Abhaya C. Nayak†
Paul Nelson‡ and Hanan Polansky‡

# Belief Change as Change in Epistemic Entrenchment

**(Running head: Entrenchment Kinematics)**

**Abstract**

In this paper, it is argued that both the belief state and its input should be represented as epistemic entrenchment (EE) relations. A belief revision operation is constructed that updates a given EE relation to a new one in light of an evidential EE relation, and an axiomatic characterization of this operation is given. Unlike most belief revision operations, the one developed here can handle both "multiple belief revision" and "iterated belief revision".

## 1 Introduction

The ability to deal with new information, in particular the ability to incorporate new information into a belief state while preserving consistency, is a crucial component of intelligent behavior. In recent years this problem has received much attention from the artificial intelligence community. Consequently different researchers have put forward different proposals, with different applications in mind, including database dynamics [1], reason maintenance systems [6, 7], reasoning about actions [14, 42], belief change in a static world [9], belief change in a dynamic world [23] and nonmonotonic logics [5, 25], to mention only a few.

An approach to rational belief change that has influenced much of this discussion is the "AGM model" [2]. In the AGM model, a belief state (epistemic state) is represented as a belief set (i.e., as a set of sentences $K$ closed under the consequence operation of a suitable logic). With each such belief set is associated a selection mechanism, say, $\sigma$. The received information is represented as a sentence, $x$, possibly inconsistent with $K$. A belief revision operation, $*$, then yields a unique belief set $*(K, \sigma, x)$ that is "well behaved" in that it satisfies the Gärdenfors postulates (AGM postulates) of belief revision (see [2] or §3 of [9]).

A well known problem with the AGM approach is that it does not cope well with the concept of iterated belief change [17, 26, 39]. Since $*(K, \sigma, x)$ is again a belief set, if the agent receives information $y$ after receiving $x$, the agent requires a selection mechanism $\sigma'$ associated with the new belief set $*(K, \sigma, x)$ in order to perform the subsequent revision by $y$. Unfortunately, the selection mechanisms are so belief set specific that, in general, the original mechanism $\sigma$ is inappropriate for the new belief set $*(K, \sigma, x)$.

Another pressing problem for the AGM approach is multiple revision. It is recognized [30] that representing the new information as a single sentence is rather problematic, especially if the new information is inconsistent with the prior content of the belief state. If the new information is represented as a set of sentences $X = \{x_1, x_2, \ldots\}$, one may argue, revision of $K$ by $X$ involves the retraction of the set $\overline{X} = \{\neg x_1, \neg x_2, \ldots\}$. But, in general, retraction of $\overline{X}$ cannot be equated either with the retraction of $\neg x_1 \wedge \neg x_2 \wedge \ldots$ or with the retraction of $\neg x_1 \vee \neg x_2 \vee \ldots$.[1] Accordingly, revision by the set $X = \{x_1, x_2, \ldots\}$ should be distinguished from the revision by the sentence $x_1 \wedge x_2 \wedge \ldots$.

The account of belief change developed in this paper solves both of these problems in a single framework. The particular AGM selection mechanism chosen for this purpose is epistemic entrenchment. Epistemic entrenchment is a relation $\preceq$ over the sentences of the agent's language, essentially specifying which beliefs of the agent are easier to give up than others. Roughly, $x \preceq y$ means that if the agent were to give up at least one of $x$ or $y$, then it would be irrational on her part to give up $y$ while still retaining belief in $x$. Rott has argued [37] that $\preceq$ is the preference of an agent revealed through choice behavior over bits of information represented by different sentences of the language.

Rott has argued [33] that a relational measure like epistemic entrenchment on the belief set is not strong enough to account for iterated belief change, and that the selection mechanism needs to be strengthened to an ordinal measure, as suggested in [39]. Williams [41] has shown that Rott's suggestion works. On the other hand, we show that it is not *necessary* to impose an ordinal measure on the body of belief in order to solve the problem of iterated belief change. We show that if the new information is represented as an epistemic entrenchment relation, then both the problem of iterated belief change and the problem of multiple belief change can be solved simultaneously. Furthermore, since there are recipes already available as to how to generate default epistemic entrenchment relations given a set of sentences [35, 36], we can deal with iterated belief change even when the new information is represented simply as a sentence or a set of sentences. The approach taken in this paper is based on a framework developed in [26]. In contrast with the approaches to iterated belief revision of Alchourrón and Makinson [3, 4], Hansson [18, 20], Rott[34] and Schlechta [38] which are philosophically rather conservative, our approach is rather "liberal" (see [27] for details).

The crux of the thesis is as follows. We represent both the belief state and the new information as epistemic entrenchment relations ($\preceq_K$ and $\preceq_E$, respectively). We then develop a generic belief revision operation which revises $\preceq_K$ to a posterior epistemic entrenchment relation ($\preceq_{K*E}$). Since the result $\preceq_{K*E}$ itself is an epistemic entrenchment relation, it can in turn be revised by any further new information, thus solving the problem of iterated belief change (see §5). Furthermore, since the "epistemic content" of $\preceq_E$ can be seen as a set of sentences, it also simultaneously solves the problem of multiple revision (see §7). Apart from giving the construction of the operation in question, we also provide its axiomatic characterization by showing that there are three intuitive postulates which are satisfied by this, and only this, operation.

We must point out that there is significant overlap between this work and a previous work by one of the present authors ([27]). However there are important differences. Although both of these works present the identical entrenchment revision operation, the motivation offered in [27] is primarily semantic, and presupposes the works of Spohn and Grove. In the current work, however, an alternative motivation behind this construction is presented which is not semantic in character and does not presuppose the works of Spohn and Grove. The construction of the entrenchment revision in this work, unlike [27], does not require the background language or its associated logic to be complete either. Furthermore, unlike [27], the current work explains how multiple belief change can be achieved in this account. However, the primary contribution of the current work comprises a representation result which is conspicuously lacking in the earlier work.

Since the model of belief change developed in this paper is an extention of the AGM model, we begin with a review of this model. Next, the importance to belief dynamics modeling of expressing a belief state as a belief set and epistemic entrenchment relation pair rather than just a belief set is discussed. Furthermore, we show in a more general setting than that of the AGM model (similar to one suggested by Rott in [33] and used by him in [34]) that a belief set may be defined by the "epistemic content" of an epistemic entrenchment relation. Consequently, a belief state may be represented simply as an epistemic entrenchment relation. We then construct a unique belief revision operation which revises $\preceq_K$ by $\preceq_E$ and results in a well behaved posterior epistemic entrenchment relation $\preceq_{K*E}$. In many ways, our approach is similar to the account of preference change developed by Hansson [19]. We close with a short summary.

## 2 The AGM Model

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 $\vee$ (*not, if... then..., if and only if, and* and *or*). Two special symbols, $\top$ and $\bot$ are used to denote "truth" and "falsity", respectively. We will represent the set of all sentences of $\mathcal{L}$ by $\mathcal{L}$ itself, and its underlying logic by a consequence operation $Cn$: $2^{\mathcal{L}} \mapsto 2^{\mathcal{L}}$ satisfying the following conditions.

| | |
|---|---|
| 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 $\Gamma$ classically implies $x$ |
| Deduction: | $y \in Cn(\Gamma \cup \{x\})$ iff $(x \rightarrow y) \in Cn(\Gamma)$ |
| Compactness: | If $x \in Cn(\Gamma)$ then $x \in Cn(\Gamma')$ for some finite $\Gamma' \subseteq \Gamma$. |

The consequence relation $\vdash$ is used in the following sense:

$$\Gamma \vdash x \text{ iff } x \in Cn(\Gamma).$$

We often write "$\alpha \vdash \beta$" instead of "$\{\alpha\} \vdash \beta$" and "$\vdash \alpha$" instead of "$\emptyset \vdash \alpha$", for any sentences $\alpha$ and $\beta$. A belief set corresponds to any set $K \subseteq \mathcal{L}$ such that $Cn(K) = K$. By $K_\bot$ we refer to the absurd belief set $\mathcal{L}$ in which everything is believed.

In the AGM literature, a belief or epistemic state is represented by a belief set. In order to deal with belief dynamics, a belief set $K$ is assumed to be accompanied by a selection mechanism, namely an epistemic entrenchment relation $\preceq_K$, when $x \preceq_K y$ means that the agent (with the body of knowledge $K$) finds it at least as easy to give up her belief in $x$ as she does $y$. (We drop the subscript $_K$ from $\preceq_K$ where no confusion is likely to arise.) This relation $\preceq$ of (standard) epistemic entrenchment is required to satisfy the following conditions (where $x$ and $y$ are two arbitrary sentences of $\mathcal{L}$ and $K_\bot = \mathcal{L}$ is the absurd theory):[2]

| | | |
|---|---|---|
| $(SEE1)$ | If $x \preceq_K y$ and $y \preceq_K z$ then $x \preceq_K z$ | (transitivity) |
| $(SEE2)$ | If $x \vdash y$ then $x \preceq_K y$ | (dominance) |
| $(SEE3)$ | For any $x$ and $y$, $x \preceq_K x \wedge y$ or $y \preceq_K x \wedge y$ | (conjunctiveness) |
| $(SEE4)$ | When $K \neq K_\bot$, $x \notin K$ iff $x \preceq_K y$ for all $y$ | (minimality) |
| $(SEE5)$ | If $y \preceq_K x$ for all $y$, then $\vdash x$ | (maximality) |

Henceforth, we refer to any relation satisfying the above conditions (SEE1)-(SEE5) as a standard epistemic entrenchment (SEE) relation. We use $\prec$ for the strict part of $\preceq$ and $x \equiv y$ as an abbreviation for $x \preceq y \wedge y \preceq x$.

The main role of this selection mechanism $\preceq$ is in deciding which beliefs to give up from the original belief set $K$ before incorporating the new information. In the AGM tradition, a belief revision operation, $*$, is defined in terms of a belief contraction operation, $-$, via the Levi Identity:

- $K^*_x = Cn((K^{\perp}_{\neg x}) \cup \{x\})$,

where $K^{\perp}_{\neg x}$ is the belief set that results when $\neg x$, the information incompatible with $x$, is (possibly only vacuously) retracted from $K$. The primary function of $\preceq$ is in fact to identify this intermediate belief set $K^{\perp}_{\neg x}$ as determined by the following constructive definition:

- $(C)$: $\quad b \in K^{\perp}_a$ iff $b \in K$ and either $a \prec (a \vee b)$ or $\vdash a$

Gärdenfors and Makinson [11] have shown that if $\preceq$ satisfies the constraints (SEE1)–(SEE5), then the ensuing revision operation $*_{\preceq}$ (constructed from $\preceq$ by $(C)$ and the Levi Identity) satisfies all the Gärdenfors postulates of belief revision. Furthermore, it has also been shown that from every well behaved belief revision operation, $*$, a binary relation $\preceq_*$ can be constructed which satisfies the conditions (SEE1)-(SEE5). However, this revision operation based on $\preceq$ is not as versatile as intended. For any belief set $K$ and any proposition $x$, the expression $K^*_x$ is assumed to be well defined, and is taken to denote a belief set (see the Gärdenfors postulate of closure, [9], p. 54). Given a belief set $K$ and its associated SEE relation $\preceq_K$, it follows that $(K^*_x)^*_y$ is a well defined belief set. But this purported belief set cannot be computed until we are furnished with an SEE relation $\preceq_{K^*_x}$ associated with the intermediate belief set $K^*_x$.

Most approaches to iterated belief change lead to a form of epistemic determinism that we find philosophically unacceptable. Alchourrón and Makinson [3, 4], Hansson [18, 20],[3] Rott[34] and Schlechta [38] roughy suggest that in order to deal with the problem of iterated belief change, we let every agent be accompanied by a family $\mathcal{R}$ of SEE relations, which contains exactly one SEE relation $\preceq_K$ corresponding to each belief set $K$ expressible in the language $\mathcal{L}$. Thus, given that an agent's belief set is $K$ and she receives information $x$ and $y$ (in that order), the agent first uses the relation $\preceq_K$ corresponding to $K$ found in her $\mathcal{R}$ in order to compute $K^*_x$, and then uses the relation $\preceq_{K^*_x}$ corresponding to $K^*_x$ found in her $\mathcal{R}$ in order to compute $(K^*_x)^*_y$.[4] The family $\mathcal{R}$ of SEE relations may be dubbed the "conceptual framework" of the agent in question.[5] It is as if the agent is assigned a conceptual framework at birth, and cannot break away from it unless she is (un)lucky enough to undergo a conceptual revolution. In effect, there are SEE relations accessible to others, but not to our agent due to the limitations imposed by her conceptual framework. The approach developed in this paper is free of this limitation. Any and every epistemic entrenchment relation is accessible to any agent given the right epistemic input.

# 3 The Nature of the Epistemic State

So far we have adhered to the AGM tradition of describing a belief set $K$ as a belief (epistemic) state. In this section we argue that such a description is misleading, and that a belief state should properly be represented as an epistemic entrenchment relation.

While a belief set $K$ contains information as to which propositions are believed with certainty, it does not distinguish in any manner between the strength (firmness, entrenchment) of belief in one sentence versus another. Consequently, representation of a belief state as a belief set works well so far as the belief statics is concerned, but is inadequate to handle belief dynamics. Consider the following example by Hansson [16]: Both Al and Bill know that there are only two restaurants, X and Y, in their small village. On the evening of a public holiday they are both out looking for a hamburger. Both of them believe that $X$ is open ($x$) and that at least one of the restaurants X or Y is open ($x \vee y$). The reason why they believe that X is open is that they noticed that the lights in X are on. However, their belief that at least one of the restaurants X or Y is open is based on different reasons. Al believes $x \vee y$ because of his belief in $x$. Bill has an independent reason for believing $x \vee y$ since he, during his hamburger hunt, met a person eating a hamburger. Thus the belief sets of both Al and Bill, $K_a$ and $K_b$, are the same, namely $Cn(x)$.[6] Hence, statically speaking, they are in the same belief state. However it is incorrect to assume that the belief states of Al and Bill are dynamically equivalent. If Al and Bill were informed that the lights in X are on because the workers are cleaning the restaurant, Bill would still be justified in believing that at least one of the restaurants is open, but Al would not.

Gärdenfors [10] explains this dynamic difference by invoking different epistemic entrenchment relations $\preceq_a$ and $\preceq_b$ associated with the belief sets $K_a$ and $K_b$. If it is assumed that both $x \not\prec_a (x \vee y)$ and $x \prec_b (x \vee y)$, then it can be shown with the help of condition $(C)$ that although $K_a$ and $K_b$ are identical, $x \vee y$ is not in $(K_a)^{\perp}_{x}$ whereas $x \vee y$ is in $(K_b)^{\perp}_{x}$. Accordingly, although $x \vee y$ is not in $(K_a)^{*}_{\neg x}$, $x \vee y$ is in $(K_b)^{*}_{\neg x}$. (We defer a systematic account of this probelm to a later section, namely §7, where we discuss multiple belief change.)

The point of this example is that epistemic entrenchment should not be viewed as extraneous to a belief state. Since the revision of a belief set must involve the associated epistemic entrenchment, it is only appropriate to view epistemic entrenchment as information retained in the belief state itself. Thus, the belief set $K$ itself does not adequately represent the belief state in question. Rather, the belief state should be represented by the pair $\langle K, \preceq_K \rangle$. Trivial as this point may appear, it is the lack of attention to such trivial detail that has led to the unwarranted conclusion that the foundations theory of belief change is superior to the coherence theory [17, 10].[7]

The relationship between the belief state, belief set and epistemic entrenchment relation can be further simplified. The epistemic entrenchment relation $\preceq_K$, apart from having information about which members of $K$ are more entrenched than others, contains enough information for the set $K$ to be extracted from it. In fact, a partial definition of $K$ by $\preceq$ is given by (SEE4). A complete definition of $K$ in terms of $\preceq$ is:

$$K = \begin{cases} \{x : \exists y \ (y \prec x)\} & \text{if there are } x \text{ and } y \text{ such that } y \prec x \\ K_\perp & \text{otherwise} \end{cases}$$

Consequently, since a belief set $K$ can be extracted from its accompanying epistemic entrenchment relation $\preceq$, the belief state $\langle K, \preceq \rangle$ may be described solely by $\preceq$. Hans Rott [34] was the first to argue for this thesis.

This discussion brings to light a possible alteration to the definition of an epistemic entrenchment relation. Since a belief set may be defined in terms of an epistemic entrenchment relation, reference to the belief set in the definition of the epistemic entrenchment relation, as in (SEE4), is unnecessary.

A second possible alteration concerns (SEE5). This condition, that the maximally entrenched beliefs be logical truths, is unnecessary, and may be inappropriate. In particular, if an agent does not mind courting inconsistency, there is no reason to assume that such an agent's maximally entrenched beliefs are logical truths.

Following these observations about (SEE4) and (SEE5), henceforth, an EE relation is defined as any relation $\preceq$ such that for any arbitrary sentences $x$, $y$, $z$ of $\mathcal{L}$,

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

An EE relation is a generalized version of the SEE relation since condition SEE4 is dispensed with and SEE5 is weakened to EE4. Rott's GEE relation [34, 12] is another such relation. We refer to an EE relation $\preceq$ such that $x \preceq y$ for all $x, y$ as an *absurd* EE relation.

It is easily noticed that the EE conditions make heavy demands on the agent's knowledge of logic. For instance, EE2 implies that logically equivalent sentences are equally entrenched. Obviously it is too much to ask from agents (natural or artificial) at the performance level. However, we must realise that the AGM theory is a normative one, and the theory is, strictly speaking, on the epistemic (doxastic) commitments of the agent, not on the actual beliefs. Another interesting property of the EE relation that may fail at the performance level, but is well regarded in the literature is that the EE relation is connected, i.e. for every pair of sentences $x$ and $y$, either $x \preceq y$ or $y \preceq x$.

The "epistemic content" $EC(\preceq)$ of an EE relation $\preceq$ is defined as:

$$EC(\preceq) = \begin{cases} \{x : \bot \prec x\} & \text{if } \bot \prec x, \text{ for some } x \\ \{x : \bot \preceq x\} & \text{otherwise.} \end{cases}$$

It is easily verified that for any EE relation $\preceq$, its epistemic content $EC(\preceq)$ is closed under the $Cn$ operation. Furthermore, the epistemic content of an absurd EE relation is the absurd belief set $K_\bot$. Intuitively, the epistemic content of an EE relation is the belief set associated with that EE relation. Consequently, the epistemic content of an EE relation may be used as an index to distinguish it from other EE relations. In particular, the epistemic content of the prior belief state is the belief set $K$, and we denote the prior belief state by the EE relation $\preceq_K$. Note that this mode of indexing EE relations is ambiguous since two or more distinct EE relations may have identical epistemic content. However, the ambiguity accrued is amply compensated for by the notational simplicity.

Following this definition of $K$ and, hence, the belief state in terms of the epistemic entrenchment relation $\preceq$, the goal of a belief revision program may be stated as follows:

- A sufficient goal of a belief change program is to give a recipe of how to revise an epistemic entrenchment relation into a new one in light of the new information acquired.

# 4 Epistemic Input as an EE Relation

In the AGM tradition, the epistemic input to a belief state is represented as a single sentence. Many researchers ([8, 17, 26, 30]) have found this to be a limitation, and have suggested that the framework be generalized to at least accept sets of sentences as the epistemic input. Such generalizations of the AGM framework can be found in [18, 20].

This suggestion, strictly speaking, is not a deviation from the AGM framework. Since a belief set $K$ is closed under $Cn$ and any reasonable input $x$ is accepted in the revised belief set $K^*_x$, effectively, $Cn(x) \subseteq K^*_x$.[8] Hence, the input may as well be viewed to be $Cn(x)$ which is a set of sentences. In this sense, our approach effectively generalizes the AGM approach to belief change.

Rather than representing the epistemic input as a set of sentences, we represent it as an epistemic entrenchment relation. This is motivated by the fact that the epistemic input, although fully accepted, is also deemed corrigible by the agent. That the epistemic input is corrigible in the AGM framework is easily demonstrated. So long as $x$ is not a logical truth, according to the Gärdenfors postulates of belief revision, $x$ is retracted from $K^*_x$ if the agent acquires the new information $\neg x$. Hence, $x$ is corrigible in $K^*_x$. Consequently, rather than assume that $x$, as an epistemic input, is incorrigible, but (magically) acquires corrigibility once it is introduced to the belief set $K$, it makes a lot more sense to assume, from the outset, that $x$, as an epistemic input, is corrigible. This viewpoint then argues for the representation of the epistemic input as an epistemic entrenchment relation.

Representing epistemic input as an epistemic entrenchment relation is surely not a commision of philosophical heresy. An important moral to be drawn from Richard Jeffrey's work on probability kinematics [22] is that the belief state and the epistemic input should be represented as objects of the same category. Jeffrey represents them both as probability measures. Wolfgang Spohn [39] argues for this same thesis and represents both the belief state and the epistemic input as ordinal conditionalization functions. Against such a backdrop, representing epistemic input as an epistemic entrenchment relation is only to be expected.

# 5 The EE Revision Operation

We have argued that the belief state of an agent is best represented as an epistemic entrenchment relation. Similarly the new information (observation, evidence) acquired by the agent is best represented by another (possibly the same) epistemic entrenchment relation. Given these representations, our goal is to develop a generic procedure of revising $\preceq_K$ by $\preceq_E$ so that the result, $\preceq_K * \preceq_E$ is another reasonable epistemic entrenchment relation. (Henceforth, for the sake of readability, we will denote by $\preceq_{K*E}$ the proposed EE relation $\preceq_K * \preceq_E$.) The epistemic content of an EE relation is used as its index to distinguish it from other EE relations. In particular, the epistemic content of the prior EE relation $\preceq_K$ is the belief set $K$, the epistemic content of the evidential EE relation $\preceq_E$ is the evidence $E$ and the epistemic content of the posterior EE relation $\preceq_{K*E}$ is the revised belief set $K * E$.

In this section, first we discuss three plausible postulates that the operation we are looking for should ideally satisfy. Then we construct an operation that satisfies these conditions. We furthermore show that no other operation satisfies these postulates.

## 5.1 Entrenchment Revision: Postulates

We desire a belief revision operation which generates a posterior belief state which is an EE relation. We call this EE revision postulate $(EE1^*)$:

$(EE1^*)$ $\preceq_{K*E}$ is an EE relation.

This closure condition which requires that the result of revising an EE relation by another must, in turn, be an EE relation is necessary in order to handle repeated belief revision using the same revision operation. The reasonableness of this requirement has been recognized before [33]. We call any operation satisfying $(EE1^*)$ an entrenchment revision operation.

However $(EE1^*)$ itself is not strong enough to guarantee a well behaved revision operation. For instance, in the revision operation,

**Def-Null:** $x \preceq_{K*E} y$ iff $x \preceq_E y$

condition $(EE1^*)$ is satisfied. However, an operation so defined is pretty badly behaved since it implies that in internalizing the acquired information, the agent loses all her old knowledge. We desire the revision operation to be well behaved in that the posterior EE relation generated is reasonable relative to the prior and evidential EE relations from which it arises. A reasonable supplement to $(EE1^*)$ is the following postulate:

$(EE2^*)$ If $x \prec_E y$ then $x \prec_{K*E} y$.

$(EE2^*)$ is based on the idea that the evidence takes priority over the prior beliefs. In the AGM approach, a prior belief (sentence) is sacrificed in order to accommodate any evidence that conflicts with it. Similarly, the prior ordering among beliefs is sacrificed in deference to the ordering suggested by the evidence.

However, both $(EE1^*)$ and $(EE2^*)$ are not jointly strong enough to ensure a well behaved revision operation since **Def-Null** satisfies both of these constraints. Another constraint is needed to deal with the case where $x \equiv_E y$. As a first approximation, we might suggest:

$(EEA)$ If $x \equiv_E y$, then $x \preceq_{K*E} y$ iff $x \preceq_K y$.

This condition looks reasonable at the first blush, but is not desirable, since it is inconsistent with $(EE1^*)$ and $(EE2^*)$ together. This is shown in the following discussion. If we have $x \equiv_K (x \wedge y) \prec_K y$ and $x \succ_E (x \wedge y) \equiv_E y$ (perfectly plausible assumptions), (EEA) results in $x \succ_{K*E} (x \wedge y) \prec_{K*E} y$. However, this means $\preceq_{K*E}$ is not an EE relation since it violates conjunctiveness $(EE3)$.

Thus (EEA) is too strong for our purpose and we need to weaken it. We take as a working hypothesis that when $x \equiv_E y$ under some special circumstances, but *not always*, the prior relation between $x$ and $y$ is preserved in the posterior EE relation. Our goal then is to identify these circumstances.

Let us introduce the following terminology:

**Definition 1** *Let $\preceq_K$ and $\preceq_E$ be the prior and the evidential EE relations respectively. A sentence $x$ is said to be "KE-special" if for all $y \succ_E x$ and for all $z$, it holds that $z \preceq_K y$ iff $z \preceq_E y$.*

This definition picks out a set of sentences $x$ which are special relative to the prior and evidential entrenchment relations $\preceq_K$ and $\preceq_E$ in the following sense: every sentence which is evidentially more entrenched than $x$ is such that its relative entrenchment in the prior and the evidence are same.

Nayak [26, 27] assumes a weakening of (EEA) which may be expressed as follows:

$(EEB)$ If $x \equiv_E y$ and both $x$ and $y$ are KE-special,
then $x \preceq_{K*E} y$ iff $x \preceq_K y$.

He constructs an EE revision operation that satisfies the three conditions ($(EE1^*)$, $(EE2^*)$ and (EEB)). He shows that the operation has many nice properties, but he fails to offer a representation theorem. Our diagnosis is that (EEB) is an over-weakened version of (EEA).

The proposal we offer, namely (EEC), sets the weakening at the right level.

**Definition 2** *Let $\preceq_K$ and $\preceq_E$ be the prior and the evidential EE relations respectively. A sentence $x$ is said to be "KE-special up Cn" if for all $y \in Cn(x)$ such that $y \succ_E x$ and for all $z$, it holds that $z \preceq_K y$ iff $z \preceq_E y$.*

This definition is more liberal than Definition 1 in the sense that the set of sentences that are KE-special is a subset of the set of sentences that are KE-special up $Cn$. A sentence $x$ is KE-special up $Cn$ as long as every consequence of $x$ which is evidentially more entrenched than $x$ is such that its relative entrenchment in the prior and the evidence are same.

$(EEC)$ If both $x \equiv_E y$ and $x \wedge y$ is KE-special up $Cn$,
then $x \preceq_{K*E} y$ iff $x \preceq_K y$.

One would expect that (EE1$^*$), (EE2$^*$) and (EEC) would not jointly offer a complete characterization of entrenchment revision, since they do not explicitly tell us what happens when $x \equiv_E y$ but $x \wedge y$ is not KE-special up Cn. However, perhaps surprisingly, as theorems (1-3) show in §6, the characterization is complete. We suggest that in presence of (EE2$^*$) and (EEC), the condition (EE1$^*$) determines how $x$ and $y$ will be related by $\preceq_{K*E}$ when $x \equiv_E y$ but $x \wedge y$ is not KE-special up Cn.

Given that both $\preceq_K$ and $\preceq_E$ are EE relations, (EEC) is equivalently expressed as ($EE3^*$), which we adopt as the official formulation of this postulate:

$$(EE3^*) \qquad \text{If } x \equiv_E y \text{ and } \forall_\alpha, \forall_{\beta \in Cn(x,y):\beta \succ_E x}(\alpha \preceq_E \beta \text{ iff } \alpha \preceq_K \beta), \\ \text{then } x \preceq_{K*E} y \text{ iff } x \preceq_K y.$$

Admittedly it would be very nice if we can equivalently replace ($EE3^*$) by a set of simple postulates. It is not obvious what set of postulates will do the trick, nor have we explored this in any detail. However, we believe the equivalence between (EEC) and ($EE3^*$) makes the meaning of the latter quite clear.

For ease of reference, we call any EE revision operation that satisfies the three postulates ($EE1^*$)-($EE3^*$) a "well behaved entrenchment revision operation". A well behaved entrenchment revision exhibits very nice properties. We state without proof that every well behaved entrenchment revision operation, more precisely a multiple belief revision operation $*$ generated by a well behaved entrenchment revision operation, satisfies the following intuitive generalizations (($G1^*$)-($G8^*$)) of the Gärdenfors postulates of belief revision.[9]

($G1^*$) $K * E$ is a theory
($G2^*$) $E \subseteq K * E$
($G3^*$) $K * E \subseteq Cn(K \cup E)$
($G4^*$) If $(K \cup E) \nvdash \perp$ then $Cn(K \cup E) \subseteq K * E$
($G5^*$) $K^*E = K_\perp$ iff $E \vdash \perp$
($G6^*$) If $Cn(E_1) = Cn(E_2)$ then $K * E_1 = K * E_2$
($G7^*$) $K * (E_1 \cup E_2) \subseteq Cn((K * E_1) \cup E_2)$
($G8^*$) If $(K * E_1) \cup E_2 \nvdash \perp$ then $Cn((K * E_1) \cup E_2) \subseteq K * (E_1 \cup E_2)$

($G1^*$) states that the epistemic content of the resulting EE relation must be a belief set. ($G2^*$) says that the epistemic content of the evidential EE relation must be accepted after the revision. This is based on the AGM assumption that the evidence takes precedence over the prior beliefs. ($G3^*$) states that the epistemic content of the posterior EE relation cannot contain any information that is not in the prior knowledge and the evidence together. ($G4^*$) maintains that no information should be lost in the process of revision unless the epistemic content of the evidential EE relation conflicts with that of the prior EE relation. $(G5^*)$ notes that the only case in which the epistemic content of the posterior EE relation is inconsistent is when the evidence itself is contradictory. $(G6^*)$ says that so long as the epistemic contents of two evidential EE relations are the same, the epistemic contents of the corresponding posterior EE relations remain the same. The epistemic ordering of the evidential beliefs only affects the ordering of the posterior beliefs, without affecting their contents in any manner. The last two conditions, in a sense, generalize the idea behind $(G3^*)$ and $(G4^*)$. If we replace the set $E_1$ by $\emptyset$ and $E_2$ by $E$ in $(G7^*)$ and $(G8^*)$, then we get, respectively, $(G3^*)$ and $(G4^*)$.

Does a well behaved entrenchment revision operation existΓ One way to answer this important question is to in fact construct an operation which can be demonstrated to be well behaved. We construct one such operation in the next section. Furthermore, we show that no other entrenchment revision operation is well behaved. In effect, we prove a representation theorem.

## 5.2 Construction

From here onwards, we assume that $\mathcal{L}$ is a finitary language, i.e., a language generated from a finite number of atomic sentences. (A truth functional language with $n$ atomic sentences has $2^{2^n}$ truth functions. So, although a finitary language has an infite number of sentences *per se*, the language has a finite mumber of sentences modulo logical equivalence.)

For any sentence $x$ of $\mathcal{L}$, denote by $\mathbf{x}$ the set of sentences $\{x' \mid x' \in \mathcal{L} \text{ and } Cn(x) = Cn(x')\}$. That is, $\mathbf{x}$ is the set of sentences in $\mathcal{L}$ that are logically equivalent to $x$. By $\mathbf{L}$ we denote the family of equivalence classes in $\mathcal{L}$. We will assume that there is a choice function $\mathcal{F} : \mathbf{L} \to \mathcal{L}$ that chooses, from an equivalence class a unique member of that class. The existence of such a function is guaranteed by the axiom of choice.

For any sentence $x$ of $\mathcal{L}$, denote by $|x|$ the set $\{x' : \ x \prec x'\}$ if this set is not empty. Otherwise, denote by $|x|$ the set $\{x' : \ x \preceq x'\}$. We refer to $|x|$ as the $x$-cut induced by $\preceq$. Furthermore, let $x^\#$ denote $\bigwedge\{\mathcal{F}(\mathbf{y}) \mid y \in |x|\}$.[10]

Notice that since the language $\mathcal{L}$ is finitary, even if the set $|x|$ is infinite, it breaks up to only a finite number of equivalence classes; so the conjunction $\bigwedge\{\mathcal{F}(\mathbf{y}) \mid y \in |x|\}$ is well-formed. Thus, for instance, if $\preceq_K$ is an EE relation defined over the members of $\mathcal{L}$,

- $|\perp|_K$ denotes the epistemic content $K$ of $\preceq_K$
- $\perp_K^\#$ denotes the content of $K$ as expressed by a single sentence.

These definitions generate four interesting observations. We will present the proof of Observation 1. Observations 2 and 4 are rather trivial; a proof of Observation 3 can be found in [27].

**Observation 1** *For any non-absurd EE relation* $\preceq$ *and any contingent sentence* $x$, $x \prec x^{\#}$

**Proof:** By definition, $x^{\#}$ is the sentence $\bigwedge\{\mathcal{F}(\mathbf{y}) \mid y \in |x|\}$. Hence $x^{\#} \equiv \bigwedge\{\mathcal{F}(\mathbf{y}) \mid y \in |x|\}$. It can be shown, with the help of $(EE3)$ and $(EE2)$ that there is some member $s \in |x|$ such that $s \preceq \bigwedge\{\mathcal{F}(\mathbf{y}) \mid y \in |x|\}$. By definition, $x \prec s$. Hence, $x \prec x^{\#}$. ■

**Observation 2** *For any EE relation* $\preceq$ *and any sentence* $x$, $|x|$ *is a belief set*

**Observation 3** *For any EE relation* $\preceq$ *and any sentence* $x$, $x \equiv (x^{\#} \rightarrow x)$

**Observation 4** $EC(\preceq) = |\perp|$ *for any EE relation* $\preceq$.

The posterior EE relation $\preceq_{K*E}$, the result of revising the prior $\preceq_K$ by the evidential $\preceq_E$, is defined as follows:

**Definition 3** *For any sentences* $x$ *and* $y$ *of* $\mathcal{L}$,

$$x \preceq_{K*E} y \text{ iff} \begin{cases} \text{either} & x \prec_E y \\ \text{or} & x \equiv_E y \text{ and } (x_E^{\#} \rightarrow x) \preceq_K (x_E^{\#} \rightarrow y) \end{cases}$$

The term $x^{\#} \rightarrow x$ plays a crucial role in the construction of the posterior EE relation. In §5.3, it is shown that terms of the type $x^{\#} \rightarrow x$ actually refer to a type of fixed point whose ordering in the prior is used to determine, to an important extent, the posterior ordering. Consequently we refer to any entrenchment revision operation constructed by Definition 3 as an "FPO (fixed point ordering) revision operation".

Before going to the motivation behind this construction, we clarify a subtle point. In Definition 3 we refer to the inequality $(x_E^{\#} \rightarrow x) \preceq_K (x_E^{\#} \rightarrow y)$. One might suspect an epistemic regress here, since, it appears as though the agent needs to know in advance what the evidence $\preceq_E$ is going to be so that she can possibly order $x_E^{\#} \rightarrow x$ and $x_E^{\#} \rightarrow y$ when she is in the prior belief state $\preceq_K$. But there is really no such regress. It is worth noting that both $\preceq_K$ and $\preceq_E$, being EE relations, are defined over the whole language $\mathcal{L}$. So, given any two sentences, say $x$ and $x'$, (1) either $x \preceq_K x'$ or $x' \preceq_K x$, and (2) either $x \preceq_E x'$ or $x' \preceq_E x$. (Of course, we would expect $\preceq_K$ to be more fine-grained that $\preceq_E$.) Hence, there should be no a priori objection to the inequality $(x_E^{\#} \rightarrow x) \preceq_K (x_E^{\#} \rightarrow y)$ in Definition 3.

The construction of the FPO revision operation is motivated by the following considerations:

(1) The evidential ordering is more reliable than the prior ordering. Hence if $x$ is strictly less entrenched than $y$ according to the evidence, than $x$ should also be strictly less entrenched than $y$ in the posterior.

(2) When the evidential ordering is indifferent between $x$ and $y$, it is interpreted as meaning that the evidence is not fine grained enough to discriminate between $x$ and $y$. In such a case, instead of $x$ and $y$ being equally entrenched in the posterior, relevant information from the prior is used to determine how $x$ and $y$ should be related in the posterior. When the posterior status of $x$ and $y$ are thus being determined, it may be assumed that the fate of all sentences that are strictly more evidentially entrenched than either of them, namely the members of $|x|_E$, is already determined, and they may be considered "known". Thus, $|x|_E$ constitutes the standard of judging $x$ and $y$.

This intuition is captured by adopting the following principle: $x$ and $y$ will be related in the posterior as are related in the prior the sentences $f(x)$ and $f(y)$, where $f(x)$ (respectively $f(y)$) encodes the information content of $x$ relative to $|x|_E$. Since a sentence $x \rightarrow y$ may be understood to encode the relative information available in the sentence $y$ given the known information $x$ [40], this intuition is captured by assuming that $x^{\#} \rightarrow x$ encodes the relative information in $x$ given that all members of $|x|_E$ are considered known. Similarly, $x^{\#} \rightarrow y$ encodes the relative information in $y$ given that all sentences in $|x|_E$ are considered known. This is the intuitive motivation behind Definition 3. Now we offer some further motivation.

## 5.3 Further Motivation: Fixed Point Ordering

In this section we offer further motivation behind the construction of entrenchment revision as defined by Definition 3. Unlike the motivation offered in [27] which is semantic in character and is heavily dependent on Spohn's account of ordinal conditionalization functions [39] and Grove's system of spheres [15], the motivation we offer is based on syntactic considerations. Its understanding does not presuppose familiarity with the work of Spohn or Grove.

We utilize a pictorial representation of an EE relation, suggested by Quine's account of web-of-belief [31, 32]. Imagine the sentences of $\mathcal{L}$ to be a sphere in which the more entrenched sentences are more centrally located than are the less entrenched sentences. Since $\mathcal{L}$ is finitary, a relation $\preceq$ induces a finite number of "cuts" in this language. According to Observation 2, each of these cuts is a belief set. Moreover, it is easily seen that in general, if $x \prec y$, then the $y$-cut is a proper subset of the $x$-cut. Another important feature of this system is that only the $\perp$-cut, namely $\mathcal{L}$

itself, contains contradictory sentences. All other cuts are consistent. Accordingly we can view $\preceq$ as generating a system of spheres (SOS) in the space of sentences. Such systems of spheres are similar to the systems of spheres constructed by Grove [15]. The main differences are that (1) while Grove's SOS's are defined over a space of worlds our SOS's are defined over the space $\mathcal{L}$ itself and (2) whereas in a Grovian SOS the points (worlds) can lie anywhere, in our SOS, the points (sentences) must satisfy certain constraints with regard to their location.

Figure 1 shows an EE relation, $\preceq$, as a system of spheres. For convenience, the

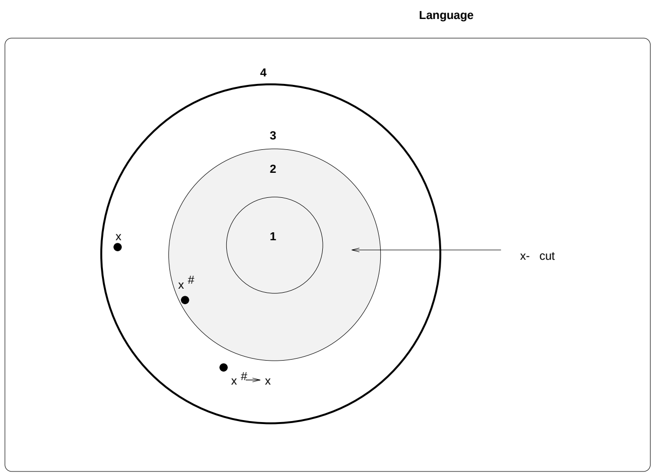


The closer a point (sentence) is to the center, the more entrenched it is understood to be. Falsity and other minimally entrenched sentences are outside the bold line. Points in the area bounded by the bold line represent the beliefs. The maximally entrenched sentences (Truth) are not shown in this figure.

Figure 1: $\preceq$ as a system of spheres

line representing $\mathcal{L}$ is shown as a rectangle instead of a circle. The numerals name the region of the space bounded by the lines around it. In this representation of $\preceq$, the spheres (cuts) in descending order of entrenchment (of sentences generating them) are 1, $1 \cup 2$, $1 \cup 2 \cup 3$ and $1 \cup 2 \cup 3 \cup 4$, respectively. Note that the most entrenched cut, namely $\top$-cut, is the null sphere $Cn(\emptyset)$. Hence, we conveniently avoid explicitly representing it in the graphical representation. Note further that given a sentence $x$ in the band $i$, for $i > 1$, the $x$-cut can be represented as the sphere $1 \cup 2 \cup \ldots \cup (i \perp 1)$. If $x$ is in band 1 or the null-sphere, then the $x$-cut is of course the null sphere. In Figure 1, $x$ is shown to be in band 3. Accordingly the $x$-cut is $1 \cup 2$. Furthermore, according to Observation 1, $x^{\#}$ is in $1 \cup 2$. In fact, it can be shown with the help of dominance and conjunctiveness that $x^{\#}$ is in band 2 as shown in the figure. Observation 3 implies that $x^{\#} \rightarrow x$ is in band 3 itself as is $x$. Observation 4 implies that the epistemic content of the EE relation depicted in this figure is the area enclosed by the bold line.

In Figure 2, we show both $\preceq_K$ and $\preceq_E$. The former is the system of spheres drawn in solid lines centered on 1, and the latter is drawn in broken lines centered on A. This figure makes explicit the reason for not graphically representing the null sphere. For

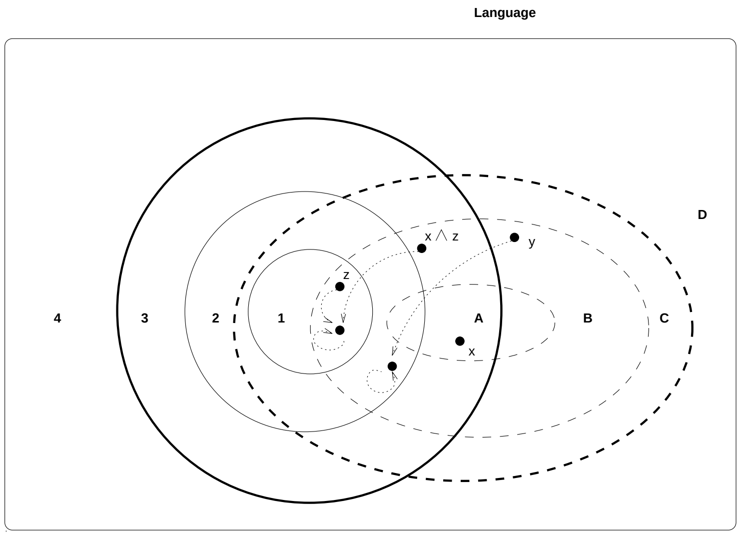


Figure 2: $\preceq_K$ and $\preceq_E$ as two SOS's

obvious reasons, the $\top$-cuts, namely $|\top|_K$ and $|\top|_E$, would be properly inside 1 and A respectively. Since they are disjoint in this picture, we cannot pictorially represent them. However, theoretically there is no problem since a null-sphere is a sphere.

Let us consider how to update $\preceq_K$ in light of $\preceq_E$. According to $(EE2^*)$, since sentences in A are evidentially (strictly) more entrenched than sentences in B, they remain so in the posterior. However, members of A can be discriminated according to their prior entrenchment. This is shown in Figure 2 by its partitioning into three disjoint subsets. It is useful to have names to denote such disjoint sectors. We do this in Figure 3. According to the prior entrenchment ordering, members of A2 are more

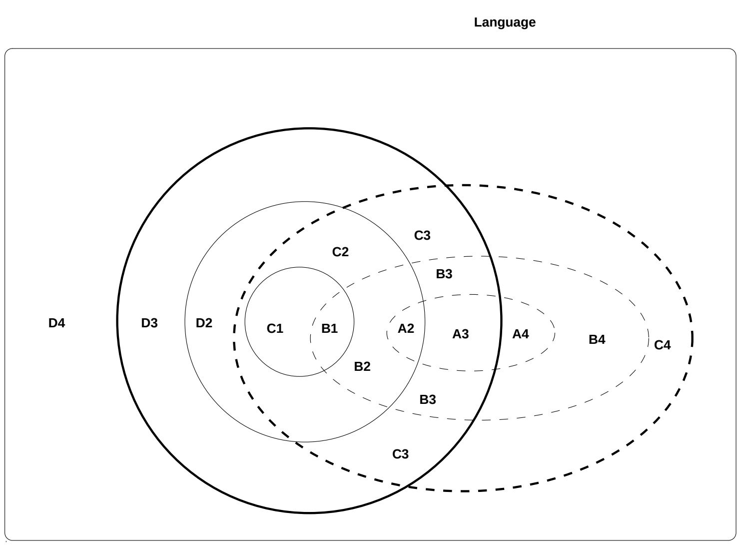


Figure 3: A partition generated by $\preceq_K$ and $\preceq_E$

entrenched than members of A3, and the latter in turn are more entrenched than the members of A4. Thus we expect the posterior system of spheres to be centered on A2 which would be immediately enveloped by A3, and A3 in turn enveloped by A4.

Since Observation 2 more or less characterizes an EE relation, it is easily verified that this approach is without any trouble so far. Observation 2 requires that A2, A2∪A3 and A2∪A3∪A4 be belief sets. And they indeed are, since (a) A, $1 \cup 2$ and $1 \cup 2 \cup 3$ are belief sets, (b) the intersection of any two belief sets is again a belief set, and (c) A2 = A∩$(1 \cup 2)$, A2∪A3 = A∩$(1 \cup 2 \cup 3)$ and A2∪A3∪A4 = A.

This success is tempting. Naturally the next step is to assume that A4 will be enveloped by B1 which in turn will be enveloped by B2, and so on, as pictured in Figure 4. Unfortunately, this does not work. For one thing, Observation 2 requires

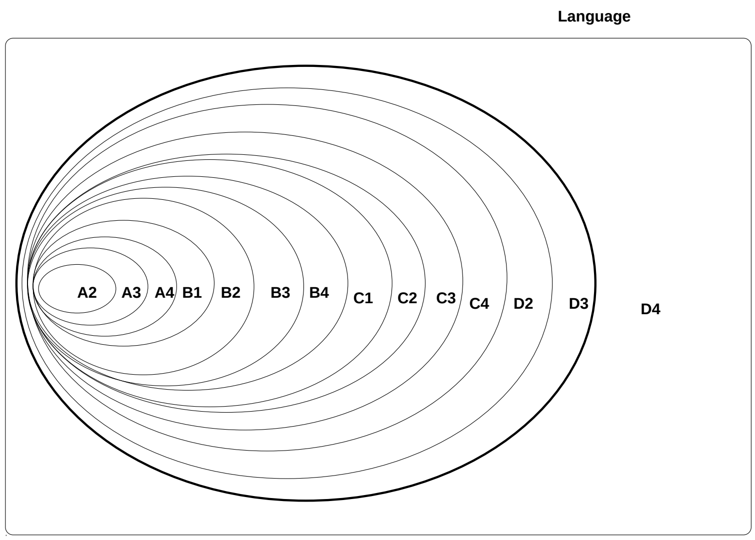


Figure 4: This is **not** the desired posterior SOS

that if A4 is to be enveloped by B1, then A2 ∪ A3 ∪ A4 ∪ B1 (= A ∪ B1) must be a belief set. It is true that A is a belief set, and that B1 too, being the intersection of the two belief sets A∪B and 1, is a belief set. However, the union of two belief sets is not necessarily a belief set. The other reason why the suggested procedure may fail is that B1 might contain elements that are inconsistent with some members of A. But it is required by an EE relation that any sentence inconsistent with a sentence in some inner band must be dumped in the outermost band!

The issue of A∪B1 not being a belief set can be analyzed pictorially by referring again to Figure 2. We first define a function $f_e : \mathcal{L} \mapsto \mathcal{L}$ as follows:

$$f_e(a) = a^{\#}_{E} \rightarrow a \text{ for every sentence } a.$$

We normally drop the index from $f_e(.)$ when it is not confusing. In Figure 2, the fact that $f(a) = b$ is shown by joining $a$ to $b$ by a dotted arrow pointing at the latter. It is easily verified that

- (1) $a$ and $f(a)$ are in the same band with respect to the evidential SOS. As pointed out earlier, this is due to Observation 3.
- (2) The $f$-arrows cannot be "centrifugal" with respect to *any* SOS in the space $\mathcal{L}$. This is so since $a \vdash f(a)$ implies by dominance, for any EE relation $\preceq$ that, $a \preceq f(a)$.
- (3) $f(f(a))$ is logically equivalent to $f(a)$. So, if we view the points in the SOS as propositions instead, for any point $a$, $f(a)$ may be viewed as a fixed point. This follows from the fact that the $a$-cut $|a|_E$ and the $f(a)$-cut $(|f(a)|_E)$ are identical, and the sentences $a^{\#}_E \rightarrow a$ and $a^{\#}_E \rightarrow (a^{\#}_E \rightarrow a)$ are logically equivalent.
- (4) The mapping $f$ is many to one. This is so as long as there is at least one non-fixed-point. In general there would be many such points.
- (5) The mapping $f_e$ may map a point from an outer $\preceq_K$ band to an inner $\preceq_K$ band. This is easily seen by examining the points $x$, $z$, $x \wedge z$, $f(z)$ and $f(x \wedge z)$ in Figure 2. (The reasoning is is based on the fact that $x \wedge \alpha \rightarrow z$ and $x \wedge \alpha \rightarrow x \wedge z$ are logically equivalent, for any $\alpha$. Note that $z^{\#} = (x \wedge x)^{\#}$ and $x$ is a logical consequence of each of them. So there exists some sentence $\alpha$ such that $x \wedge \alpha$ is logically equivalent to $z^{\#}$.)
- (6) For any point $b = f(a)$, it holds that $|b|_E \cup \{b\} \vdash a$.

Now consider the point $z$ in B1. By (1) above, $f(z)$ is in B. Hence, it is not $\top$. By (2), then, it is in band 1. Therefore, it has to be in B1. Thus both $z$ and $f(z)$ are in A∪B1. If the latter, as desired, is a belief set, then by (6) above, $x \wedge z$ is in A∪B1. But clearly this is not so. Hence, A∪B1 is not a belief set.

The other reservation we had, that some member of B1 might be in conflict with A is not a serious possibility since, *ex hypothesi*, B1 and A are parts of A∪B which is consistent. But such a problem might arise later in this process. For instance some member of D2 might be in conflict with A∪B∪C.

An important point to be learned from this fairly technical exercise is the following. If a set $\mathbf{S}_n$ is to be used as the nth band in the posterior SOS, then $\mathbf{S}_n$ must satisfy the following closure property:

- **Band-Closure:** If any sentence $x$ is in $\mathbf{S}_n$, then all sentences connected with it by $f$-arrows (whether to or from $x$) are also in $\mathbf{S}_n$. Furthermore, $x$ is consistent with all sentences in $\mathbf{S}_m$ for all $m \leq n$, if there is an $(n+1)$-th band.

We suggest that the posterior SOS be constructed by minimally modifying the bands shown in Figure 4 such that the modified bands satisfy the **Band-Closure** property. Let $S_1, S_2, \ldots, S_n$ be the bands shown in Figure 4 (from center to periphery). Then the posterior SOS is given by $BC(S_1), BC(S_2), \ldots, BC(S_n)$ where $BC(S_i)$ is constructed in the following manner:

$$FP(S_i) = \{x \in S_i | f(x) = x\}$$
$$BC(S_i) = FP(S_i) \cup \{x | f(x) \in FP(S_i)\}$$

Thus $BC(S_i)$ is constructed by first identifying the fixed points in $S_i$ (namely $FP(S_i)$), and then inserting into it all sentences which are mapped into it by $f$. We could equivalently define $BC(S_i)$ simply as:

$$BC(S_i) = \{x | f(x) \in S_i\}.$$

We present a quick proof to this effect.

**Proof:** We need to show that $FP(S_i) \cup \{x | f(x) \in FP(S_i)\} = \{x | f(x) \in S_i\}$. First assume that $y \in LHS$. Then, either both $y \in S_i$ and $f(y) = y$ or both $f(y) \in S_i$ and $f(f(y)) = f(y)$. In either case $f(y) \in S_i$ whereby $y \in RHS$. Next, assume that $y \in RHS$. Hence $f(y) \in S_i$. Furthermore, obviously $f(f(y)) = f(y)$. Hence $f(y) \in FP(S_i)$ whereby $y \in LHS$. ■

It is worth noting that there might be $S_j$ such that $FP(S_j)$ are empty, and hence $BC(S_j)$ are empty too. This holds for all $j > j'$ such that $\neg \bigwedge E \in S_{j'}$. (By convention, such empty bands are deleted from the final configuration in the posterior SOS.) The reason why this holds is as follows. For any $x \notin E$, it holds that $x^{\#}_{E}$ is the sentence $\bigwedge E$. Furthermore, since $\neg \bigwedge E \vdash (\bigwedge E \rightarrow x)$ it follows that $\neg \bigwedge E \preceq_K (\bigwedge E \rightarrow x)$. Hence no fixed point exists outside the band in which $\neg \bigwedge E$ (which, not incidentally, is the fixed point $\bot^{\#}_{E} \rightarrow \bot$) exists. This explains how the consistency condition is satisfied by the construction of the FPO revision operation.

We will give here a concrete example just to show how it works. (The reader is recommended to draw systems of spheres – we avoid it to save some space.) Consider a language with just two atomic sentences, $a$ and $b$. Let $\preceq_K$ consist of the bands $B^K_0, \ldots, B^K_3$ where

$$B^K_0 = \{\top\}$$
$$B^K_1 = \{a \vee b\}$$
$$B^K_2 = \{a,\ b,\ a \leftrightarrow b,\ a \rightarrow b,\ b \rightarrow a,\ a \wedge b\} \text{ and}$$
$$B^K_3 = \{\bot,\ \neg a,\ \neg b,\ \neg a \vee \neg b,\ a \wedge \neg b,\ \neg a \wedge b,\ \neg(a \leftrightarrow b),\ \neg a \wedge \neg b\}.$$

(We are actualy simplifying the story – for instance, $B_0^K$ will also contain sentences like $a \rightarrow a$ which are logically equivalent to $\top$. But this will serve our purpose.) Let $\preceq_E$ consist of the bands $B_0^E, \ldots, B_2^E$ where

$$
\begin{aligned}
&B_0^E = \{\top\} \\
&B_1^E = \{\neg a,\ a \rightarrow b,\ \neg a \vee \neg b\} \text{ and} \\
&B_2^E = \{a \vee b,\ a,\ b,\ a \leftrightarrow b,\ b \rightarrow a,\ a \wedge b,\ \bot,\ \neg b,\ a \wedge \neg b,\ \neg a \wedge b, \\
&\qquad \neg(a \leftrightarrow b),\ \neg a \wedge \neg b\}.
\end{aligned}
$$

(It is easily verified that these two SOS's satisfy the appropriate constraints. The former may be viewed as the natural EE relation satisfying the constraint $\bot \prec a \wedge b \prec a \vee b \prec \top$ whereas the latter is the natural EE relation given the constraint $\bot \prec \neg a \prec \top$.) Given these two SOS's, the intermediary SOS (analogous to Figure 4) is given by the bands $S_0, \ldots, S_5$ where

$$
\begin{aligned}
&S_0 = \{\top\} \\
&S_2 = \{a \rightarrow b\} \\
&S_3 = \{\neg a,\ \neg a \vee \neg b\} \\
&S_4 = \{a \vee b\} \text{ and} \\
&S_5 = \{a,\ b,\ a \leftrightarrow b,\ b \rightarrow a,\ a \wedge b\ \bot,\ \neg b,\ a \wedge \neg b,\ \neg a \wedge b,\ \neg(a \leftrightarrow b),\ \neg a \wedge \neg b\}.
\end{aligned}
$$

Now, noting that the sentences $\neg a^{\#}$, $(a \rightarrow b)^{\#}$, $(\neg a \vee \neg b)^{\#}$ and $\top^{\#}$ are all logically equivalent to $\top$, and $\alpha^{\#}$ for any other sentence $\alpha$ (i.e., not equivalent to $\neg a$, $a \rightarrow b$ or $\neg a \vee \neg b$) is logically equivalent to $\neg a$, we get the following $f$-values.

| $\alpha$ | $f(\alpha)$ | $\alpha$ | $f(\alpha)$ |
|---|---|---|---|
| ====== | ===== | ===== | ===== |
| $\top$ | $\top$ | $\neg a$ | $\neg a$ |
| $a \rightarrow b$ | $a \rightarrow b$ | $\neg a \vee \neg b$ | $\neg a \vee \neg b$ |
| $a \vee b$ | $a \vee b$ | $a$ | $a$ |
| $b$ | $a \vee b$ | $a \leftrightarrow b$ | $b \rightarrow a$ |
| $b \rightarrow a$ | $b \rightarrow a$ | $a \wedge b$ | $a$ |
| $\bot$ | $a$ | $\neg b$ | $b \rightarrow a$ |
| $a \wedge \neg b$ | $a$ | $\neg a \wedge b$ | $a \vee b$ |
| $\neg(a \leftrightarrow b)$ | $a \vee b$ | $\neg a \wedge \neg b$ | $b \rightarrow a$ |

Accordingly, we get

$$
\begin{aligned}
&BC(S_0) = \{\top\} \\
&BC(S_1) = \{a \rightarrow b\} \\
&BC(S_2) = \{\neg a,\ \neg a \vee \neg b\}
\end{aligned}
$$

$BC(S_3) = \{a \vee b,\ b,\ \neg a \wedge b,\ \neg(a \leftrightarrow b)\}$
$BC(S_4) = \{a,\ a \leftrightarrow b,\ b \rightarrow a,\ a \wedge b,\ \bot,\ \neg b,\ a \wedge \neg b,\ \neg a \wedge \neg b\}$ and
$BC(S_5) = \emptyset$.

By convention, the empty set $BC(S_5)$ is deleted, and the rest, namely, $BC(S_0), \ldots, BC(S_4)$ determine the resultant EE relation. (This resultant EE relation is the natural completion of the constraint $\bot \prec \neg a \wedge b \prec \neg a \prec a \rightarrow b \prec \top$.)

To summarize, then, as a first approximation to the revision of $\preceq_K$ by $\preceq_E$, we construct the (analogue of) Figure 4 via the (analogue of) Figure 3. Then we let each fixed point in each band pull the sentences that point at it into their respective bands. That results in the desired system of bands (spheres) representing the posterior EE relation. Note in particular that, given any two arbitrary sentences $x$ and $y$, (1) if $x$ is closer to the center than $y$ is in the evidential system of spheres, then the "fixed point of $x$" (i.e. $f(x)$) will be closer to the center in Figure 4 than the "fixed point of $y$" would be. Hence, if $x \prec_E y$ then $x \prec_{K*E} y$. (2) On the other hand, if $x$ and $y$ are equidistant from the center in the evidential system of spheres, it is possible that even if $x$ is closer to the center than $y$ in Figure 4, their respective fixed points are not so related. Hence we are not allowed any short-cut; we have to compute the relation between the fixed points $f(x)$ and $f(y)$ in Figure 4. In other words, when $x \equiv_E y$ then $x \preceq_{K*E} y$ iff $f(x)$ is not farther from the center than $f(y)$ is in Figure 4. However, $f(x)$ is simply $x^\#_E \rightarrow x$ and $f(y)$ is simply $x^\#_E \rightarrow y$. Hence, this condition translates to: when $x \equiv_E y$ then $x \preceq_{K*E} y$ iff $(x^\#_E \rightarrow x) \preceq_K x^\#_E \rightarrow y$. Thus, combining these two results, we get Definition 3.

# 6 A Representation Theorem

In this section we will present our primary result. Two results follow directly from definition 3:

**Observation 5** *If $x \prec_E y$ then $x \prec_{K*E} y$*

**Observation 6** *$x \equiv_{K*E}\ y$ iff both $x \equiv_E y$ and $(x^\#_E \rightarrow x) \equiv_K\ (x^\#_E \rightarrow y)$.*

Of the two above, proof of the former is obvious; proof of the latter can be found in [27].

The following two theorems jointly show that every FPO entrenchment revision operation is a well behaved entrenchment revision operation. Again, we omit their proofs since they can be found in [27], although in the case of the Theorem (2) a slight modification is necessary.

**Theorem 1** *Let* $\preceq_K$ *and* $\preceq_E$ *be two EE relations, and* $\preceq_{K*E}$ *be a relation constructed via Definition 3. Then,* $\preceq_{K*E}$ *is again an EE relation.*

**Theorem 2** *Any EE relation constructed by Definition 3 satisfies* $(EE2^*)$ *and* $(EE3^*)$.

Furthermore, Theorem 3 shows that the FPO entrenchment revision operation is the only well behaved entrenchment revision operation. Hence, the constraints $(EE1^*) \perp (EE3^*)$ offer an axiomatization of the FPO entrenchment revision operation.

**Lemma 1** *If* $\neg(x_E^{\#}) \vdash z$ *and* $(x_E^{\#} \rightarrow x) \prec_E z$*, then* $\vdash z$.

**Proof**: We show only the non-trivial case where $\preceq_E$ is non-absurd, and $\not\vdash x$. Let $\neg(x_E^{\#}) \vdash z$ and $x_E^{\#} \rightarrow x \prec_E z$. Since, by Observation 3, $x_E^{\#} \rightarrow x \equiv_E x$ and $x_E^{\#} \rightarrow x \prec_E z$, it follows that $x \prec_E z$. It follows then from the construction of $x_E^{\#}$ that $x_E^{\#} \vdash z$. Thus both $x_E^{\#} \vdash z$ and $\neg(x_E^{\#}) \vdash z$ whereby $\vdash z$. ■

**Theorem 3** *Every well behaved entrenchment revision operation is an FPO entrenchment revision operation.*

**Proof:** Assume that the revision operation * satisfies constraints $(EE1^*)$ - $(EE3^*)$. For reductio, assume further that either (1) both $x \prec_E y$ and $x \not\prec_{K*E} y$ or (2) $x \equiv_E y$, $(x_E^{\#} \rightarrow x) \preceq_K (x_E^{\#} \rightarrow y)$ but $x \succ_{K*E} y$. It will be sufficient to show that these assumptions lead to a contradiction. The first case is trivial – it conflicts with the constraint $(EE2^*)$. As to the second case, first we note that since $x \equiv_E y$, it follows by Observation 3 that $(x_E^{\#} \rightarrow x) \equiv_E (x_E^{\#} \rightarrow y)$. First we claim that (3) $\forall_z, \forall_{z' \in Cn(x_E^{\#} \rightarrow x, x_E^{\#} \rightarrow y): z' \succ_E (x_E^{\#} \rightarrow x)} z \preceq_E z'$ iff $z \preceq_K z'$. The proof goes as follows: since $z' \in Cn(x_E^{\#} \rightarrow x, x_E^{\#} \rightarrow y)$, surely $\neg(x_E^{\#}) \vdash z'$. Furthermore, by assumption, $x_E^{\#} \rightarrow x \prec_E z'$. Hence, by lemma 1, $\vdash z'$. Since both $\preceq_E$ and $\preceq_K$ are EE relations, it follows that both $z \preceq_E z'$ and $z \preceq_K z'$ (by EE4). Hence, $z \preceq_E z'$ iff $z \preceq_K z'$. It follows from (3) and the fact that $(x_E^{\#} \rightarrow x) \equiv_E (x_E^{\#} \rightarrow y)$, by $(EE3^*)$, that $(x_E^{\#} \rightarrow x) \preceq_{K*E} (x_E^{\#} \rightarrow y)$ iff $(x_E^{\#} \rightarrow x) \preceq_K (x_E^{\#} \rightarrow y)$. Since by assumption, $(x_E^{\#} \rightarrow x) \preceq_K (x_E^{\#} \rightarrow y)$, it follows that (4) $(x_E^{\#} \rightarrow x) \preceq_{K*E} (x_E^{\#} \rightarrow y)$. By assumption, $x \succ_{K*E} y$. By $(EE1^*)$, $\preceq_{K*E}$ is an EE relation. Hence by dominance, $x \preceq_{K*E} (x_E^{\#} \rightarrow x)$. It follows then, from (4), that $y \prec_{K*E} x \preceq_{K*E} (x_E^{\#} \rightarrow x) \preceq_{K*E} (x_E^{\#} \rightarrow y)$. Hence, by transitivity of $\preceq_{K*E}$ it follows that (5) $y \prec_{K*E} x_E^{\#} \rightarrow y$. By the conjunctiveness of $\preceq_{K*E}$ it follows then that either $(x_E^{\#} \rightarrow y) \preceq_{K*E} y$ or $(x_E^{\#} \vee y) \preceq_{K*E} y$. From (5), it follows then that $(x_E^{\#} \vee y) \preceq_{K*E} y$, from which, dominance, it follows that $(x_E^{\#} \vee y) \equiv_{K*E} y$. Since, by assumption, $x \succ_{K*E} y$, it follows that (6) $x \succ_{K*E} (x_E^{\#} \vee y)$. However, by Observation 1, $x \prec_{K*E} x_E^{\#}$, and by dominance, $x_E^{\#} \preceq_{K*E} x_E^{\#} \vee y$, from which, by transitivity, it follows that (7) $x \prec_{K*E} x_E^{\#} \vee y$. But (7) contradicts (6). ■

Theorems 1-3 jointly entail the representation theorem

- An entrenchment revision operation is well behaved if and only if it is an FPO entrenchment revision operation.

Thus we have constructed an entrenchment revision operation, namely an FPO entrenchment revision operation, which maps a pair of EE relations into a posterior EE relation, thus solving the problem of iterated belief change. Furthermore, this operation is "better" than other possible entrenchment revision operations since it is the only "well behaved" entrenchment revision operation in the sense that no other entrenchment revision operation satisfies the conditions $(EE1^*)$–$(EE3^*)$.

The epistemic input for this operation is represented as an EE relation. Since the epistemic content of an EE relation is a set of sentences, it is a solution to the problem of multiple belief change. Moreover, since FPO entrenchment revision operations satisfy $(G1^*$-$G8^*)$ which are intuitive generalizations of the Gärdenfors postulates of belief revision, the solution to the multiple belief change problem offered is a nice one. That is the object of discussion in the next section.

# 7 Multiple Belief Change

It may be argued, based on the work of Niederée [30], that revision by a set of sentences $E$ should not be equated either with revision by $\bigwedge(E)$ or with revision by $\bigvee(E)$. Does the account developed in this paper give any alternativeΓ At first blush, the result may not look satisfactory. According to $(G6^*)$, given the same prior EE relation, if two evidential EE relations have the same epistemic content, then the corresponding posterior EE relations will also have the same epistemic content. When applied to this problem, what this means is that revision of a given belief state by $\bigwedge(E)$ and by $E$ will result in belief states with the same epistemic content. To that extent, the result is disappointing. But, all is not lost. In fact, the intuition that accepting $e_1 \wedge e_2$ is different from accepting $\{e_1, e_2\}$ need not mean that accepting them leads to belief states with different epistemic contents. The intuition in question is easily explained within the framework of this paper. Given a suitable representation of these two pieces of evidence as two different EE relations, it can be shown that the corresponding posteriors represent two different belief *states*.

A reasonable principle to follow while representing these naked pieces of evidence as EE relations is:

> in the absence of any other information, $e_1 \wedge e_2$ is represented as an EE relation $\preceq$ such that:

- $\bot \prec e_1 \equiv e_2 \prec \top$

and $\{e_1, e_2\}$ is represented as an EE relation $\preceq$ such that:

- $\bot \prec e_1 \equiv e_2 \prec e_1 \vee e_2 \prec \top$

Hans Rott [35, 36] has suggested the following recipe to construct an SEE relation $\preceq_E$ given a set of sentences $E$.

**Rott1:** The evidential SEE relation $\preceq_E$ generated by the evidence $E$ is given by: $x \prec_E y$ iff $\not\vdash x$ and for every $E' \subseteq E$ such that $E' \not\vdash y$ there exists an $E''$ such that $E' \subset E'' \subseteq E$ and $E'' \not\vdash x$.

However, since the evidence is required to be an EE relation, we might slightly modify Rott's recipe in order to construct the evidential EE relation in the following manner:

**Rott2:** $x \prec_E y$ iff $E \not\vdash \bot$, $\not\vdash x$ and for every $E' \subseteq E$ such that $E' \not\vdash y$ there exists an $E''$ such that $E' \subset E'' \subseteq E$ and $E'' \not\vdash x$.

(Note that if $E \vdash \bot$ then using (Rott2) we get that $x \preceq_E y$ for all $x$ and $y$.) We will assume that when the evidence acquired is just a set $E$ of sentences (without any constraints), Rott's revised proposal (Rott2) is used to construct the evidential EE relation.[11]

It is easily verified that $\preceq_{e_1 \wedge e_2}$ and $\preceq_{\{e_1, e_2\}}$ constructed from $\{e_1 \wedge e_2\}$ and $\{e_1, e_2\}$ respectively, following (Rott2), satisfy the two constraints suggested earlier in this paragraph. In particular, according to the relation $\preceq_{e_1 \wedge e_2}$, all sentences (including $\bot$) that are not logical consequences of $e_1 \wedge e_2$ are minimally entrenched, all nontautological consequences of $e_1 \wedge e_2$ are equally entrenched (and more entrenched than $\bot$) and $\top$ is the only sentence to be maximally entrenched. On the other hand, according to the latter EE relation, namely $\preceq_{\{e_1, e_2\}}$, all sentences (including $\bot$) that are not logical consequences of $e_1 \wedge e_2$ are minimally entrenched, all consequences of $e_1 \wedge e_2$ that are not consequences $e_1 \vee e_2$ are equally entrenched (and more entrenched than $\bot$), all nontautological consequences of $e_1 \vee e_2$ are equally entrenched (and more entrenched than $e_1 \wedge e_2$) and $\top$ is the only sentence to be maximally entrenched.

The tangible effect of such different representations, namely, $\preceq_{e_1 \wedge e_2}$ and $\preceq_{\{e_1, e_2\}}$ is as follows. According to $\preceq_{e_1 \wedge e_2}$, if $e_1 \wedge e_2$ is the sole cause of introducing $e_1 \vee e_2$ into one's belief set, then $e_1 \wedge e_2$, $e_1$, $e_2$ and $e_1 \vee e_2$ are all equally corrigible. Accordingly, if later experience leads one to give up $e_1 \wedge e_2$, all these beliefs in question are lost. On the other hand, according to $\preceq_{\{e_1, e_2\}}$, if $\{e_1, e_2\}$ is the sole cause of introducing $e_1 \vee e_2$ into one's belief corpus, then, although $e_1 \wedge e_2$, $e_1$ and $e_2$ are equally corrigible,

$e_1 \vee e_2$ is less corrigible than them. Accordingly, giving up $e_1 \wedge e_2$ results in giving up both $e_1$ and $e_2$ individually, but $e_1 \vee e_2$ survives the assault.

This means that although revision of a given belief state by $e_1 \wedge e_2$ and $\{e_1, e_2\}$ leads to belief states that have the same epistemic content, they may exhibit different dynamic behavior. This can be shown by a simple example.

Consider the prior EE relation $\preceq_\top$. That is, the agent in this state has no empirical knowledge. Consider the revision of this belief state by $e_1 \wedge e_2$ and $\{e_1, e_2\}$ respectively, following the principle outlined above. The resultant belief states are $\preceq_{e_1 \wedge e_2}$ and $\preceq_{\{e_1, e_2\}}$ respectively [26]. These two states are statically equivalent since they have the same epistemic content $(Cn(\{e_1 \wedge e_2\}) = Cn(\{e_1, e_2\})$. However, as the following argument shows, these two states are not dynamically equivalent. The point made in this argument was informally stated earlier in the context of Hansson's "hamburger example" (§3). However, we present it here for the sake of completeness.

Consider the revision of these two states by the EE relation $\preceq_{\neg e_1} = \preceq_{\{\neg e_1\}}$. Whereas $e_2$ is believed in $\preceq_{\{e_1,e_2\} * \neg e_1}$, it is not believed in $\preceq_{(e_1 \wedge e_2) * \neg e_1}$. Therefore, these resulting states are statically different. Note that since neither $e_1$ nor $e_2$ is a logical consequence of $\neg e_1$, $\bot \equiv_{\neg e_1} e_1 \equiv_{\neg e_1} e_2$. Hence, $e_2 \in (e_1 \wedge e_2) * \neg e_1$ just in case $(\bot^{\#}_{\neg e_1} \rightarrow \bot) \prec_{e_1 \wedge e_2} (\bot^{\#}_{\neg e_1} \rightarrow e_2)$ and $e_2 \in \{e_1, e_2\} * \neg e_1$ if $(\bot^{\#}_{\neg e_1} \rightarrow \bot) \prec_{\{e_1, e_2\}} (\bot^{\#}_{\neg e_1} \rightarrow e_2)$. Surely $\bot^{\#}_{\neg e_1}$ is logically equivalent to $\neg e_1$. Hence $e_2 \in (e_1 \wedge e_2) * \neg e_1$ just in case $e_1 \prec_{e_1 \wedge e_2} (e_1 \vee e_2)$ and $e_2 \in \{e_1, e_2\} * \neg e_1$ if $e_1 \prec_{\{e_1, e_2\}} (e_1 \vee e_2)$. It follows then that $e_2$ is lost in the former process, but retained in the latter process. This result is hardly surprising. In the former case, the only justification for $e_2$, namely $e_1 \wedge e_2$, was retracted when the new evidence conflicted with $e_1$. Accordingly, $e_2$ was also lost in the process. However, in the latter case, $e_1 \wedge e_2$ was not a justification either for $e_1$ or $e_2$ since both $e_1$ and $e_2$ had independent warrants. Hence the evidence against $e_1$ did not affect the status of $e_2$. This is as it should be, and squares well with the literature on multiple belief change [16, 21, 36].

# 8 Summary

We started with the goal of constructing a belief revision operation that, apart from being "well-behaved", is capable of handling belief revision when the new information comes in the form of a set of sentences and leaves enough room for subsequent belief revision using the same operation. We began with a discussion of the AGM framework and how it deals with belief revision. In the next two sections we argued that both the epistemic state and epistemic input should be represented as epistemic entrenchment relations. Finally, we constructed a belief revision operation which, and only which, satisfies three conditions. The first of these is a closure condition which requires that the result of revising an EE (epistemic entrenchment) relation by another must also be an EE relation. This allows iterated belief revision. The other two conditions require the posterior EE relation (belief state) to be well behaved relative to both the prior belief state and the information which necessitated the belief revision. A simple generalization of the AGM framework also allows this revision operation to handle multiple belief revision, as discussed in the last section.

# Endnotes

* The authors are thankful to John G. Bennett, Henry E. Kyburg, Jr., Norman Foo, Pavlos Peppas, Maurizio Pagnucco, Bob Hadley, members of the Knowledge Systems Group at the University of Sydney and participants of the AI'93 workshop on "Belief Revision: Bridging the Gap Between Theory and Practice" at Melbourne where a version of this paper was presented, for their suggestions. The authors also express their gratitude to the two anonymous referees for *Synthese* for their excellent comments and suggestions. Of course it is only the authors that are responsible for the errors that remain.

[1]See however Hansson's [17] p.19 where he argues that under fairly weak assumptions contraction by $\overline{X}$ may be equated with the contraction by $\{\neg x_1 \vee \neg x_2 \vee \ldots\}$.

[2] We say these conditions are the properties of a standard epistemic entrenchment relation in order to distinguish it from other related entrenchment relations such as Rott's GEE relation [34] and our EE relation (this paper, §3).

[3]It should be noted that unlike the other approaches mentioned here, Hansson's account is about belief corpora that need not be closed under the consequence operation $Cn$.

[4]This summarization is only approximate. For instance, instead of the family $\mathcal{R}$, Hansson uses a super selector whereas Rott and Schlechta use $K$-independent background relations. The summary presented here however captures the spirit of the works in question. For a slightly more detailed account see [27]. pp. 379-380.

[5]As one of the anonymous referees has pointed out, the talk of conceptual frameworks here is misleading since the AGM literature concerns belief change at a sentential level whereas concepts would be located at a sub-sentential level. We agree. However, our loose talk here is not completely off the mark. Isaac Levi [24] (§2.2) equates the conceptual framework of an agent with a family of potential states of full belief and represents a potential state of full belief at the linguistic level as a belief corpus. Since a belief corpus for Levi is, more or less, a belief set for AGM, and we on the other hand use an EE relation from which the belief set in question can be extracted, for our purpose, the conceptual framework of an agent is represented as a family of EE relations.

[6]Note that since belief sets are simply sets of sentences, a belief set cannot in itself contain information about the epistemic pedegree of its members. From the informal presentation of the above example, it is clear that Al's and Bill's beliefs have different epistemic pedigrees, so, in a sense their beliefs are different. But this discrimination cannot be made in the AGM framework at the "knowledge level".

[7]We are *not* claiming, however, that the foundations theory of belief change is never more appropriate than its coherence counterpart. For instance, in the hamburger example in question, it might still be better to use the belief base dynamics formalism because of its intuitive appeal and adequacy. However, the fact remains that the particular argument – since the distinction between the basic beliefs and the inferred beliefs is obliterated in the coherence approach, the belief base dynamics formalism approach is superior – is no longer available to its proponents. See [29] for some related results.

[8]For convenience, we occasionally abbreviate $Cn(\{\ldots\})$ as $Cn(\ldots)$.

[9]For the statement of the eight Gärdenfors postulates, see, for example, [9]. We show later that every well behaved entrenchment revision operation is an FPO entrenchment revision operation (Theorem 3). It has been shown [26, 27] that every FPO entrenchment revision operation satisfies $(G1^*)$-$(G8^*)$. Hence it follows that every well behaved entrenchment revision operation satisfies $(G1^*)$-$(G8^*)$.

[10]We are indebted to Pavlos Peppas for this construction. In [27] the construction of the operation # is simpler, but presumes a complete language with a complete logic. The current construction effectively achieves the same without commiting to infinite conjunctions and infinite disjunctions.

[11]A special case of this is when $E$ is a singleton, say $\{e\}$. This is the case one should consider if one were to compare our approach with the AGM approach. This question is explored to some extent in [28] (see the construction of the operation $\odot$ there).

# References


[1] Abiteboul, S., "Updates: A New Frontier, in M. Gyssens", *et al.* (eds), *Proceedings of the Second International Conference on Database Theory,* LNCS 326, pp. 1-18, Springer 1988.

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

[3] Alchourrón, C. E. and D. Makinson, "On the Logic of Theory Change: Safe Contractions", *Studia Logica,* vol. 44, pp.405-422, 1985.

[4] Alchourrón, C. E. and D. Makinson, "Maps Between Some Different Kinds of Contraction Function: The Finite Case", *Studia Logica,* vol. 45, pp.187-198, 1986.

[5] Brewka, G, *Nonmonotonic Reasoning: Logical Foundations of Commonsense,* Cambridge University Press, 1991.

[6] Doyle, J, "A Truth Maintenance System", *Artificial Intelligence,* vol. 12, pp.231-272, 1979.

[7] Doyle, J, "The Ins and Outs of Reason Maintenance", in *Proceedings of the Eighth International Joint Conference on Artificial Intelligence,* pp.349-351, 1983.

[8] Fuhrmann, André, "Theory Contraction Through Base Revision", *Journal of Philosophical Logic,* vol. 20, pp.175-203, 1991.

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

[10] Gärdenfors, Peter "The Dynamics of Belief Systems: Foundations vs. Coherence Theories", *Revue Internationale de Philosophie* vol. 44, pp.24-46, 1990.

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

[12] Gärdenfors, Peter and Hans Rott, "Belief Revision", in *Handbook of Logic in Artificial Intelligence and Logic Programming* vol.IV, Epistemic and Temporal Reasoning, D. M. Gabbay, C. J. Hogger and J. A. Robinson (eds.), Oxford, 1995.

[13] Ghosh, A., P. Hadjinian, A. Sattar, J. You and R. Goebel, "Iterated Belief Change", Griffith University CIT Research Report no 64.

[14] Ginsberg, M. L. and D. E. Smith, "Reasoning About Action I: A Possible Worlds Approach", in F. M. Brown (ed), *The Frame Problem in Artificial Intelligence: Proceedings of the 1987 Workshop*, Morgan Kaufmann, pp. 133-185, 1987.

[15] Grove, Adam "Two Modelings for Theory Change", *Journal of Philosophical Logic,* vol. 17, pp. 157-170, 1988.

[16] Hansson, Sven Ove "New Operators for Theory Change", *Theoria,* vol. 55, pp. 114-136, 1989.

[17] Hansson, Sven Ove *Belief Base Dynamics*, Uppsala, 1991.

[18] Hansson, Sven Ove "A Dydadic Representation of Belief", in Peter Gärdenfors ed. *Belief Revision*, CUP, pp. 89-121, 1992.

[19] Hansson, Sven Ove "Changes in Preference", *Theory and Decision*, **38**, pp. 1-28, 1995.

[20] Hansson, Sven Ove "Reversing the Levi Identity", *Journal of Philosophical Logic,* **22**, pp. 637-669, December 1993.

[21] Hansson, Sven Ove "Changes on Disjunctively Closed Bases", *Journal of Logic, Language and Information*, **2**, pp. 255-284, October 1993.

[22] Jeffrey, Richard C. *The Logic of Decision,* 2nd ed, Chicago University Press, 1983.

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

[24] Isaac Levi, *Fixation of Belief and its Undoing,* Cambridge University Press, 1991.

[25] Makinson, D. and P. Gärdenfors, "Relations Between the Logic of Theory Change and Nonmonotonic Logic", in A. Fuhrmann and M. Mourreau, (eds.), *The Logic of Theory Change,* Lecture Notes in Artificial Intelligence 465, Springer Verlag, 1991.

[26] Nayak, Abhaya C. *Studies in Belief Change*, Ph.D. Dissertation, Department of Philosophy, University of Rochester, Rochester, New York, 1993.

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

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

[29] Nayak, A. C., M. Pagnucco, N. Foo and R. Kwok, "Entrenchment and Retractability: A Preliminary Report", *Proceedings of the 8th Australian Joint Conference on Artificial Intelligence,* pp. 219-226, Xin Yao (Ed.), World Scientific, 1995.

[30] Niederée, R. "Multiple Contraction. A Further Case against Gärdenfors' Principle of Recovery", in A. Fuhrmann and M. Mourreau, eds., *The Logic of Theory Change,* Lecture Notes in Artificial Intelligence 465, Springer Verlag, pp. 322-334, 1991.

[31] Quine, W. V. "Two Dogmas of Empiricism", in A. P. Martinich (ed.) *The Philosophy of Language*, OUP 1985, pp. 26-39.

[32] Quine, W. V. and S. Ullian, *Web of Belief,* Random House, 1970.

[33] Rott, Hans "Two Methods of Constructing Contractions and Revisions of Knowledge Systems", *Journal of Philosophical Logic,* vol. 20, pp. 149-173, 1991.

[34] Rott, Hans "Preferential Belief Change Using Generalized Epistemic Entrenchment", *Journal of Logic, Language and Information*, vol. 1, 45-78, 1992.

[35] Rott, Hans "Modellings for Belief Change: Base Contraction, Multiple Contraction and Epistemic Entrenchment (Preliminary Report)", in *Logics in AI*, European Workshop JELIA '92, D. Pearce and G. Wagner eds., Lecture Notes in Artificial Intelligence 633, Berlin: Springer, pp.139-153.

[36] Rott, Hans "Modellings for Belief Change: Prioritization and Entrenchment", *Theoria*, **LVIII**, pp. 21-57, 1992.

[37] Rott, Hans "Coherent Choice and Epistemic Entrenchment", (manuscript).

[38] Schlechta, Karl "Theory Revision and Probability", *Notre Dame Journal of Formal Logic*, vol. 32, 307-319.

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

[40] Stalnaker, Robert "A Theory of Conditionals", in *Ifs*, (eds) W. L. Harper, *et al.*, Reidel, pp. 41-55, 1981.

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

[42] Winslett, M. S., "Reasoning about Action using a Possible Models Approach" in *Proceedings of the 7th National Conference of the Americal Association for Artificial Intelligence,* St. Paul, pp. 89-93, 1988.

†Knowledge Systems Group
Department of Computer Science, F09
University of Sydney
NSW 2006, Australia

‡Simon School of Business
University of Rochester
Rochester, N.Y. 14627, USA