interpreted as using a dynamic belief revision operation in order to make sense. It is not necessary to go into the detailed structure

[7] For an alternative proof of this theorem, see [23, pp. 134–135]. We thank Hans Rott for pointing out a minor error in the proof originally provided in [17].

[8] . . . or more, depending on what counts as an *alternative* explanation. Many works in the area [18,19,25,26] directly deal with the change in the belief *state* (as opposed to belief set) that incorporates a selection mechanism. Change in the belief state in this sense corresponds to the change in the revision operation, and hence is subsumed under the second item below. We choose to underplay this explanation since we have not yet officially introduced the notion of a belief state, and the notion of belief state played no role in generating the contradiction at hand.

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.10 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 10

of Boutilier's account; all we need is the following consequence of his assumptions [4, Theorem 7, p. 524]:

Theorem 4 (Boutilier, 1993). *Let $A_1, \ldots, A_n$ be a revision sequence with one incompatible update A_n. Then $((K^*_{A_1})^*_{A_2}\ldots)^*_{A_n} = K^*_{A_1\wedge\cdots A_k\wedge A_n}$ 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 Theorem 4.

Observation 1. *Let $*$ be a fixed revision operation satisfying Boutilier's 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' no matter how she came to the beliefs in 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. Many works in the literature including Spohn [25], Nayak [18], Nayak et al. [19], Williams [26] and of course, Darwiche and Pearle [7] provide models of belief state change to account for this. We aim at attaining effectively the same goal by changing the belief sets instead; but this cannot be done if the belief revision operation is fixed once and for all. Hence the need for a dynamic belief revision operation $*$. It is important, however, to keep in mind that we are not studying a different phenomenon, but offering a different perspective of the same phenomenon.

2.4. Nature of dynamic revision

In light of the above discussion, we submit 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:

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.11 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 11

- **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 constraints on $|$ rather than on $*$, and in fact jointly defining the meta-revision operation $|$ itself.

It is important to be clear what the last two postulates really imply. Consider for instance a sequence of observations $e_1, e_2, \ldots, e_n, \neg e_n$ where $e_1, e_2, \ldots, e_n$ are atomic. It would appear that the following is sound reasoning: $((K^*_{e_1})\ldots)^{((*|e_1)\ldots)|e_{n-1}}_{e_n} = K^*_{e_1\wedge\cdots\wedge e_n}$, by (7^*new). Hence K revised by the sequence $e_1, e_2, \ldots, e_n, \neg e_n$, using (8^*new), results in $K^*_{\neg e_n}$, thus washing out all the previous evidence. The first step in this reasoning is legal (by induction). The second step is however fallacious reasoning since it uses indiscriminate use of Extensionality that dynamic revision is designed to block. On the other hand, we can see that the desired result is $K^*_{e_1\wedge\cdots\wedge e_{n-1}\wedge\neg e_n}$. Let us take a simpler example. Consider the observation sequence $a, b, \neg b$, where a and b are atoms. We need to compute $((K^*_a)^{*'}_b)^{*''}_{\neg b}$ where $*' = *|a$ and $*'' = *'|b$. We can use (7^*new) first, and get

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.12 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 12

$((K_a^*)_b^{*'})_{\neg b}^{*''} = (K_{a\wedge b}^*)_{\neg b}^{(*|a)|b}$. But we cannot use (8*new) to simplify it any further now, since $*|a \wedge b$ is different from $(*|a)|b$. On the other hand, we can apply (8*new) as the first step and infer that $((K_a^*)_b^{*'})_{\neg b}^{*''} = (K_a^*)_{\neg b}^{*|a}$. We can now apply (7*new) to simplify it further and obtain $(K_a^*)_{\neg b}^{*|a} = K_{a\wedge\neg b}^*$ whereby $((K_a^*)_b^{*'})_{\neg b}^{*''} = K_{a\wedge\neg b}^*$.

2.5. Dynamic revision: binary or unary?

So far we have argued that there are good reasons to view the revision operation $*$ as a dynamic operation. This conclusion presupposed two important assumptions: (1) the revision operation $*$ is a binary operation, and (2) the revision operation $*$ should be well equipped to deal with iterated revision. The second assumption is not controversial, since iterated belief revision is a declared goal of this enterprise. However, as we pointed out in Section 1, opinion is divided in the research community as to whether or not $*$ should be viewed as a binary operation. We revisit the issue here, since, as we will soon see, the dynamic revision must be interpreted as a unary operation.

Let us see why the assumption that $*$ is a binary operation is so intuitively appealing. The job of the operation $*$ is to provide a mechanism to accommodate new information. It is tempting to think that each epistemic agent is blessed with her own revision operation, and no matter what her current beliefs are, can use this operation to revise them in light of new information. Since theoretically any member $K \in \mathbf{K}$ could be the agent's current belief set, there is good reason to assume that $*$ must be a binary operation taking an arbitrary belief set as its first argument.

This motivation, however, runs counter to the fact that $*$ is a dynamic operation. In order to illustrate this point, let us consider a trivial language $\mathcal{L}_T$ generated by two atoms a and b. Accordingly, there are sixteen members in $\mathbf{K}$ (corresponding to the sixteen possible sentences in this language modulo truth-functional equivalence), say $K_1, \ldots, K_{16}$. Let us assume that the agent's initial belief set is $K_1 = Cn(\{a, b\})$, and her initial revision operation is $*$. Now, since $*$ is a binary operation, it maps any belief set in $\mathbf{K}$ and any sentence in $\mathcal{L}_T$ to a belief set. Let's say the first piece of evidence the agent accepts is $\neg a$, and her revision operation $*$ happens to map $\langle Cn(\{a, b\}), \neg a\rangle$ to $K_2 = Cn(\{\neg a, b\})$, i.e., $(K_1)_{\neg a}^* = K_2$. In the process, of course, the revision operation $*$, being a dynamic one, evolves to $*|\neg a$. We of course know that this new operation $*|\neg a$ can be used to revise the new belief set $(K_1)_{\neg a}^* = K_2$. However, we have to grant that $*|\neg a$ is also a binary operation, and that this operation also maps every belief set, sentence pair to a belief set. For instance, we should be able to meaningfully talk about the belief-set $(K_{12})_b^{*|\neg a}$. But we cannot do that without ensuring that K_{12} is the result of revising some "initial" belief set by $\neg a$—and obviously K_1 was *not* this initial belief set. The emerging picture is that, while migrating from $*$ to $*|\neg a$, we not only revise the *actual* initial belief set K_1, but also other members $K_2, \ldots, K_{16}$ of $\mathbf{K}$ which potentially could have been the initial belief set. How we can systematically revise such potential belief sets appears to be problematic.

That, however, is not the end of our troubles. Since the operation $*|\neg a$ can take *any* member of $\mathbf{K}$ as its first argument, we must ensure that for every member $K_j \in \mathbf{K}$ there is a member $K_i \in \mathbf{K}$ such that $(K_i)_{\neg a}^* = K_j$. In other words, the unary function $*_{\neg a}$ constructed from $*$ by fixing its second argument to $\neg a$ must be a 1–1 mapping from

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.13 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 13

K to **K**. But that is a big ask! There is nothing special about our choice of $\neg a$ as the first piece of evidence we chose. It could have been any other sentence of the language $\mathcal{L}_T$. So what is required is that, for any sentence x of $\mathcal{L}_T$, the initial revision operation $*$ must be such that the unary function $*_x$ constructed from it by fixing its second argument to x must be a 1–1 mapping from **K** to **K**. But that is not possible, specifically because $*_\bot$ is a constant function mapping every member of **K** to $K_\bot$! We hasten to add that no special appeal was made in this argument to the fact that we chose $\mathcal{L}_T$ for the purpose of this illustration—the argument will go through no matter which language is at issue.

The above argument conclusively shows that, contrary to our initial assumption, a dynamic belief revision operation cannot be a binary revision operation. Further reasons for choosing unary revision operations over binary operations can be found in Rott's [24]. From here onwards, we assume that the revision operation $*$ is actually a unary operation, taking only an evidential sentence as its argument, and that the K in the postulates of belief revision is a fixed belief set; i.e., $* : \mathcal{L} \to \mathbf{K}$. Dropping the reference to K is not a handicap, since it can be obtained from the (initial) unary operation $*$ as follows: $*(\top) = K$. Thus one may view the operation $*$ as the belief state itself, since it not only provides the means to change one's beliefs, but encodes the beliefs as well.

It is worth noting that there are two operations involved when an agent revises her beliefs. There is the (unary) revision operation $*$ that determines what the agent's beliefs would be after the revision step: $*(x)$ is the new belief set. There is a (binary) meta-revision operation $|$ that determines what the next revision operation would be: $*|x$ is the new revision operation. It is tempting to draw an analogy with Bayesian Updating which predates (and has significant connection with) belief revision.

According to the Bayesian view, an agent's belief state is represented as a Probability function, P. According to the view we are espousing, it is represented by the unary revision operation $*$. In both the frameworks, the (full) beliefs of the agent can be obtained from the respective belief states: it is the set of sentences that have probability 1 (full probability)[9] in the former, and $*(\top)$ in the latter. Both the frameworks also allow for iterated belief-state change. There is however a slight twist.

The way iterated belief change is dealt with in the probabilistic framework can be viewed in two different ways. One way of looking at it is that the belief state is represented as an absolute probability function, say initially P_0. When a piece of evidence e_i (for $i > 0$) is received, the relevant belief state P_{i-1} gets updated to the new belief state P_i as follows: $P_i(x) = P_{i-1}(x \mid e_i)$. Thus the job of the probability function P is to determine what is believed (to what extent), whereas a state-transition function $|$ is used to modify the belief state P. There is however another way of looking at it. The belief state is represented not by an absolute measure, $P(.)$, but by a conditional probability measure $P(. \mid -)$, where the second argument is any (possibly empty) sequence of evidential sentences. The rules of probability calculus determine the probability due to any sentence, given any sequence of evidence (with the standard proviso necessary to avoid division by 0). The initial (absolute) probability measure P_0 would correspond, in this view, to the state $P(. \mid \mathbf{0})$

[9] $\{x \mid P(x) = 1\}$ if absolute probability measure is being used, $\{x \mid P(x \mid \Omega) = 1\}$ if a conditional probability measure is being used.

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.14 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 14

where **0** represents the empty sequence. The (absolute) probability measure P_i in this view will correspond to the state $P(.\mid\langle e_1,\ldots,e_{i-1}\rangle)$. The interesting thing to observe here is the notational simplicity achieved, and the way it hides the dynamic nature of the belief state under the veneer of the fixed (but parametrised) probability function. According to one school of thought [22], it is the notion of *conditional probability* that is the primitive notion of probability; the notion of absolute or unconditional probability is a derived notion obtained from the conditional probability by conditionalising with respect to the whole event-space. On the other hand, traditionally the notion of absolute probability, due to its intuitive simplicity, is taken to be the primitive one, and that of the conditional probability as a derived notion. Both ways of looking at probability are equivalent.

Analogously, there are alternative ways of looking at the dynamic belief revision operation $*$. The notation we have been using is suggestive of the former perspective: the belief state is represented as a unary revision function, say initially $*_0$, that takes an evidential sentence as its argument. When a piece of evidence e_i (for $i > 0$) is received, the then current belief state $*_{i-1}$ gets updated to the new belief state $*_i$ as follows: $*_i = *_{i-1}|e_i$. A revision step is carried out in two sub-steps. Given current belief set $K_{i-1} = *_{i-1}(\top)$ and evidence e_i,

$$K_i := K_{e_i}^{*_{i-1}},$$

$$*_i := *_{i-1}|e_i.$$

The other perspective of belief revision has been advocated by Rott [24]. According to this view, the analogue of the latter perspective on probabilistic representation of belief states, a belief state is represented by a unary revision operation $*$ that takes a sequence of evidential sentences as its argument. The initial belief set K_0 in this framework corresponds to the set $*(\mathbf{0})$, and the belief set K_i corresponds to $*(\langle e_1,\ldots,e_i\rangle)$. Again, here the notation hides the fact that the revision operation undergoes change.

One might wonder if we cannot unofficially drop the | notation after officially introducing it for the sake of readability? Surely, we can read off from the unofficial expression $(K_x^*)_y^*$ that the second occurrence of $*$ is a short-hand for $*|x$! It is easy to notice that tempting though this proposal is, it is still worthwhile to stick to our official notation. Imagine what would happen if we allowed the proposed unofficial notation in the example we considered earlier at the end of Section 2.4, namely $((K_a^*)_b^{*'})_{\neg b}^{*''}$. This expression becomes $((K_a^*)_b^*)_{\neg b}^*$ in the unofficial notation. Applying (8*new) followed by (7*new) we get from this expression $K_{a\wedge\neg b}^*$ as desired. But on the other hand, if we apply (7*new) first, we obtain from $((K_a^*)_b^*)_{\neg b}^*$ the expression $(K_{a\wedge b}^*)_{\neg b}^*$. But now there is nothing to prevent us from applying (8*new) and obtain the incorrect result $K_{\neg b}^*$—unless, at least, we introduce extra rules to make such application of (8*new) illegal. Thus, although from the original (unofficial) expression, $((K_a^*)_b^*)_{\neg b}^*$, we can read off the intent behind different occurrences of the operation $*$, book-keeping becomes difficult in the inferential process. Accordingly we suggest that the official notation should be retained.

Let us now reconsider what we have done. We assumed a binary belief revision operation from the start. We noticed that there are associated problems in dealing with iterated belief change and argued for viewing this operation as a dynamic one, and argued for a set of constraints that this operation should satisfy. But later we noticed that this

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.15 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 15

operation cannot be a binary operation after all, and decided to reinterpret it as a unary operation. But if we started from a unary operation from the beginning, we probably would not have faced any problem in the first place. So what good are the constraints that we imposed on the dynamic revision operation? The fact is, if we interpreted $*$ as a unary operation right from the start, then the Classical AGM account does not tell us *anything* at all about iterated belief change. Even **AGM-It** would not follow from the eight AGM postulates. So we would need extra postulates to deal with iterated belief change. The Darwiche–Pearle postulates are definitely good candidates to consider. Even if we consider the revision operation $*$ to be a unary operation, the D–P postulates remain over-permissive, and need to be strengthened. The extra constraints we imposed do precisely that.

The account we have provided, though principled, is not constructive in nature. The value of the later approach by Darwiche and Pearl lies in that it attempts to provide a construction for this dynamic revision operator. It however fails because it does not fully specify how this new operator is to be constructed. In the next two sections we provide two different ways of constructing the dynamic belief revision operator. In Section 3 we will examine how the dynamic belief revision can be constructed via revision of a belief state represented as an entrenchment relation. In Section 4 we will analogously do the same when a belief state is represented as a total preorder over models of a language, albeit in an indirect manner. For the sake of readability, although $*$ is intended to be a unary operation, we still use expressions such as K^*_x. In such expressions, K should be construed as the belief set associated with $*$, namely, $*(\top)$.

3. Character of entrenchment kinematics

In this section we provide a constructive account of dynamic belief change presented in the last section via the account of entrenchment kinematics offered by Nayak et al. [19]. In particular, we will look at how the account of dynamic belief revision offered in the previous section can be characterised in terms of change in belief state, with the assumption that a belief state is represented as an epistemic entrenchment relation. We will also see how the problematic Example 2 is dealt with in terms of change in epistemic entrenchment.

3.1. Entrenchment kinematics

In [19] a belief state is represented as an epistemic entrenchment (EE) relation which is slightly different from the standard epistemic entrenchment (SEE) relation introduced in AGM [10]. An EE relation, $\preceq$, is defined as any relation over the language $\mathcal{L}$ that satisfies the following four constraints:

(*EE*1) If $x \preceq y$ and $y \preceq z$ then $x \preceq z$ (transitivity)

(*EE*2) If $x \vdash y$ then $x \preceq y$ (dominance)

(*EE*3) For any x and y, $x \preceq x \wedge y$ or $y \preceq x \wedge y$ (conjunctiveness)

$$(EE4)\quad \text{Given that } \bot \prec z \text{ for some } z, \quad \text{if } y \preceq x \text{ for all } y, \text{ then } \vdash x \qquad \text{(maximality).}$$

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

Definition 5 (*Epistemic content*).

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

We note that the epistemic content of an EE relation $\preceq$ is inconsistent, i.e., $EC(\preceq) = \mathcal{L}$, when for every sentences $x, y \in \mathcal{L}$: $x \preceq y$. We denote this absurd EE relation by $\preceq_\bot$.

Belief change in this framework is captured by providing an account of how the EE relation $\preceq$ changes in light of new information. 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 6 (*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

$$\text{either } \vdash y$$
$$\text{or both } \nvdash x \text{ and } x \preceq y$$

(c) $e \nvdash x$ and $(e \to x) \preceq (e \to y)$.

The following consequences of this definition may be taken to provide the justification for it:

(1) If $e \vdash \bot$ then $x \preceq_e^\odot y$ for every sentences x and y. This leads to an inconsistent epistemic content, and hence accords well with the AGM account. In fact it follows that $\preceq_e^\odot = \preceq_\bot$ iff $e \vdash \bot$.
(2) Given that $\preceq$ is absurd, i.e., $\preceq = \preceq_\bot$, and $e \nvdash \bot$,
    (a) If $e \nvdash x$ then $x \preceq_e^\odot y$ for all y. This is so because e is the only knowledge the agent has so far. Accordingly, non-consequences of e are unknown, hence minimally entrenched.
    (b) If $e \vdash x$, $\nvdash x$ and $\nvdash y$ then $x \preceq_e^\odot y$ iff $e \vdash y$. This condition essentially says that non-trivial consequences of the evidence are now believed on an equal footing. This is rational given that there was no knowledge in $\preceq_\bot$.
(3) Given that $\preceq$ is non-absurd, $\nvdash e$, $\nvdash x$ and $\nvdash y$,
    (a) If $e \vdash x$ and $e \vdash y$ then $x \preceq_e^\odot y$ iff $x \preceq y$. That is, the relative entrenchment among the consequences of the evidence e are not affected.

ARTICLE IN PRESS
S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.17 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 17

(b) If $e \nvdash x$ and $e \nvdash y$ then $x \preceq_e^{\odot} y$ iff $e \rightarrow x \preceq e \rightarrow y$. This condition says that the entrenchment among the non-consequences x and y of the evidence is determined by the prior entrenchment of the relative information in x and y with respect to the evidence e.

(c) If $e \nvdash x$ and $e \vdash y$ then $x \prec_e^{\odot} y$. This condition tells us that the consequences of the evidence have higher priority than its non-consequences.

3.2. *Example 2 revisited*

As a litmus test, let us apply this account of entrenchment revision to Example 2. Towards this end 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 7 (*EE to* $*$). Let $\preceq$ be an *EE* relation and $K = EC(\preceq)$. Define $*_{\preceq}$ as: $y \in K_x^{*_{\preceq}}$ iff either $(x \rightarrow \neg y) \prec (x \rightarrow y)$ or $x \vdash y$.

Definition 8 ($*$ *to EE*). Let $*$ be a belief revision operation satisfying the eight AGM postulates where $K = *(\top)$. Define $\preceq_*$ as: $x \preceq_* y$ iff either $x \notin K^*_{\neg x \vee \neg y}$ or $\vdash x \wedge y$ or $K = K_{\perp}$.

$\preceq_*$ is demonstrably an EE relation with epistemic content K. On the other hand, the operation $*_{\preceq}$ is also demonstrably a revision operation satisfying the AGM postulates 1*–8*.

Entrenchment Kinematics and Example 2. We assume a language generated by two atoms *bird* and *sing*. Our initial belief state K is $Cn(\{bird, sing\})$. We assume that the EE relation associated with K is $\preceq$. Our goal is to show that if $\preceq$ is first revised by $\neg bird$ and then by $\neg sing$, the sentence $\neg bird$ will still be accepted in the resultant belief state, i.e., $\neg bird$ will be a member of the belief set associated with the EE relation $(\preceq^{\odot}_{\neg bird})^{\odot}_{\neg sing}$. By Definition 7, and the fact that $\preceq^{\odot}_{\neg bird}$ is a connected relation, it will be sufficient to show that $(\neg sing \rightarrow \neg bird) \npreceq^{\odot}_{\neg bird} (\neg sing \rightarrow \neg\neg bird)$. Since $\neg bird \nvdash \perp$, $\neg bird \vdash (\neg sing \rightarrow \neg bird)$ and $\neg bird \nvdash (\neg sing \rightarrow bird)$, the desired result follows from Definition 6.

Thus we notice that the entrenchment revision mechanism advocated here has the desirable feature we need. We will now state some formal results to the effect that this revision mechanism has the appropriate formal properties.

3.3. *Dynamic belief revision and entrenchment kinematics*

It is shown in [19] 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 [19].

Theorem 9. *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. Then* $(\preceq_*)^{\odot}_e = \preceq_{*|e}$.

This result shows that given two revision operations $*$ and $*|e$ satisfying 0^*–8^*(new), Definition 8 generates from them two entrenchment relations, the latter of which is the revision of the former by the evidence e in accordance with Definition 6.

Theorem 10. *Let $\preceq$ be an EE relation 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 $K = *(\top) = EC(\preceq)$.*

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 6, we can generate from them revision operations $*$ and $*|e$ with the help of Definition 7 which satisfy postulates 0^*–8^*(new).

Two Theorems 9 and 10, together, provide what may be seen as a representation result connecting the account we offer in this paper and the account of entrenchment kinematics offered in [19]. 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 Definitions 7 and 8 consecutively).

Theorem 11. *Let $*$ and $*|e$ be two revision operations satisfying the postulates* 0^*–8^*(*new*). *Let $K = *(\top)$ and $K' = K^*_e$. Then the revision operations $*_{\preceq_*}$ and $*_{\preceq_{*|e}}$ are belief revision operations that satisfy* 0^*–8^*(*new*).

As expected, when we reconstruct a pair of revision operations from another via an entrenchment relation, the appropriate relation between the reconstituted revision operations hold.

Theorem 12. *Let $\preceq$ be an EE relation, and $\preceq^{\odot}_e$ the result of revising $\preceq$ by a sentence e. Then $\preceq_{*_{\preceq}} = \preceq$ and $\preceq_{*_{\preceq^{\odot}_e}} = \preceq^{\odot}_e$.*

This result is not surprising, since it easily follows from results in the AGM literature [10], but we list it here only for the sake of completeness in exposition. On the basis of these results, we conclude that the postulates 0^*–8^*(new) characterise the account of entrenchment kinematics discussed in [19].

4. A semantic characterisation

In the last section we examined how the account of iterated belief revision we offer can be characterised in terms of change in epistemic entrenchment. In the literature [5,7], belief revision is often also given a semantic characterisation in terms of a plausibility ordering over the interpretations generated by the background language. Accordingly, our approach can also be semantically captured by an account of how the plausibility ordering over the interpretations is modified in light of new evidence. We will borrow the semantic framework adopted by Darwiche and Pearl for ease of comparison. In the process we will explain what it is that prevents Darwiche and Pearl offering the intuitive solution to Example 2, but allows us to do exactly that.

4.1. Semantics of belief change

Definition 13. Let Ω be the set of possible worlds (interpretations) of the background language $\mathcal{L}$ and $\sqsubseteq$ a total preorder (a connected, transitive and reflexive relation) over Ω. For any set $\Sigma \subseteq \Omega$ and world ω we will say ω is a $\sqsubseteq$-minimal member of Σ if and only if both $\omega \in \Sigma$ and $\omega \sqsubseteq \omega'$ for all $\omega' \in \Sigma$.

By $\omega_1 \sqsubseteq \omega_2$ we will understand that ω_2 is not more plausible than ω_1. The expression $\omega_1 \equiv \omega_2$ will be used as a shorthand for ($\omega_1 \sqsubseteq \omega_2$ *and* $\omega_2 \sqsubseteq \omega_1$). The symbol $\sqsubset$ will denote the strict part of $\sqsubseteq$. For any set $S \subseteq \mathcal{L}$ we will denote by $[S]$ the set $\{\omega \in \Omega \mid \omega \models s$ for every $s \in S\}$. For readability, we will abbreviate $[\{s\}]$ by $[s]$.

Intuitively, the preorder $\sqsubseteq$ will be the semantic analogue of an EE relation. They both represent the belief states of an agent. Just as $EC(\preceq)$, the epistemic content of $\preceq$, captures the belief set associated with $\preceq$, similarly we will say that $K_{\sqsubseteq}$ is the belief set associated with the preorder $\sqsubseteq$. It is defined as the set of sentences satisfied by the $\sqsubseteq$-minimal worlds, i.e.,

$$K_{\sqsubseteq} = \{x \in \mathcal{L} \mid \omega \models x \text{ for all } \sqsubseteq\text{-minimal } \omega \in \Omega\}.$$

A special case has to be made to represent the belief state of an agent who has courted inconsistency. From this agent's point of view, effectively there are no "most plausible" scenarios—the set of most plausible worlds is the null set. We will represent the belief state of this agent not by a total preorder $\sqsubseteq$ but by the empty relation $\sqsubseteq_{\perp}$: for every pair $\omega, \omega' \in \Omega$, $\omega \not\sqsubseteq_{\perp} \omega'$. Since $\sqsubseteq_{\perp}$ does not have any minimal element, it is clear that the belief set associated with it is $K_{\perp}$, as we would expect. As far as we are concerned, once the agent has reached the state $\sqsubseteq_{\perp}$, she has lost all information, and needs to start her epistemic life from scratch.

Now that we know how to extract the belief set associated with a belief state $\sqsubseteq$, we might want to know how belief revision itself can be captured in this semantic framework. It is well known that an AGM-rational belief revision operation $*_{\sqsubseteq}$ based on a total preorder $\sqsubseteq$ can be constructed as follows [12]

$$(\textbf{Grove}) \quad x \in K_e^{*_{\sqsubseteq}} \text{ iff } \omega \models x \text{ for every world } \omega \text{ that is } \sqsubseteq\text{-minimal in } [e].$$

Note that this definition does not take care of the situation where the belief state is the absurd one $\sqsubseteq_{\perp}$. To deal with this scenario, we modify Grove as follows:

Definition 14 ($\sqsubseteq$ *to* $*$).

$$x \in K_e^{*_{\sqsubseteq}} \quad \text{iff} \quad \begin{cases} [e] \subseteq [x] & \text{if } \sqsubseteq \; = \; \sqsubseteq_{\perp}, \\ \omega \models x \text{ for every } \omega \sqsubseteq\text{-minimal in } [e] & \text{otherwise.} \end{cases}$$

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.20 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 20

4.2. Iterated belief change: two received accounts

The reason the AGM account of belief revision does not cope well with iterated belief change is that in general the resultant belief set $K_e^{*_\sqsubseteq}$ is not appropriately associated with the original preorder $\sqsubseteq$. In particular, since $[K_e^{*_\sqsubseteq}]$ is not the set of $\sqsubseteq$-minimal elements in Ω, a more appropriate belief state must be constructed to be associated with $K_e^{*_\sqsubseteq}$. Different accounts of iterated belief change are effectively accounts of generating a new preorder $\sqsubseteq_e^\circ$ based on the original preorder $\sqsubseteq$ and evidence e that can be appropriately associated with the new belief set $K_e^{*_\sqsubseteq}$, i.e., $[K_e^{*_\sqsubseteq}]$ is exactly the $\sqsubseteq_e^\circ$-minimal elements of Ω. In other words, all these approaches attempt to supplement the following condition on the preorder revision operation $\circ$:

Faith. If ω_1 is $\sqsubseteq$-minimal in $[e]$, then ω_1 is $\sqsubseteq_e^\circ$-minimal in Ω

with further constraints so that a total preorder $\sqsubseteq_e^\circ$ is fully specified. It is easily noticed that the condition Faith is mandated by Definition 14 presented earlier together with the view that the belief set $K_\sqsubseteq$ associated with any preorder $\sqsubseteq$ is the set of sentences satisfied by the $\sqsubseteq$-minimal worlds. Hence Faith is indeed required by the classical AGM account of belief change. Darwiche and Pearl [6] impose the following additional constraints on $\circ$:

DPS1. If $\omega_1 \models e$ and $\omega_2 \models e$ then $\omega_1 \sqsubseteq_e^\circ \omega_2$ iff $\omega_1 \sqsubseteq \omega_2$.

DPS2. If $\omega_1 \models \neg e$ and $\omega_2 \models \neg e$ then $\omega_1 \sqsubseteq_e^\circ \omega_2$ iff $\omega_1 \sqsubseteq \omega_2$.

DPS3. If $\omega_1 \models e$ and $\omega_2 \models \neg e$ then $\omega_1 \sqsubset_e^\circ \omega_2$ if $\omega_1 \sqsubset \omega_2$.

DPS4. If $\omega_1 \models e$ and $\omega_2 \models \neg e$ then $\omega_1 \sqsubseteq_e^\circ \omega_2$ if $\omega_1 \sqsubseteq \omega_2$.

It is easily noticed that these four constraints together with Faith are not sufficient to uniquely determine the revised preorder $\sqsubseteq_e^\circ$. In fact Darwiche and Pearl point out in [7] that Boutilier's account of Natural Revision [4] can be obtained by adding further constraints (and fault it as being a bit of an overkill). However, the reason Darwiche and Pearl's account fails to satisfactorily deal with Example 2 is precisely the under-specification of the new preorder. To revisit Example 2, let us assume that the language $\mathcal{L}$ is generated by two atoms, *bird* and *sing*, and accordingly supports four interpretations (worlds). Let these four worlds be $\omega_1 - \omega_4$ as follows:

$$\omega_1 \models bird \wedge sing, \quad \omega_2 \models \neg bird \wedge sing, \quad \omega_3 \models bird \wedge \neg sing \quad \text{and}$$
$$\omega_4 \models \neg bird \wedge \neg sing$$

so that $[bird] = \{\omega_1, \omega_3\}$ and $[sing] = \{\omega_1, \omega_2\}$. According to Example 2, then, $K = Cn(\{bird, sing\})$, $K^*_{\neg bird} = Cn(\{\neg bird, sing\})$ and $K^*_{\neg sing} = Cn(\{bird, \neg sing\})$. It is easily verified that the following are the only three preorders that are consistent with this example, given Faith:

Preorder1. $\omega_1 \sqsubset \omega_2 \sqsubset \omega_3 \sqsubset \omega_4$.

Preorder2. $\omega_1 \sqsubset \omega_3 \sqsubset \omega_2 \sqsubset \omega_4$.

Preorder3. $\omega_1 \sqsubset \omega_2 \equiv \omega_3 \sqsubset \omega_4$.

ARTICLE IN PRESS
S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 aij1970 P.21 (1-36) by:violeta p. 21

Since Darwiche and Pearl's account is built on top of Faith, the above inference is in accordance with their account as well. Now we are in a position to examine Example 2 from the perspective of the D–P account.

Example 2 and the D–P Account. Example 2 and the D–P Account In order to get the desired result that $(K^*_{\neg bird})^*_{\neg sing} = Cn(\{\neg bird, \neg sing\})$ we need ω_4 to be the only $\sqsubset^{\circ}_{\neg bird}$-minimal world in $[\neg sing]$, that is we require that $\omega_4 \sqsubset^{\circ}_{\neg bird} \omega_3$. We further note that no matter which of the three preorders, Preorder1–Preorder3, is the correct one, $\omega_3 \sqsubset \omega_4$. So while we have $\omega_3 \sqsubset \omega_4$, we need $\omega_4 \sqsubset^{\circ}_{\neg bird} \omega_3$ for the example to go through. A brief look at the four constraints imposed by Darwiche and Pearl shows that such reversal of ordering among worlds is not guaranteed (although permitted). In thus not tightly specifying $\sqsubseteq^{\circ}_{e}$, Darwiche and Pearl's account suffers from the burden of over-permissiveness.

That Darwiche and Pearl's account is over-permissive is evidenced by the fact that it is consistent with Boutilier's Natural Revision which has undesirable consequences. Boutilier's Natural Revision is obtained by supplementing Faith and (DPS1–DPS4) with the following two additional constraints:

NR1. Given that $\omega_1, \omega_3 \models e$ and $\omega_2 \models \neg e$, if $\omega_3 \sqsubset \omega_1$ and $\omega_2 \sqsubseteq \omega_1$ then $\omega_2 \sqsubseteq^{\circ}_{e} \omega_1$.

NR2. Given that $\omega_1, \omega_3 \models e$ and $\omega_2 \models \neg e$, if $\omega_3 \sqsubset \omega_1$ and $\omega_2 \sqsubset \omega_1$ then $\omega_2 \sqsubset^{\circ}_{e} \omega_1$.

It is easily verified that together with Faith and (DPS1–DPS4), these two additional constraints fully specify the new preorder $\sqsubseteq^{\circ}_{e}$. However, as it turns out, Natural Revision does not give the right answer to Example 2 either!

Example 2 and Natural Revision. We note that $\omega_2, \omega_4 \models \neg bird$ and $\omega_3 \models bird$. Moreover, in each of the three preorders (Preorder1–Preorder3), both $\omega_2 \sqsubset \omega_4$ and $\omega_3 \sqsubset \omega_4$. It is now easily verified that if we apply Natural Revision to Example 2, no matter which of the three preorders is the prior, by NR2, after revision by $\neg bird$ we get $\omega_3 \sqsubset^{\circ}_{\neg bird} \omega_4$. This is exactly the opposite of what we need, namely, $\omega_4 \sqsubset^{\circ}_{\neg bird} \omega_3$. Hence Natural Revision fails to provide the appropriate solution to Example 2.

In fact, irrespective of which preorder (among Preorder1–Preorder3) is the prior, by Faith we get $\omega_2 \sqsubset^{\circ}_{\neg bird} \omega_1$, by DP1 we get $\omega_2 \sqsubset^{\circ}_{\neg bird} \omega_4$, and by DP2 we get $\omega_1 \sqsubset^{\circ}_{\neg bird} \omega_3$. Thus we get $\omega_2 \sqsubset^{\circ}_{\neg bird} \omega_1 \sqsubset^{\circ}_{\neg bird} \omega_3$. From NR2 we get $\omega_3 \sqsubset^{\circ}_{\neg bird} \omega_4$. Thus, altogether, we get $\omega_2 \sqsubset^{\circ}_{\neg bird} \omega_1 \sqsubset^{\circ}_{\neg bird} \omega_3 \sqsubset^{\circ}_{\neg bird} \omega_4$. By Faith, then, ω_3 is $\sqsubset^{\circ}_{\neg bird}$-minimal in $\neg sing$. Accordingly, we get the undesirable result that $(K^*_{\neg bird})^{*|\neg bird}_{\neg sing} = Cn(\{bird, \neg sing\})$.

4.3. Another proposal: simple lexicography

We have noticed that the primary problem with the Darwiche Pearl account is that it, by under-specifying the revised preorder $\sqsubseteq^{\circ}_{e}$, makes room for undesirable accounts of preorder evolution such as Natural Revision. Our goal is to strengthen the D–P account so as to get a more appropriate account of preorder revision.

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.22 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 22

Let us first set up the machinery to deal with the special cases.

SimpSpecial1. If $[e] = \emptyset$ then, and only then, $\sqsubseteq^\circ_e = \sqsubseteq_\bot$.

SimpSpecial2. Else, if $\sqsubseteq = \sqsubseteq_\bot$, then $\omega_1 \sqsubseteq^\circ_e \omega_2$ iff either $\omega_1 \models e$ or $\omega_2 \models \neg e$.

It is clear from SimpSpecial1 that the revised relation $\sqsubseteq^\circ_e$ is the empty relation $\sqsubseteq_\bot$ if and only if $e \vdash \bot$. This is vindicated given that in the AGM account, an agent courts inconsistency exactly if the accepted information itself is inconsistent. The second condition, SimpSpecial2, says that if the agent is in the absurd state $\sqsubseteq_\bot$, and receives some informative evidence e, then the only structure that its belief state would have is that worlds satisfying e are positively preferred to the worlds failing to satisfy e. The agent would be indifferent among worlds in $[e]$ (respectively $[\neg e]$) so that the new state becomes a total preorder.

For the general case, when the prior preorder is nonempty ($\sqsubseteq \neq \sqsubseteq_\bot$) and the evidence is satisfiable ($[e] \neq \emptyset$), we propose the following constraints:

SimpLex1. If $\omega_1 \models e$ and $\omega_2 \models e$ then $\omega_1 \sqsubseteq^\circ_e \omega_2$ iff $\omega_1 \sqsubseteq \omega_2$.

SimpLex2. If $\omega_1 \models \neg e$ and $\omega_2 \models \neg e$ then $\omega_1 \sqsubseteq^\circ_e \omega_2$ iff $\omega_1 \sqsubseteq \omega_2$.

SimpLex3. If $\omega_1 \models e$ and $\omega_2 \models \neg e$ then $\omega_1 \sqsubset^\circ_e \omega_2$.

The SimpLex conditions are specialisations of the lexicographic ordering proposed in [18].[10] It is easily noticed that conditions SimpLex (1–3) completely specify the preorder $\sqsubseteq^\circ_e$. Note that SimpLex1 and SimpLex2 are simply DPS1 and DPS2, and DPS3 and DPS4 are immediate consequences of SimpLex3. Finally, SimpLex guarantees Faith. Hence, it is a strengthening of the D–P account.

Let us now have another look at Example 2, keeping our preferred account of preorder change:

Example 2 and SimpLex. We note that $[bird] = \{\omega_1, \omega_3\}$ and $[\neg bird] = \{\omega_2, \omega_4\}$. It immediately follows from SimpLex3 that $\omega_4 \sqsubset^\circ_{\neg bird} \omega_3$ no matter which of the preorders Preorder1 to Preorder3 is the prior preorder. As we noted earlier, that guarantees that $\neg bird \in (K^*_{\neg bird})^{*|\neg bird}_{\neg sing}$, as desired.

4.4. SimpLex and entrenchment kinematics

The SimpLex approach solves the problem raised by Example 2. However we have yet to formally show that this actually captures the semantics of the dynamic revision operation presented in Section 2. We could proceed in one of two ways for demonstrating this. We can show that, in a manner analogous to what we did in Section 3, the SimpLex method semantically characterises the dynamic belief revision operation discussed in Section 2.

[10] The revision method proposed in [18], dubbed lexicographic revision by many researchers now (see, for instance, [3]), assumed that the evidence itself is a preorder as well. If a piece of "naked evidence", such as a sentence e, is represented as: $\omega \sqsubseteq \omega'$ iff either $\omega \notin [e]$ or $\omega' \in [e]$, then the lexicographic revision conditions reduce to SimpLex.

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.23 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 23

Otherwise, we can exploit the known semantics of epistemic entrenchment. We will adopt the latter strategy.

Definition 15. The trivial EE relation $\preceq_\bot$ and the empty relation $\sqsubseteq_\bot$ are each other's counterpart, and are counterparts of no other relations.

When the *EE* relation $\preceq \neq \preceq_\bot$ and $\sqsubseteq$ is a total preorder over Ω, we will say that $\preceq$ and $\sqsubseteq$ are each other's counterpart just in case the following condition holds:

for every sentences x and y, $x \preceq y$ iff either $[y] = \Omega$[11] or there exists world ω_1 which is $\sqsubseteq$-minimal in $[\neg x]$ and $\omega_1 \sqsubseteq \omega_2$ for every world ω_2 that is $\sqsubseteq$-minimal in $[\neg y]$.[12]

Note that since $\sqsubseteq$ is total, if two distinct worlds ω and ω' are $\sqsubseteq$-minimal in some subset $S \subseteq \Omega$, then $\omega \equiv \omega'$. Hence, given $\preceq$ and its counterpart $\sqsubseteq$, from $x \preceq y$ and any worlds ω_1 and ω_2 that are respectively $\sqsubseteq$-minimal in $[\neg x]$ and $[\neg y]$, it follows that $\omega_1 \sqsubseteq \omega_2$.

It is well known that given an EE relation $\preceq$ and its counterpart preorder $\sqsubseteq$, the revision operations $*_\preceq$ and $*_\sqsubseteq$ generated from them respectively using Definition 7 and Definition 14 are identical [21].[13] In this context, the following result is interesting:

Theorem 16. *Let $\preceq$ be an EE relation and $\sqsubseteq$ either an empty relation $\sqsubseteq_\bot$ or a total preorder over Ω. For every sentence e, the revised EE relation $\preceq_e^\odot$ and the revised preorder $\sqsubseteq_e^\circ$ are each other's counterpart, given that $\preceq$ and $\sqsubseteq$ are each other's counterpart.*

In Section 3 we have shown that entrenchment kinematics provides the recipe for modifying an EE relation in a way that captures the dynamic belief revision operation we want. Now we have shown that entrenchment kinematics (as captured by $\odot$) and preorder revision as represented by $\circ$ take two counterpart relations ($\preceq$ and $\sqsubseteq$, respectively) to relations that are counterparts of each other. Hence, we get the semantic analogues of the results in Section 3 for free. In other words, semantic analogues of Theorems 3–7 hold.

5. Discussion and conclusion

In this paper we have argued that a reasoned account of 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 provided two constructive accounts of dynamic belief revision via entrenchment kinematics and Simple Lexicography.

---

[11] We need this condition to deal with special situations involving empty sets of worlds.

[12] Normally the second disjunct of this definition is put as: "there exists world $\omega_1 \in [\neg x]$ such that $\omega_1 \sqsubseteq \omega_2$ for every world $\omega_2 \in [\neg y]$". We make the definition slightly complicated (but equivalent) in order to simplify some of the proofs.

[13] That is not entirely true, since we have adapted $\preceq$ and $\sqsubseteq$ to deal with inconsistent evidence $\bot$. But that should not create any problem here.

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.24 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 24

5.1. Comparison with natural revision

Questions still linger as to how good the proposal we advocate is. We present below two examples that Darwiche and Pearl [7] use to illustrate that Boutilier's account of Natural Revision is too stringent. It turns out that our account of dynamic belief revision blocks both counterexamples.

Example 17 (*Darwiche and Pearl*, 1997). We encounter a strange new animal and it appears to be a bird, so we believe the animal is a bird. As it comes closer to our hiding place, we see clearly that the animal is red, so we believe that it is a red bird. To remove further doubt about the animal's birdness, we call a bird expert who takes it for examination and concludes that it is not a bird but some sort of mammal. The question now is whether we should still believe that the animal is red. (Intuitively we should, according to Natural Revision we should not.)

Assuming atoms *bird* and *red*, and the initial belief set $K = Cn(\{bird\})$ this example asks whether or not $red \in (K^*_{red})^{*|red}_{\neg bird}$. It immediately follows from (7*new) that $(K^*_{red})^{*|red}_{\neg bird} = K^*_{red\wedge\neg bird}$ whereby $red \in (K^*_{red})^{*|red}_{\neg bird}$. This example can be easily examined via entrenchment kinematics or SimpLex, leading to the same result.

Example 18 (*Darwiche and Pearl*, 1997). We face a murder trial with two main suspects, John and Mary. Initially it appears that the murder was committed by one person, so our initial belief set is $K = Cn(\{john \leftrightarrow \neg mary\})$. As the trial unfolds, however, we receive a very reliable testimony incriminating John, followed by another reliable testimony incriminating Mary. At this point, it is only reasonable to believe that both suspects took part in the murder, thus dismissing the one-person theory altogether .... (Natural Revision will dismiss the testimony against John, no matter how compelling.)

From the point of view of dynamic revision, this example is very similar to Example 17. We are trying to compute $(K^*_{john})^{*|john}_{mary}$. Since $john \wedge mary \nvdash \bot$, by a simple application of (7*new) again we get that $(K^*_{john})^{*|john}_{mary} = K^*_{john\wedge mary}$ whereby both John and Mary are believed to have participated in the murder. Again, this example can be examined via entrenchment kinematics and SimpLex, confirming the same result.

Thus we notice that our account proves to be quite robust against clear counterexamples. In contrast, although Darwiche and Pearl's account does not give the wrong result in these examples, it does not block these counterexamples either.

5.2. Revising inconsistent belief sets

Of the postulates introduced here, one of the more contentious is the postulate of Absurdity that deals with revision of the inconsistent belief set $K_\bot$. This postulate can be criticised as being too severe. In particular, it can be argued that adoption of this postulate goes against the Principle of Informational Economy advocated by the AGM and other approaches to belief revision. That is, "retain as much as possible of our old beliefs" [8,

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.25 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 25

p. 49]. However, under our approach both beliefs and the structure of $\sqsubseteq$ are lost once revision results in the absurd belief set $K_\perp$. The basis for our approach is a closer adherence to the standard AGM constructions.

If we would like to retain some of the structural information, a modification to our approach is to introduce an additional element $\omega_\perp$ besides the set of possible worlds Ω. The preorder $\sqsubseteq$ is now defined over $\Omega \cup \{\omega_\perp\}$. The world $\omega_\perp$ can be viewed as an "impossible" world and has the following properties:

(1) $\omega_\perp \notin [x]$ whenever $x \nvdash \perp$; and,
(2) $[x] = \{\omega_\perp\}$ whenever $\vdash \neg x$.

In other words, the impossible world $\omega_\perp$ is "consistent" only with logical falsehoods (and is the only world so consistent). Moreover, in our semantic construction $\sqsubseteq$ we ensure that $\omega_\perp$ is confined to its own rank and is the only element at this rank. That is,

$$\text{Either}\quad \omega_\perp \sqsubset \omega \quad\text{or}\quad \omega \sqsubset \omega_\perp \quad \text{for all } \omega \in \Omega.$$

If this condition is satisfied by an ordering $\sqsubseteq$ over $\Omega \cup \{\omega_\perp\}$, then it is easy to show that it is also satisfied by a revised preorder $\sqsubseteq_e^\circ$ (satisfying SimpLex1–SimpLex3).

In this scheme it is no longer required to represent the empty relation $\sqsubseteq_\perp$. The absurd belief set $K_\perp$ is associated with any relation $\sqsubseteq$ in which $\omega_\perp$ is the (only) $\sqsubseteq$-minimal element. Moreover Definition 14 is now simplified. With the understanding that $\omega_\perp \models \perp$, the latter part of the condition is all that is required.

A further requirement that one might consider is that $\omega_\perp$ be either the (sole) $\sqsubseteq$-minimal element when $K = K_\perp$ or be the (sole) $\sqsubseteq$-maximal element otherwise. This latter condition makes some intuitive sense as one can argue that $\omega_\perp$ should not be considered more plausible than any other possible world when K is consistent. However, in doing so one would pay the price by needing to complicate the conditions SimpLex1–SimpLex3 to stipulate the special handling of $\omega_\perp$. Without this further requirement conditions SimpLex1–SimpLex3 can be used as they stand.

Along analogous lines, an epistemic entrenchment relation can be also modified so that we can provide a more sophisticated account of escaping from inconsistency.

While going to these lengths would allow us to retain the structure inherent in $\sqsubseteq$, it does so at the expense of introducing the rather arbitrary construct $\omega_\perp$. In this paper we have chosen to remain faithful to the traditional AGM constructs and forsake such artificial considerations. In summary, each of these proposals has its strengths and weaknesses. We have chosen what we consider to be a reasonable compromise to deal with this boundary case.

5.3. Conclusion

In this paper we have argued for the need to enhance the traditional account of AGM belief revision with an account of dynamic belief change. While we are by no means the first to decry this deficiency in the AGM account, there are several novel aspects to our approach. Of the existing approaches to this problem, we have argued that the Darwiche–

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.26 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 26

Pearl account is too permissive. On the other hand, the approach offered by Boutilier does not deal well with some important cases.

We have introduced an approach to iterated belief change based on dynamic belief revision operators. Postulates for dynamic belief revision are motivated and supplied. These give quite intuitive constraints on the dynamics of belief revision operations. We have also supplied two constructive modellings for our account of dynamic belief revision. The first is based on the notion of epistemic entrenchment. The second is based on possible worlds and Simple Lexicography.

The nature of commonsense reasoning is such that we can never establish beyond doubt that our account will not fall prey to the next counterexample. However we can take reasonable care to see whether it blocks the known, potential counterexamples. To that extent, our proposal survives exceptionally well.

Acknowledgements

The authors would like to thank Norman Foo, Rex Kwok, Daniel Lehman, Hans Rott, Abdul Sattar, Mary-Anne Williams, Wayne Wobcke and two anonymous referees for *Artificial Intelligence* whose comments have greatly helped to improve this paper. The first author would also like to thank Fangzhen Lin for hosting him at the Hong Kong University of Science and Technology during the period in which a substantial part of this paper was written.

Appendix 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 $\prec^{\odot}_{\alpha}$ to denote the strict part of $\preceq^{\odot}_{\alpha}$. This should not be confused with the revision of $\prec$ by α! Similarly, for instance, $(\prec_*)^{\odot}_{\alpha}$ is used to denote the strict part of the relation $(\preceq_*)^{\odot}_{\alpha}$.

More importantly, $x\ (\not\preceq_*)^{\odot}_{\alpha}\ y$ will, for instance, abbreviate $\neg(x\ (\preceq_*)^{\odot}_{\alpha}\ y)$. The notation $x \not\preceq^{\odot}_{\alpha} y$ should be similarly interpreted.

This applies to semantic notation such as $\omega \not\sqsubseteq^{\circ}_{e} \omega'$ mutatis mutandis.

In the proofs of Theorem 9 and onwards, the revision operation $*$ (with possible decorations) is taken to be a unary operation. However for readability the traditional AGM notation is maintained throughout. In an expression of the form K^*_x, the belief set K should be interpreted as the contextually fixed belief set, namely $*(\top)$, and the set K^*_x should be interpreted as $*(x)$.

Theorem 3. *No belief revision operation $*$ that satisfies the AGM postulates satisfies #Recalcitrance.*

Proof. (In the text.)

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.27 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 27

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 4, $(K^*_{\neg y})^*_x = K^*_x = K$. 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 4 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 $x \wedge y \nvdash \bot$. 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 9. *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. Then* $(\preceq_*)^{\odot}_e = \preceq_{*|e}$.

Proof. Let $K = *(\top)$ as agreed, and $K' = *(e) = K^*_e$.

($\subseteq$) First we show that if $x\ (\preceq_*)^{\odot}_e\ y$ then $x \preceq_{*|e} y$. Assume that $x\ (\preceq_*)^{\odot}_e\ y$. Then, by Definition 6, either (a) $e \vdash \bot$ or (b) $e \vdash x$, $e \vdash y$ AND either $\vdash y$ or both $\nvdash x$ and $x \preceq_* y$ or (c) $e \nvdash x$ and $e \to x \ \preceq_* \ e \to y$.

Case (a). $e \vdash \bot$. Hence $K' = K^*_e = K_\bot$. Then, trivially, $x \preceq_{*|e} y$ by Definition 8.

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 8, we get $x \preceq_{*|e} y$. 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 8 it follows that $x \preceq_{*|e} y$.

Now consider the other relevant situation, i.e., both $\nvdash x$ and $x \preceq_* y$. (Also note the active assumptions that $e \vdash x$, $e \vdash y$.) By Definition 8 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 8 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 8 we get the desired result that $x \preceq_{*|e} y$. 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 8 the desired result follows.

Case (c). Assume that $e \nvdash x$ and $e \rightarrow x \preceq_* e \rightarrow y$. 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 8 to the assumption that $e \rightarrow x \preceq_* e \rightarrow y$, 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 $x \preceq_{*|e} y$ then $x (\preceq_*)^{\odot}_e y$. Assume that $x \preceq_{*|e} y$. Then, by Definition 8, 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 6 that $x\ (\preceq_*)^{\odot}_e\ y$.

Case 2. $K^*_e = K' = K_\bot$. By (Consistency) it follows that $e \vdash \bot$. So from part (a) of Definition 6 it follows that $x\ (\preceq_*)^{\odot}_e\ y$.

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 6).

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 8 it follows that $x \preceq_* y$. However, since $e \vdash x \wedge y$, the desired result easily follows from it with the help of Definition 6(b).

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 8 it then follows, among other things, that $e \wedge \neg(x \wedge y) \rightarrow x \preceq_* e \wedge \neg(x \wedge y) \rightarrow \neg x$. By substituting logical equivalents we get $(e \rightarrow x \wedge y) \vee x \preceq_* (e \rightarrow x \wedge y) \vee \neg x$. 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 $e \rightarrow x \wedge y \equiv_* (e \rightarrow x \wedge y) \vee x$. However, since $(e \rightarrow x \wedge y) \vee x$ is logically equivalent to $e \rightarrow x$, we get $e \rightarrow x \wedge y \equiv_* e \rightarrow x$. Using the EE postulates again (since $e \rightarrow x \wedge y$ is logically equivalent $(e \rightarrow x) \wedge (e \rightarrow y)$) we get $e \rightarrow x \preceq_* e \rightarrow y$.

To get the desired result with the help of Definition 6(c) 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, $p \preceq_* e \rightarrow x$ for all sentences p (by Dominance). Since $e \rightarrow x \preceq_* e \rightarrow y$, it follows by Transitivity and (EE4) that if $\bot \prec_* z$ for some z, then $\vdash e \rightarrow y$. But since (as we saw above) $\nvdash e \rightarrow y$, it follows that $\bot \not\prec_* z$ for every sentence z. Since $\preceq_*$ is connected (as follows from EE1–EE3), it follows that $z \preceq_* \bot$ for all z. Thus $\bot$ is $\preceq_*$-maximal. Hence, by Definition 5, 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 10. *Let $\preceq$ be an EE relation 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 $K = *(\top) = EC(\preceq)$.*

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).

From Definition 7 we get that (i) $y \in K^*_x$ iff either $x \rightarrow \neg y \prec x \rightarrow y$ or $x \vdash y$, and (ii) $y \in K^{*|e}_x$ iff either $x \rightarrow \neg y \prec^{\odot}_e x \rightarrow y$ or $x \vdash y$. From Theorem 1(a) of [19] 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 5 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 $x \vdash y$. I.e., $K^*_x = Cn(x)$. Thus, $*$ satisfies (0^*).

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.30 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 30

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 5 and EE1–EE4 that $a \preceq b$ for all a and b. Hence, it follows from Definition 7 that RHS = $Cn(\{e \wedge x\})$. On the other hand, from Definition 6 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 7 it follows that $y \in (K_e^*)_x^{*|e}$ iff either $x \to \neg y \prec_e^{\odot} x \to y$ 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 $x \vdash y$. Obviously, if $x \vdash y$, then we get the desired result, that $y \in Cn(e \wedge x)$ trivially. Hence we need only consider the other 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 6 (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 $x \nvdash y$. 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 5 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 6 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 7, we get either $x \to \neg y \prec_e^{\odot} x \to y$ or $x \vdash y$. The second case is 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 6 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 8. 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 8, 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 7 that $y \in (K_e^*)_x^{*|e}$ if and only if either $x \vdash y$ or $x \to \neg y \prec_e^{\odot} x \to y$. So we assume that $x \nvdash y$. It will be sufficient to show that $x \to y \npreceq_e^{\odot} x \to \neg y$. It follows from Definition 6 that we need to demonstrate the following three claims:

(A) $e \nvdash \bot$, (B) If $e \vdash \neg x$ then $x \to y \npreceq x \to \neg y$ and (C) If $e \nvdash x \to y$ then $e \to (x \to \neg y) \prec e \to (x \to y)$. 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 7 that either $e \wedge x \vdash y$ or $e \wedge x \to \neg y \prec e \wedge x \to y$. In the first case, the claim (C) follows trivially, since, if $e \wedge x \vdash y$, then by Deduction, $e \vdash x \to y$ (which contradicts the antecedent of (C)). Similarly, in the second case (C) follows trivially, since, $e \wedge x \to \neg y \prec e \wedge x \to y$ is logically equivalent to the consequent, $e \to (x \to \neg y) \prec e \to (x \to y)$, 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 7, $y \in \mathit{LHS}$ iff either $x \to \neg y \prec^{\odot}_e x \to y$ or $x \vdash y$. On the other hand, $y \in \mathit{RHS}$ iff either $x \to \neg y \prec x \to y$ or $x \vdash y$. Now, right to left is easy. Assume that $y \in \mathit{RHS}$, i.e., either $x \to \neg y \prec x \to y$ or $x \vdash y$. Since $K = K_\bot$, it follows from Definition 5 that $a \preceq b$ for all a, b whereby $x \to \neg y \nprec x \to y$. Hence, $x \vdash y$, whereby $y \in \mathit{LHS}$. Hence consider the proof for left to right. Assume that $y \in \mathit{LHS}$. Hence either $x \to \neg y \prec^{\odot}_e x \to y$ or $x \vdash y$. The second case is trivial. So consider the principal case: $x \to \neg y \prec^{\odot}_e x \to y$. I.e., $x \to y \npreceq^{\odot}_e x \to \neg y$. Applying Definition 6 we get, among other things, that either $e \nvdash x \to y$ or $e \nvdash x \to \neg y$ or both $\vdash x \to \neg y$ and either $\vdash x \to y$ or $x \to y \npreceq x \to \neg y$. (This particularly follows from Definition 6(b).) Now, by assumption $e \wedge x \vdash \bot$ (i.e., $e \vdash \neg x$) from which it follows both that $e \vdash x \to y$ and $e \vdash x \to \neg y$. Hence, both $\vdash x \to \neg y$ and either $\vdash x \to y$ or $x \to y \npreceq x \to \neg y$. However, since $K = K_\bot$ it follows that $x \to y \preceq x \to \neg y$. Hence both $\vdash x \to \neg y$ and $\vdash x \to y$ from which it follows that $\vdash \neg x$. From this it trivially follows that $y \in \mathit{RHS}$.

Case 2. Assume that $K \neq K_\bot$. Again, $y \in \mathit{LHS}$ iff either $x \to \neg y \prec^{\odot}_e x \to y$ or $x \vdash y$. From Definition 6, the connectedness of $\preceq$, and the fact that $K \neq K_\bot$, it follows that $x \to \neg y \prec^{\odot}_e x \to y$ iff (a) $e \nvdash \bot$, (b) if $e \vdash \neg x$ then $x \to \neg y \prec x \to y$ and (c) if $e \nvdash x \to y$ then $e \to (x \to \neg y) \prec e \to (x \to y)$. 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, $x \to \neg y \prec^{\odot}_e x \to y$ iff $x \to \neg y \prec x \to y$. Thus, $y \in \mathit{LHS}$ iff either $x \to \neg y \prec x \to y$ or $x \vdash y$ iff, by Definition 7, $y \in \mathit{RHS}$.

Theorem 11. *Let $*$ and $*|e$ be two revision operations satisfying the postulates* 0*–8*(*new*)*. Let $K = *(\top)$ and $K' = K^*_e$. Then the revision operations $*_{\preceq_*}$ and $*_{\preceq_{*|e}}$ are belief revision operations that satisfy* 0*–8*(*new*)*.*

Proof. Let $*$ and $*|e$ be two revision operations satisfying the postulates 0*–8*(new). Let $K = *(\top)$ and $K' = K^*_e$. It follows from Theorem 9 that $\preceq_{*|e} = (\preceq_*)^{\odot}_e$. Hence it would be sufficient to show that the operations $*_{\preceq_*}$ and $*_{(\preceq_*)^{\odot}_e}$ satisfy 0*–8*(new).

It is easily seen ([10]; note the extra clause, $K = K_\bot$ in Definition 7 to compensate for the difference between SEE and EE relations) that $\preceq_*$ is an EE relation. Furthermore, we know from Theorem 1(a) of [19] that, since $\preceq_*$ is an EE relation, the relation $(\preceq_*)^{\odot}_e$ is also an EE relation. Hence, surely $*_{\preceq_*}$ and $*_{(\preceq_*)^{\odot}_e}$ satisfy the basic AGM revision postulates 1*–6* (see [10]; note the extra clause, $x \vdash y$, in Definition 7 to compensate for the difference between SEE and EE relations). So we need only show that $*_{\preceq_*}$ and $*_{(\preceq_*)^{\odot}_e}$

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.32 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 32

satisfy 0*, 7*new and 8*new. 0*. (i) Suppose $K = K_\perp$. Then using Definitions 8 and 7 sequentially, and noting that $x \preceq_* y$ for all x and y (by Definition 8), it is easily seen that $y \in K_x^{*_{\preceq_*}}$ iff $x \vdash y$. Hence, $K_x^{*_{\preceq_*}} = Cn(x)$. Thus, the operation $*_{\preceq_*}$ satisfies 0*. (ii) By similar argument as above, it is easily seen that the operation $*_{(\preceq_*)_e^\odot}$ satisfied 0*. 7*new.

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

(Only if) Assume that $y \in (K_e^{*_{\preceq_*}})_x^{*_{(\preceq_*)_e^\odot}}$. It follows from Definition 7 then that either $x \to \neg y \ (\prec_*)_e^\odot \ x \to y$ or $x \vdash y$. The second case is easy: we know that $*_{\preceq_*}$ satisfies 1*–6*, in particular, Closure and Success; hence, surely $y \in K_{e\wedge x}^{*_{\preceq_*}}$. So, consider the principal case. Assume that $x \to y \ (\npreceq_*)_e^\odot \ x \to \neg y$. Hence, by Definition 6, (a) $e \nvdash \perp$, (b) If $e \vdash x \to y$ and $e \vdash x \to \neg y$, then $x \to y \npreceq_* x \to \neg y$, and (c) If $e \nvdash x \to y$ then $e \to (x \to y) \npreceq_* e \to (x \to \neg y)$. Now, by Definition 7, we need to show that either $e \wedge x \vdash y$ or $(e \wedge x) \to \neg y \prec_* (e \wedge x) \to y$. We assume $e \wedge x \nvdash y$. It will be sufficient to show that $(e \wedge x) \to y \npreceq_* (e \wedge x) \to \neg y$.

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

(If) Assume that $y \in K_{e\wedge x}^{*_{\preceq_*}}$. Hence, by Definition 7, either (i) $e \wedge x \vdash y$ or (ii) $(e \wedge x) \to \neg y \prec_* (e \wedge x) \to y$. We need to show that $y \in (K_e^{*_{\preceq_*}})_x^{*_{(\preceq_*)_e^\odot}}$. In light of Definition 7 we assume that $x \nvdash y$. It will be sufficient to show that $x \to y \ (\npreceq_*)_e^\odot \ x \to \neg y$. By Definition 6, it follows then that it will be sufficient to show that (a) $e \nvdash \perp$, (b) If both $e \vdash x \to y$ and $e \vdash x \to \neg y$, then $x \to y \npreceq_* x \to \neg y$, and (c) If $e \nvdash x \to y$ then $e \to (x \to y) \npreceq_* e \to (x \to \neg y)$. 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 $e \vdash x \to y$ and $e \vdash x \to \neg y$, whereby (b) is vacuously satisfied. So we need to show only (c).

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

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

8*(new).

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

(Only if) Assume that $y \in (K_e^{*_{\preceq_*}})_x^{*_{(\preceq_*)_e^\odot}}$. Hence, by Definition 7, either (i) $x \to \neg y \ (\prec_*)_e^\odot \ x \to y$ or (ii) $x \vdash y$. Cases (ii) is trivial; so we consider only the principal case (i). So assume that $x \to \neg y \ (\prec_*)_e^\odot \ x \to y$. It follows from Definition (6, particularly 6(b)) then, among other things, that if $e \vdash \neg x$ (i.e., if both $e \vdash x \to y$ and $e \vdash x \to \neg y$), then $x \to \neg y \prec_* x \to y$. It follows from Definition 7 then that $y \in K_x^{*_{\preceq_*}}$.

(If) Conversely, assume that $y \in K_x^{*_{\preceq_*}}$. Hence, by Definition 7, either (i) $x \to \neg y \prec_* x \to y$ or (ii) $x \vdash y$. Case (ii) is trivial. So we consider only case (i). Assume then that $x \to \neg y \prec_* x \to y$. Hence, vacuously, we get that if both $e \vdash x \to y$ and $e \vdash x \to \neg y$ then $x \to y \npreceq_* x \to \neg y$ (use part (b) of Definition 6). Furthermore, by assumption, $\nvdash \neg e$. Besides, since by assumption $e \vdash \neg x$, we vacuously get that if $e \nvdash x \to y$ then $e \to (x \to y) \npreceq_* e \to (x \to \neg y)$ (use part (c) of Definition 6). From Definition 6 we then get that $x \to \neg y (\prec_*)_e^\odot x \to y$. It follows from Definition 7 then that $y \in (K_e^{*_{\preceq_*}})_x^{*_{(\preceq_*)_e^\odot}}$.

Theorem 12. *Let $\preceq$ be an EE relation, and $\preceq_e^\odot$ the result of revising $\preceq$ by a sentence e. Then $\preceq_{*_\preceq} = \preceq$ and $\preceq_{*_{\preceq_e^\odot}} = \preceq_e^\odot$.*

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.33 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 33

Proof. (We omit the proof of this theorem. The proof can be easily constructed by slightly modifying the analogous proof in [10].)

Theorem 16. *Let $\preceq$ be an EE relation and $\sqsubseteq$ either an empty relation $\sqsubseteq_\bot$ or a total preorder over Ω. For every sentence e, the revised EE relation $\preceq_e^\odot$ and the revised preorder $\sqsubseteq_e^\circ$ are each other's counterpart, given that $\preceq$ and $\sqsubseteq$ are each other's counterpart.*

Proof. Assume that the EE relation $\preceq$ and the preorder $\sqsubseteq$ are each other's counterparts, and e is some arbitrary sentence. Let us first look at the two special cases.

Special case 1. Assume that $e \vdash \bot$, i.e., $[e] = \emptyset$. It follows from Definition 6 of Entrenchment Revision that $x \preceq_e^\odot y$ for all $x, y \in \mathcal{L}$. Hence $\preceq_e^\odot$ is the trivial relation $\preceq_\bot$. On the other hand, since $[e] = \emptyset$, it follows from SimpSpecial1 that $\sqsubseteq_e^\circ = \sqsubseteq_\bot$. We know from Definition 15 that these two resultant relations are counterparts of each other.

Special case 2. Assume that $e \nvdash \bot$ but $\preceq = \preceq_\bot$. Accordingly, we take its counterpart to be $\sqsubseteq_\bot$.

Since $e \nvdash \bot$ we know that $\preceq_e^\odot$ and $\sqsubseteq_e^\circ$ are not absurd. Hence, in order to show that $\preceq_e^\odot$ and $\sqsubseteq_e^\circ$ are counterparts of each other, we need to show that $x \preceq_e^\odot y$ iff either $[y] = \Omega$ or there exists ω_1 that is $\sqsubseteq_e^\circ$-minimal in $[\neg x]$ and for every ω_2 that is $\sqsubseteq_e^\circ$-minimal in $[\neg y]$, $\omega_1 \sqsubseteq_e^\circ \omega_2$.

(Left to Right) Assume that $[y] \neq \Omega$. From Definition 6 together with the assumption that $\preceq = \preceq_\bot$ we infer that the three cases, (a) through (c) below, are jointly exhaustive. We show for each case that either it leads to the desired conclusion, or is impossible.

Case (a). $e \nvdash x$. So $[e] \cap [\neg x] \neq \emptyset$. Let $\omega_1 \in [e] \cap [\neg x]$. It follows from SimpSpecial2 that $\omega_1 \sqsubseteq_e^\circ \omega$ for every world ω. Hence ω_1 is $\sqsubseteq_e^\circ$-minimal in $[\neg x]$ and $\omega_1 \sqsubseteq_e^\circ \omega_2$, for any ω_2 that is $\sqsubseteq_e^\circ$-minimal in $[\neg y]$.

Case (b). $e \vdash x$ and $\vdash y$. Impossible since by assumption $[y] \neq \Omega$.

Case (c). $e \vdash x$, $e \vdash y$ and $\nvdash x$. Since $\nvdash x$ it follows that $[\neg x] \neq \emptyset$. Since $[y] \neq \Omega$ it follows that $[\neg y] \neq \emptyset$. So both $[\neg x]$ and $[\neg y]$ have $\sqsubseteq_e^\circ$-minimal elements. Pick two such elements ω_1 and ω_2. Since $e \vdash x$ and $e \vdash y$ it follows that both $\omega_1, \omega_2 \in [\neg e]$. It follows from SimpSpecial2 then that $\omega_1 \sqsubseteq_e^\circ \omega_2$.

(Right to Left) Assume that case (a) and (b) above are false. (Hence, $e \vdash x$ and $\nvdash y$.) We need to show that case (c) above holds, i.e., $e \vdash x$, $e \vdash y$ and $\nvdash x$. We have to consider two cases:

Case (1). $[y] = \Omega$. Impossible since we assumed that $\nvdash y$.

Case (2). There exists ω_1 that is $\sqsubseteq_e^\circ$-minimal in $[\neg x]$ and if any ω_2 is $\sqsubseteq_e^\circ$-minimal in $[\neg y]$ then $\omega_1 \sqsubseteq_e^\circ \omega_2$. Denote some minimal element of $[\neg x]$ by ω_1 and let ω_2 be $\sqsubseteq_e^\circ$-minimal in $[\neg y]$. By assumption $e \vdash x$. Furthermore, since $[\neg x]$ has a $\sqsubseteq_e^\circ$-minimal element, clearly $[\neg x] \neq \emptyset$ whereby $\nvdash x$. So it will be sufficient to show that $e \vdash y$. Suppose to the contrary that $e \nvdash y$. Then $[e] \cap [\neg y] \neq \emptyset$ whereas $[e] \cap [\neg x] = \emptyset$. Clearly, then, $\omega_1 \in [e]$. On the other hand, follows from SimpSpecial2 that $\omega_2 \in [e]$ whereby $\omega_1 \not\sqsubseteq_e^\circ \omega_2$ contradicting our original hypothesis.

Principal case. Assume that $e \nvdash \bot$ and $\preceq \neq \preceq_\bot$. Accordingly we assume that $[e] \neq \emptyset$ and $\sqsubseteq$ is a total preorder over Ω. Furthermore, since $[e] \neq \emptyset$ it follows that $\sqsubseteq_e^\circ$ is not the empty relation $\sqsubseteq_\bot$. We need to show that $\preceq_e^\odot$ and $\sqsubseteq_e^\circ$ are each other's counterparts.

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) P.34 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 aij1970 by:violeta p. 34

(Left to Right) Assume that $x \preceq_e^{\odot} y$. Assume further that $[y] \neq \Omega$. We need to show that there exists ω_1 that is $\sqsubseteq_e^{\circ}$-minimal in $\neg x$ and if any ω_2 is $\sqsubseteq_e^{\circ}$-minimal in $\neg y$ then $\omega_1 \sqsubseteq_e^{\circ} \omega_2$. Since $e \nvdash \bot$ and $[y] \neq \Omega$, according to Definition 6, effectively, either: $e \vdash x, e \vdash y, \nvdash x$ and $x \preceq y$, or: $e \nvdash x$ and $e \to x \preceq e \to y$. We will consider these two cases separately.

(Case 1) $e \vdash x, e \vdash y$, $\nvdash x$ and $x \preceq y$. Since $\nvdash x$ it follows that $[\neg x] \neq \emptyset$. Since $[y] \neq \Omega$ it follows that $[\neg y] \neq \emptyset$. It follows then that both $[\neg x]$ and $[\neg y]$ have $\sqsubseteq_e^{\circ}$-minimal elements. Let ω and ω' respectively be such minimal elements in $[\neg x]$ and $[\neg y]$. Since $e \vdash x$ and $e \vdash y$ it follows that $[\neg x] \subseteq [\neg e]$ and $[\neg y] \subseteq [\neg e]$. It follows from SimpLex2 then that ω is a $\sqsubseteq$-minimal element of $[\neg x]$ and ω' is a $\sqsubseteq$-minimal element of $[\neg y]$. Since $x \preceq y$, and $\preceq$ and $\sqsubseteq$ are counterparts, it follows than that $\omega \sqsubseteq \omega'$, as desired.

(Case 2) $e \nvdash x$ and $e \to x \preceq e \to y$. Since $e \nvdash x$ it follows that $[e] \cap [\neg x] \neq \emptyset$. Clearly then both $[\neg x]$ and $[\neg y]$ are nonempty. Let ω_1 and ω_2 be, respectively, two $\sqsubseteq_e^{\circ}$-minimal members of $[\neg x]$ and $[\neg y]$. It will be sufficient to show that $\omega_1 \sqsubseteq_e^{\circ} \omega_2$. Since $e \to x \preceq e \to y$, it follows that either: $[e \to y] = \Omega$, i.e., $[e] \cap [\neg y] = \emptyset$, or: there exists a $\sqsubseteq$-minimal member ω of $[e] \cap [\neg x]$ such that $\omega \sqsubseteq \omega'$ for every $\sqsubseteq$-minimal member ω' of $[e] \cap [\neg y]$. We consider these cases separately.

- (Case 2.1) Assume that $[e \to y] = \Omega$, i.e., $[e] \cap [\neg y] = \emptyset$. Since $[e] \cap [\neg x] \neq \emptyset$ it follows from SimpLex3 that $\omega_1 \in [e]$. On the other hand, since $[e] \cap [\neg y] = \emptyset$ it follows that $\omega_2 \in [\neg e]$. From another application of SimpLex3 it follows that $\omega_1 \sqsubseteq_e^{\circ} \omega_2$.
- (Case 2.2) assume that there exists a $\sqsubseteq$-minimal member ω of $[e] \cap [\neg x]$ such that $\omega \sqsubseteq \omega'$ for every $\sqsubseteq$-minimal member ω' of $[e] \cap [\neg y]$. Now, either $\omega_2 \in [\neg e]$ or $\omega_2 \in [e]$. If $\omega_2 \in [\neg e]$ then by the same argument used in case 2.1, we get $\omega_1 \sqsubseteq_e^{\circ} \omega_2$. So without loss of generality assume that $\omega_2 \in [e]$. Now, by assumption, ω_1 is $\sqsubseteq_e^{\circ}$-minimal in $[\neg x]$. Since $\omega_1 \in [e]$ it follows that ω_1 is $\sqsubseteq_e^{\circ}$-minimal in $[e] \cap [\neg x]$. By SimpLex1, it follows that ω_1 is $\sqsubseteq$-minimal in $[e] \cap [\neg x]$. By our assumption then, $\omega_1 \sqsubseteq \omega'$ for any $\omega' \in [e] \cap [\neg y]$, including ω_2, as desired.

(Right to Left) There are two cases to consider. Either: $[y] = \Omega$ or: there exists a world ω_1 that is $\sqsubseteq_e^{\circ}$-minimal in $[\neg x]$ such that $\omega_1 \sqsubseteq_e^{\circ} \omega_2$ for any ω_2 that is $\sqsubseteq_e^{\circ}$-minimal in $[\neg y]$. We need to show that from either of these assumptions, it follows that $x \preceq_e^{\odot} y$.

(Case 1) Since $[y] = \Omega$, it follows that $\vdash y$ whereby $e \vdash y$ and $e \to x \preceq e \to y$. By taking the disjunctive cases: $e \vdash x$ and $e \nvdash x$ we get the disjunction of conditions (b) and (c) in Definition 6.

(Case 2) Without loss of generality assume that $[\neg y] \neq \emptyset$ and ω_2 is $\sqsubseteq_e^{\circ}$-minimal in it. Hence $\omega_1 \sqsubseteq_e^{\circ} \omega_2$. Now we consider two sub-cases: $e \vdash x$ and $e \nvdash x$.

- (Case 2.1) $e \vdash x$, i.e., $[e] \subseteq [x]$. Since $\omega_1 \in [\neg x]$ it follows that $\omega_1 \in [\neg e]$. Since $\omega_1 \sqsubseteq_e^{\circ} \omega_2$, it follows from SimpLex (particularly SimpLex2) that $\omega_2 \in [\neg e]$. Since $\omega_2 \sqsubseteq_e^{\circ} \omega$ for any arbitrary member $\omega \in [\neg y]$, it follows from SimpLex again that every member $\omega \in [\neg y]$ is in $[\neg e]$; thus, $[\neg y] \subseteq [\neg e]$ whereby $e \vdash y$. Since $\omega_1 \in [\neg x]$ of course $\nvdash x$. So, by part (b) of Definition 6, it would be sufficient to show that $x \preceq y$. Assume, without loss of generality that ω is $\sqsubseteq$-minimal in $[\neg x]$ and ω' is $\sqsubseteq$-minimal in $[\neg y]$. It will be sufficient then to show that $\omega \sqsubseteq \omega'$. Now, since $\omega_1 \in [\neg x]$ and

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.35 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 35

$\omega_2 \in [\neg y]$ it follows that $\omega \sqsubseteq \omega_1$ and $\omega' \sqsubseteq \omega_2$. On the other hand, since $[e] \subseteq [x]$ and $[e] \subseteq [y]$, it follows that $\omega, \omega' \in [\neg e]$. Clearly, then, $\omega_1 \sqsubseteq^\circ_e \omega$ and $\omega_2 \sqsubseteq^\circ_e \omega'$. From SimpLex2 it follows that $\omega_1 \sqsubseteq \omega$ and $\omega_2 \sqsubseteq \omega'$. Hence $\omega \equiv \omega_1 \sqsubseteq \omega_2 \equiv \omega'$ whereby $\omega \sqsubseteq \omega'$ as desired.

– (Case 2.2) $e \nvdash x$; thus $[e] \cap [\neg x] \neq \emptyset$. By part (b) of Definition 6 it will be sufficient to show that $e \to x \preceq e \to y$. If $e \vdash y$ then it is trivially shown; hence, without loss in generality assume that $e \nvdash y$, i.e., $[e] \cap [\neg y] \neq \emptyset$. Since $[e] \cap [\neg x] \neq \emptyset$, let ω be a $\sqsubseteq$-minimal element of $[e] \cap [\neg x]$. Similarly, let ω' be a $\sqsubseteq$-minimal element of $[e] \cap [\neg y]$. So it will be sufficient to show that $\omega \sqsubseteq \omega'$. Now, we know that ω_1 is $\sqsubseteq^\circ_e$-minimal in $\neg x$, ω_2 is $\sqsubseteq^\circ_e$-minimal in $\neg y$, and $\omega_1 \sqsubseteq^\circ_e \omega_2$. Since $[e] \cap [\neg x]$ and $[e] \cap [\neg y]$ are nonempty, it follows from SimpLex3 that ω_1 and ω_2 are in $[\neg e]$ as well. From SimpLex1 it then follows that ω_1 is $\sqsubseteq$-minimal in $[e] \cap [\neg x]$, and ω_2 is $\sqsubseteq$-minimal in $[e] \cap [\neg y]$. Hence, it follows that $\omega \equiv \omega_1$ and $\omega_2 \equiv \omega'$. Furthermore, since $\omega_1 \sqsubseteq^\circ_e \omega_2$, it follows from SimpLex that $\omega_1 \sqsubseteq \omega_2$, whereby $\omega \sqsubseteq \omega'$, as desired.

References

[1] C.E. Alchourrón, P. Gärdenfors, D. Makinson, On the logic of theory change: Partial meet contraction and revision functions, J. Symbolic Logic 50 (1985) 510–530.
[2] C. Areces, V. Becher, Iterable AGM functions, in: M.-A. Williams, H. Rott (Eds.), Frontiers in Belief Revision, Kluwer Academic, Dordrecht, 2001.
[3] R. Booth, The lexicographic closure as a revision process, J. Appl. Non-Classical Logics 11 (2001) 35–58.
[4] C. Boutilier, Revision sequences and nested conditionals, in: R. Bajcsy (Ed.), Proc. IJCAI-93, Chambéry, France, 1993, pp. 519–525.
[5] C. Boutilier, Iterated revision and minimal revision of conditional beliefs, J. Philos. Logic 25 (1996) 262–304.
[6] A. Darwiche, J. Pearl, On the logic of iterated belief revision, in: R. Fagin (Ed.), Proc. Fifth Conference on the Theoretical Aspects of Reasoning about Knowledge, 1994, pp. 5–23.
[7] A. Darwiche, J. Pearl, On the logic of iterated belief revision, Artificial Intelligence 89 (1997) 1–29.
[8] P. Gärdenfors, Knowledge in Flux. Modeling the Dynamics of Epistemic States, Bradford/MIT, Cambridge, MA, 1988.
[9] M. Freund, D. Lehmann, Belief revision and rational inference, Technical Report 94-16, Leibniz Center for Research in Computer Science, Institute of Computer Science, Hebrew University of Jerusalem.
[10] P. Gärdenfors, D. Makinson, Revision of knowledge systems using epistemic entrenchment, in: M. Vardi (Ed.), Proc. Second Conference on Theoretical Aspects of Reasoning about Knowledge, Morgan Kaufmann, Los Altos, CA, 1988, pp. 83–95.
[11] P. Gärdenfors, D. Makinson, Nonmonotonic inference based on expectation, Artificial Intelligence 65 (1994) 197–245.
[12] A. Grove, Two modellings for theory change, J. Philos. Logic 17 (1988) 157–170.
[13] S.O. Hansson, A Textbook of Belief Dynamics: Theory Change and Database Updating, Kluwer Academic, Dordrecht, 1999.
[14] H. Katsuno, A.O. Mendelzon, On the difference between updating a knowledge base and revising it, in: P. Gärdenfors (Ed.), Belief Revision, Cambridge University Press, Cambridge, 1992, pp. 183–203.
[15] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, Artificial Intelligence 44 (1990) 167–207.
[16] D. Lehmann, Belief revision, revised, in: Proc. IJCAI-95, Montreal, Quebec, 1995, pp. 1534–1540.
[17] A. Nayak, N. Foo, M. Pagnucco, A. Sattar, Changing conditional beliefs unconditionally, in: Y. Shoham (Ed.), Proc. Sixth Conference on Theoretical Aspects of Reasoning about Knowledge, 1996, pp. 119–135.
[18] A.C. Nayak, Iterated belief change based on epistemic entrenchment, Erkenntnis 41 (1994) 353–390.

S0004-3702(03)00017-1/FLA AID:1970 Vol.●●●(●●●) aij1970 P.36 (1-36)
ELSGMLTM(ARTINT):m1a 2003/02/07 Prn:11/02/2003; 11:44 by:violeta p. 36

[19] A. Nayak, N. Foo, M. Pagnucco, A. Sattar, Entrenchment kinematics 101, in: C. Zhang, J. Debenham, D. Lukose (Eds.), Proc. Seventh Australian Joint Conference on Artificial Intelligence, World Scientific, Singapore, 1994, pp. 157–164.
[20] O. Papini, Iterated revision operations stemming from the history of an agent's observations, in: M.-A. Williams, H. Rott (Eds.), Frontiers in Belief Revision, Kluwer Academic, Dordrecht, 2001.
[21] P. Peppas, M.-A. Williams, Constructive modellings for theory change, Notre Dame J. Formal Logic 36 (1) (1995) 120–133.
[22] A. Reńyi, Foundations of Probability, Holden-Day, 1970.
[23] H. Rott, Change, Choice and Inference. A Study of Belief Revision and Nonmonotonic Reasoning, Oxford University Press, Oxford, 2001.
[24] H. Rott, Coherence and conservatism in the dynamics of belief. Part I: Finding the right framework, Erkenntnis 50 (1999) 387–412.
[25] W. Spohn, Ordinal conditional functions: A dynamic theory of epistemic states, in: W.L. Harper, B. Skyrms (Eds.), Causation in Decision, Belief Change and Statistics, Vol. 2, 1988, pp. 105–134.
[26] M.-A. Williams, Transmutations of knowledge systems, in: J. Doyle, E. Sandewall (Eds.), Proc. Fourth International Conference on Principles of Knowledge Representation and Reasoning, 1994, pp. 619–629.