Iterated Belief Contraction from First Principles *

Abhaya C. Nayak
Department of Computing
Macquarie University
Sydney, NSW
Australia 2109
abhaya@ics.mq.edu.au

Randy Goebel
Department of Comp. Sc.
University of Alberta
Edmonton, Alberta
Canada T6G 2H1
goebel@cs.ualberta.ca

Mehmet A. Orgun
Department of Computing
Macquarie University
Sydney, NSW
Australia 2109
mehmet@ics.mq.edu.au

Abstract

Importance of contraction for belief change notwithstanding, literature on iterated belief change has by and large centered around the issue of iterated belief revision, ignoring the problem of iterated belief contraction. In this paper we examine iterated belief contraction in a principled way, starting with Qualified Insertion, a proposal by Hans Rott. We show that a judicious combination of Qualified Insertion with a well-known Factoring principle leads to what is arguably a pivotal principle of iterated belief contraction. We show that this principle is satisfied by the account of iterated belief contraction modelled by Lexicographic State Contraction, and outline its connection with Lexicographic Revision, Darwiche-Pearl's account of revision as well as Spohn's Ordinal ranking theory.

Keywords: Belief Change, Information State Change, Iterated Belief Contraction.

Literature on belief change deals with the problem of how new evidence impinges upon the current knowledge of a rational agent. The pioneering works in the area such as [Alchourrón *et al.*, 1985] provide the formal framework, and the solution for "one-shot belief change". Follow up work in the area, e.g. [Williams, 1994; Nayak, 1994; Boutilier, 1996] have explored the related issue of *Iterated Belief Change* that deals with sequential changes in belief. This latter research has by and large been confined to iterated belief revision: *Given belief corpus K, two sequential pieces of evidence x, y and revision operator $*$, how do we construct the resultant corpus $(K^*_x)^*_y$?* However, the accompanying problem of iterated belief contraction, namely, *Given belief corpus K, two beliefs x, y that are to be sequentially removed from K, and contraction operator $-$, how do we construct the resultant corpus $(K^-_x)^-_y$?* has, for no obvious reason, been rather neglected. Very few reseach works, e.g., [Bochman, 2001; Chopra *et al.*, 2002; Rott, 2001; 2006] and [Nayak *et al.*, 2006], have addressed this problem. The primary aim of this paper is to examine this problem from what we may call the "first principles" of belief change.

In Section 1, we introduce the problem of iterated belief contraction, and as a starting point, take up a proposal by Hans Rott [Rott, 2001] called Qualified Insertion. In the next section, we examine this principle and its variations, and show that when combined with Factoring, a well-known result in belief change, Qualified Insertion can lead to a good account of iterated belief contraction. In Section 3 we study some interesting properties of this contraction, followed by its semantic modelling via Lexicographic state contraction in Section 4. We conclude with a short summary.

1 Background

The theory of belief change purports to model how a current theory or body of beliefs, K, can be rationally modified in order to accommodate a new observation x. An observation, such as x, is represented as a sentence in a propositional language $\mathcal{L}$, and a theory, such as K, is assumed to be a set of sentences in $\mathcal{L}$, closed under a supra-classical consequence operation, Cn. Since the new piece of information x may contravene some current beliefs in K, chances are, some beliefs in K will be discarded before x is eased into it. Accordingly, three forms of belief change are recognised in the belief change framework:

1. CONTRACTION: K^-_x is the result of discarding some unwanted information x from the theory K
2. EXPANSION: K^+_x is the result of simple-mindedly incorporating some information x into the theory K, and
3. REVISION: K^*_x is the result of incorporating some information x into the theory K in a manner so as to avoid internal contradiction in K^*_x.

The intuitive connection among these operators is captured by the following two identities named, respectively, after Isaac Levi and William Harper:

LEVI IDENTITY: $K^*_x = (K^-_{\neg x})^+_x$, and

HARPER IDENTITY: $K^-_x = K^*_{\neg x} \cap K$.

A belief change (revision, contraction and expansion) operation is AGM-RATIONAL if it satisfies the corresponding set of *AGM rationality postulates*. The three sets of postulates, along with motivation and interpretation for them, may be found in [Gärdenfors, 1988]. It is

*The authors acknowledge the research support from the Australian Research Council (ARC), and thank Maurice Pagnucco and the referees for very helpful suggestions.

well known that a revision operator constructed from an AGM-rational contraction operation via the Levi Identity is AGM rational; and conversely AGM-rational revision operations can be used to construct AGM-rational contraction operations via the Harper Identity. For convenience we reproduce the AGM postulates for contraction below.

$(1-)$	K_x^- is closed under Cn	CLOSURE
$(2-)$	$K_x^- \subseteq K$	INCLUSION
$(3-)$	If $x \notin K$, then $K_x^- = K$	VACUITY
$(4-)$	If $\not\vdash x$ then $x \notin K_x^-$	SUCCESS
$(5-)$	$K \subseteq (K_x^-)_x^+$	RECOVERY
$(6-)$	If $\vdash x \leftrightarrow y$ then $K_x^- = K_y^-$	PRESERVATION
$(7-)$	$K_x^- \cap K_y^- \subseteq K_{x \wedge y}^-$	INTERSECTION
$(8-)$	If $x \notin K_{x \wedge y}^-$ then $K_{x \wedge y}^- \subseteq K_x^-$	CONJUNCTION

The expansion operation is very easily constructed: $K_x^+ = Cn(K \cup \{x\})$. Contraction and Revision operations are relatively more sophisticated operations since they deal with choice.

It is easily noticed that the AGM postulates of contraction deal with "one-shot" belief contraction. The only constraint of interest on iterated contraction that they impose is:

Vacuity Specialised If $y \notin K_x^-$, then $(K_x^-)_y^- = K_x^-$

which is not really much of a constraint.[1] In particular, postulates $(7-)$ and $(8-)$ impose very little. Consider the following example to see how ineffective the AGM postulates are in dealing with sequential contraction.

Example 1 *Suppose I believe on independent grounds that both a:* `Argentina will qualify for the FIFA world-cup final` *and b:* `Brazil will qualify for the FIFA world-cup final`.[2] *Thus my knowledge corpus is $K = Cn(a, b)$, whereby $a \vee b \in K$.*

Since I believe a and b on independent grounds, arguably $K_a^- = Cn(b)$ and $K_b^- = Cn(a)$. Moreover, since my belief in $a \vee b$ is assumed to have no independent standing, we would expect a "correct" belief contraction operation to behave in a way such that both $a \vee b \notin (K_a^-)_b^-$ and $a \vee b \notin (K_b^-)_a^-$.

This however is not mandated by the AGM postulates. We can plausibly argue that $K_{a \wedge b}^- = Cn(b)$ given Brazil's past record.This will satisfy all the eight postulates of contraction.[3] *However, since $b \in K_a^-$, the postulates do not constrain the content of $(K_a^-)_b^-$ in any interesting way at all. In particular, the postulates cannot insist on the absence of $a \vee b$ in the resultant belief corpus, contrary to our intuition.*

Not many researchers have examined iterated belief contraction in a principled way, an exception being Hans Rott who has advocated the following principle [Rott, 2001](p.221):

Qualified Insertion

$$(K_x^-)_y^- = \begin{cases} K_x^- & \text{if } y \notin K_x^- \\ K_x^- \cap K_y^- & \text{otherwise.} \end{cases}$$

This is indeed an appealing suggestion. Its first part is simply Vacuity Specialised, justified by the AGM rationality postulates. As to the second part of this principle, since both x and y are individually removed from K sequentially, it stands to reason that when the second contraction is non-vacuous, the result should consist of the beliefs that resist removal in individual contractions of K by x and y.

Despite the intuitive plausibility of Qualified Insertion, its inadequacy is established by the Example 2 below.

Example 2 *Consider again the scenario in Example 1. We know that $b \in K_a^-$. In accordance with Qualified Insertion, then, $(K_a^-)_b^- = K_a^- \cap K_b^-$. Moreover, $a \vee b \in K_a^-$ and $a \vee b \in K_b^-$ whereby $a \vee b \in K_a^- \cap K_b^-$. Hence $a \vee b \in (K_a^-)_b^-$. However, since a and b are the only reasons for $a \vee b$ to be in the knowledge corpus, $a \vee b$ should not be in $(K_a^-)_b^-$.*

We find that Qualified Insertion, though intuitive, fails to offer an acceptable account of Iterated Belief Contraction. We take this to be the starting point of this paper. In the next section we analyse this principle to arrive at what we consider to be a pivotal principle of iterated belief contraction.

2 Qualified Insertion – Qualified

We break Qualified Insertion into two sub-principles for convenience in referring to them as follows:[4]

Vacuity Specialised If $y \notin K_x^-$, then $(K_x^-)_y^- = K_x^-$.

Insertion If $y \in K_x^-$, then $(K_x^-)_y^- = K_x^- \cap K_y^-$.

Clearly, Vacuity Specialised is not controversial, and our examination in Example 2 indicates that Insertion is the culprit. Before we claim what is wrong with Insertion, we digress to another well known result [Gärdenfors, 1988](p.65):

Factoring Either

1. $K_{x \wedge y}^- = K_x^-$, or
2. $K_{x \wedge y}^- = K_y^-$, or
3. $K_{x \wedge y}^- = K_x^- \cap K_y^-$.

Given the naming of this principle, it is only appropriate that we refer to the three "values" of $K_{x \wedge y}^-$ as the three **factors**, here as well as in other variants of this principle below. Noting the logical equivalence between y and $(x \vee y) \wedge (x \rightarrow y)$, and making appropriate substitutions in FACTORING we get:

Factoring_2 Either

1. $K_y^- = K_{x \vee y}^-$, or
2. $K_y^- = K_{x \rightarrow y}^-$, or
3. $K_y^- = K_{x \vee y}^- \cap K_{x \rightarrow y}^-$.

From INSERTION and FACTORING_2, in turn, follows the following principle:

Factored Insertion If $y \in K_x^-$, then either

[1] If $\emptyset \vdash y$, then by RECOVERY $(K_x^-)_y^- = K_x^-$. We assume throughout the paper that y is a removable belief, i.e., $\not\vdash y$.

[2] Both Argentina and Brazil had qualified to play in the quarters-finals by the time this paper was being written. Italy played France in the final and won the 2006-world-cup.

[3] Readers acquainted with the semantics involved can verify it by considering the following plausibility rank-order over the worlds: $ab \sqsubset \overline{a}b \sqsubset a\overline{b} \sqsubset \overline{ab}$ where ab is the most preferred world.

[4] Our nomenclature is a bit inaccurate – Rott refers to the consequent of the second item as Insertion.

1. $(K_x^-)_y^- = K_x^- \cap K_{x\vee y}^-$, or
2. $(K_x^-)_y^- = K_x^- \cap K_{x\rightarrow y}^-$ or
3. $(K_x^-)_y^- = K_x^- \cap K_{x\vee y}^- \cap K_{x\rightarrow y}^-$.

It is our view that FACTORED INSERTION together with VACUITY SPECIALISED jointly constitute a central principle for iterated belief contraction. Before arguing our case, we make the following observation:

Observation 1 INSERTION *implies, but is not implied by,* FACTORED INSERTION*, in presence of* FACTORING_2.

The import of this observation may not be immediately obvious. How is it that FACTORED INSERTION fails to imply INSERTION in presence of FACTORING_2? The important point here is which "factor" kicks in when. Let's continue with our FIFA example.

Example 3 *Let us model the scenario in Example 2 as follows. Let the truth vectors* $11, 01, 10, 00$ *represent the four worlds over the atoms* $\{a, b\}$ *(thus* 10 *is the interpretation that assigns* 1 *to* a *and* 0 *to* b*), and* $11 \sqsubset 01 \sqsubset 10 \sqsubset 00$ *where* $\alpha \sqsubset \beta$ *means* α *is strictly preferred to* β*. Under this modelling,* $min([\neg a]) = min(01, 00) = 01$ *whereby* $[K_a^-] = \{11, 01\}$*; thus* $K_a^- = Cn(b)$*. Similarly,* $K_b^- = Cn(a)$*. More importantly,* $[K_{a\rightarrow b}^-] = \{11, 10\}$ *whereby* $K_{a\rightarrow b}^- = Cn(a) = K_b^-$*, and* $[K_{a\vee b}^-] = \{11, 00\}$ *whereby* $K_{a\vee b}^- = Cn(a \leftrightarrow b)$*. Note that Since* $K_b^- = K_{a\rightarrow b}^-$*,* FACTORING_2 *is satisfied.*[5] *Similarly, our intuitive expectation, that* $(K_a^-)_b^- = Cn(\texttt{true})$ *is in accordance with* FACTORED INSERTION *since* $K_a^- \cap K_{a\vee b}^- \cap K_{a\rightarrow b}^- = Cn(b) \cap Cn(a \leftrightarrow b) \cap Cn(a) = Cn(\texttt{true})$*.*[6] *However, since* $a \vee b$ *is in* $K_a^- \cap K_b^-$ *but not in* $Cn(\texttt{true})$*,* INSERTION *is not satisfied.* INSERTION *is not satisfied because, while the "third factor" kicked in to satisfy* FACTORED INSERTION*, the "second factor" kicked in to satisfy* FACTORING_2.

This example highlights an important point, namely, *apart from factorising iterated contraction as in* FACTORED INSERTION*, we should also examine the "trigger conditions" for these factors*. Towards this end, we start with FACTORING, and provide the following result.

Theorem 1 [7] *Given belief set* K*, sentences* x *and* y *and an AGM-rational belief contraction operation* $-$*,*

1. *If* $y \in K_{x\wedge y}^-$ *then* $K_{x\wedge y}^- = K_x^-$
2. *If* $x \in K_{x\wedge y}^-$ *then* $K_{x\wedge y}^- = K_y^-$
3. *If* $x \notin K_{x\wedge y}^-$ *and* $y \notin K_{x\wedge y}^-$ *then* $K_{x\wedge y}^- = K_x^- \cap K_y^-$.

Proof Sketch: The third claim easily follows from the AGM postulates $(7-)$ and $(8-)$. The first two claims are symmetric; so it is sufficient to show only claim (1). We use the standard properties of the epistemic entrenchment relation $\leq$, and the well known construction of contraction via it, namely, $y \in K_x^-$ iff both $y \in K$ and either that $\vdash y$ or $x < x \vee y$. We use $x < y$ and $x \sim y$ as short-hand for $(x \leq y) \wedge (y \nleq x)$ and $(x \leq y) \wedge (y \leq x)$ respectively. We will use $\equiv$ for logical equivalence. From the assumption $y \in K_{x\wedge y}^-$, in the principal case it follows that $x \wedge y < (x \wedge y) \vee y \equiv y$. From this it follows that $(x \wedge y) \sim x < y$. Now consider some arbitrary sentence z. It will be in K_x^- just in case (in the principal case) $x < x \vee z$. Similarly, it will be in $K_{x\wedge y}^-$ just in case (in the principal case) $x \wedge y < (x \wedge y) \vee z$. It is sufficient to show that these two conditions are equivalent. (i) Suppose $x < x \vee z$. We know $x < y$ where by $x < y \vee z$. Since $(x \wedge y) \vee z \equiv (x \vee z) \wedge (y \vee z)$ and $x \sim x \wedge y$ it follows $x \wedge y < (x \wedge y) \vee z$. (ii) Suppose $x \wedge y < (x \wedge y) \vee z \equiv (x \vee z) \wedge (y \vee z)$. It trivially follows that $x < x \vee z$. ■

Theorem 1 tells us how to reduce contraction by a conjunction to known contractions by the individual conjuncts. An immediate corollary of it is applicable to the principle FACTORING_2.

Corollary 1 *Given belief set* K*, sentences* x *and* y *and an AGM-rational belief contraction operation* $-$*,*

1. *If* $x \rightarrow y \in K_y^-$ *then* $K_y^- = K_{x\vee y}^-$
2. *If* $x \vee y \in K_y^-$ *then* $K_y^- = K_{x\rightarrow y}^-$
3. *If both* $x \rightarrow y \notin K_y^-$ *and* $x \vee y \notin K_y^-$ *then* $K_y^- = K_{x\vee y}^- \cap K_{x\rightarrow y}^-$.

Now, our goal was to examine the conditions that trigger different factors of FACTORED INSERTION. If we qualify each item in FACTORED INSERTION by the corresponding condition in Corollary 1 as is:

Naive Factored Insertion Given $y \in K_x^-$,

1. If $x \rightarrow y \in K_y^-$, then $(K_x^-)_y^- = K_x^- \cap K_{x\vee y}^-$
2. If $x \vee y \in K_y^-$, then $(K_x^-)_y^- = K_x^- \cap K_{x\rightarrow y}^-$
3. If both $x \rightarrow y \notin K_y^-$ and $x \vee y \notin K_y^-$, then $(K_x^-)_y^- = K_x^- \cap K_{x\vee y}^- \cap K_{x\rightarrow y}^-$.

then, we will be no better off, since it will lead us to all the problems associated with INSERTION:

Observation 2 NAIVE FACTORED INSERTION *is logically equivalent to* INSERTION.

However, we notice that in Theorem 1, in reducing contraction by a conjunction to contraction by individual conjuncts, the conditions in question perform a "check" on the result of the contraction by the conjunction. In FACTORED INSERTION, the intention is to reduce iterated contraction to one-shot contractions. Accordingly we would expect the appropriate conditions to perform checks on the result of the sequential contraction, not on the result of the one-shot contractions. Accordingly we suggest the following principle instead.

Principled Factored Insertion Given $y \in K_x^-$,

1. If $x \rightarrow y \in (K_x^-)_y^-$, then $(K_x^-)_y^- = K_x^- \cap K_{x\vee y}^-$

[5] That's a bit too quick. We need also to verify for other values of x and y. But that's not the point we are trying to make here.

[6] This intuition is, of course, debatable. One might say that since Brazil has a better track record than Argentina, we should still have $a \rightarrow b \in (K_a^-)_b^-$. That is fine. $K_a^- \cap K_{a\vee b}^- = Cn(b) \cap Cn(a \leftrightarrow b) = Cn(\mathtt{a} \rightarrow \mathtt{b})$, again is in accord with FACTORED INSERTION.

[7] Note that the first two items are equivalent. We nonetheless treat them separately since they eventually lead to distinct principles in PRINCIPLED FACTORED INSERTION below.

2. If $x \vee y \in (K_x^-)_y^-$, then $(K_x^-)_y^- = K_x^- \cap K_{x\to y}^-$
3. If both $x \to y \notin (K_x^-)_y^-$ and $x \vee y \notin (K_x^-)_y^-$, then $(K_x^-)_y^- = K_x^- \cap K_{x\vee y}^- \cap K_{x\to y}^-$

Let us revisit our favourite FIFA example again.

Example 4 *Let the scenario be as in previous examples. We are trying to construct $(K_a^-)_b^-$. We know from Example 1 that both $b \in K_a^-$ and $a \vee b \notin (K_a^-)_b^-$. Since $b \in K_a^-$, we consider* PRINCIPLED FACTORED INSERTION. *Since we already know that $a \vee b \notin (K_a^-)_b^-$, the second case does not apply. So the result of this iterated contraction will crucially depend on whether or not $a \to b$ is in $(K_a^-)_b^-$. If $a \to b \in (K_a^-)_b^-$, then the first case will apply, and we will get $(K_a^-)_b^- = K_a^- \cap K_{a\vee b}^- = Cn(b) \cap Cn(a \leftrightarrow b) = Cn(a \to b)$. Else, the third case will apply, and we will get $K_a^- \cap K_{a\vee b}^- \cap K_{a\to b}^- = Cn(b) \cap Cn(a \leftrightarrow b) \cap Cn(a) = Cn(\texttt{true})$.*

Note that it is the "slight" difference in the conditions attached to different factors that made the difference. We assumed that $a \in K_b^-$, hence $a \vee b \in K_b^-$. If we used NAIVE FACTORED INSERTION, then the the second factor would have kicked in. But since, by assumption, $a \vee b \notin (K_a^-)_b^-$, in case of PRINCIPLED FACTORED INSERTION, it will be a factor other than the second one that kicks in, making all the difference.

3 Some Properties of Principled Iterated Contraction

We have noticed that PRINCIPLED FACTORED INSERTION performs reasonably well in our toy example. In this, short, section we examine some properties that this principle entails.

By a PRINCIPLED ITERATED CONTRACTION OPERATION we understand any AGM-rational contraction operation that also satisfies PRINCIPLED FACTORED INSERTION. It is understood that such contraction operations will automatically satisfy VACUITY SPECIALISED since, being AGM-rational, they satisfy VACUITY. We assume in the following observations that the contraction operation $-$ is a PRINCIPLED ITERATED CONTRACTION OPERATION.

Observation 3 *Given that $y \in K_x^-$, if $x \vdash y$ then $(K_x^-)_y^- = K_x^- \cap K_y^-$.*

The following immediate corollary to Observation 3 indicates that things are as they should be.

Corollary 2 $(K_x^-)_\top^- = K_x^-$

Observation 4 *Given that $y \in K_x^-$, if $\neg x \vdash y$ then $(K_x^-)_y^- = K_x^- \cap K_y^-$.*

From Observation 4 we straightforwardly get the following corollary that lends PRINCIPLED ITERATED CONTRACTION further credence:

Corollary 3 *If $\neg x \in K_x^-$ then $(K_x^-)_{\neg x}^- = K_{\neg x}^-$.*

Observations 3 and 4 are simple, yet quite interesting. On the first count, they identify conditions under which INSERTION holds in general, i.e., $(K_x^-)_y^-$ is identified with $K_x^- \cap K_y^-$. On the second count, these two properties appear to be the analogues of the first two of the four well known postulates advocated by Darwiche and Pearl [Darwiche and Pearl, 1997] for iterated belief revision:

$DP1$ If $y \vdash x$ then $(K_x^*)_y^* = K_y^*$
$DP2$ If $y \vdash \neg x$ then $(K_x^*)_y^* = K_y^*$

Postulates DP1 and DP2 indicate what happens in the case of iterated revision when the two pieces of evidence x and y are logically related. The result is given purely in terms of K_y^*. Similarly, Observations 3 and 4 tell us what happens when the two beliefs to be removed, x and y are logically related. The result is provided purely in terms of K_x^- and K_y^-.

Another interesting fact about these two simple observations is that they record properties of LEXICOGRAPHIC STATE CONTRACTION introduced in [Nayak *et al.*, 2003]. This suggests a deeper connection between PRINCIPLED ITERATED CONTRACTION and LEXICOGRAPHIC STATE CONTRACTION. We take up this issue in the next section.

4 Lexicographic Contraction

LEXICOGRAPHIC CONTRACTION is the contraction-counterpart of of LEXICOGRAPHIC REVISION introduced in [Nayak, 1994; Nayak *et al.*, 2003]. Like most approaches to iterated belief change, this approach has been semantically motivated. In this section we revisit LEXICOGRAPHIC REVISION and LEXICOGRAPHIC CONTRACTION as a prelude to showing that the iterated belief contraction operation generated by Lexicographic Contraction must be a PRINCIPLED ITERATED CONTRACTION operation.

4.1 Lexicographic State Revision

The lexicographic approach to iterated belief revision is captured by a particular account of state revision [Nayak *et al.*, 2003]. The semantics of Lexicographic Revision is given in terms of an evolving belief state, where a belief state is represented as a plausibility ordering over the interpretations generated by the background language.

Definition 1 *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 $\omega \in \Omega$ 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 $\omega' \in \Sigma$, and denote it as $\omega \in min_{\sqsubseteq}(\Sigma)$. We will often drop the parentheses for readability.*

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 $\{\omega \in \Omega \mid \omega \models s \text{ for every } s \in S\}$. For readability, we will abbreviate $[\{s\}]$ to $[s]$. Intuitively, the preorder $\sqsubseteq$ will be the semantic analogue of the revision operation $*$, and will represent the belief state of an agent. 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.

$$K_{\sqsubseteq} = \{x \in \mathcal{L} \mid \omega \models x \text{ for all } \omega \in min_{\sqsubseteq}(\Omega)\}.$$

An inconsistent belief state is represented by an empty relation $\sqsubseteq_\perp$: for every pair $\omega, \omega' \in \Omega, \omega \not\sqsubseteq_\perp \omega'$. Note that this violates connectedness, and hence the plausibility relation $\sqsubseteq$ is, strictly speaking, no longer a total preorder. However, this is a special case, and merits special treatment.

A modified Grove-Construction [Grove, 1988] is used to construct the revision from a given plausibility relation:

Definition 2 **($\sqsubseteq$ to *)**

$$x \in K_e^{*\sqsubseteq} \textit{ iff } \begin{cases} [e] \subseteq [x] & \textit{if } \sqsubseteq \, = \, \sqsubseteq_\bot \\ min_\sqsubseteq[e] \subseteq [x] & \textit{otherwise.} \end{cases}$$

The plausibility ordering (belief state) $\sqsubseteq$, in light of new evidence e, is stipulated to evolve to the new ordering $\sqsubseteq_e^\circledast$ via the use of a state revision operator $\circledast$ as follows.

1. SPECIAL CASES:
   (a) If $[e] = \emptyset$ then, and only then, $\sqsubseteq_e^\circledast = \sqsubseteq_\bot$
   (b) Else, if $\sqsubseteq \, = \, \sqsubseteq_\bot$,
   then $\omega_1 \sqsubseteq_e^\circledast \omega_2$ iff either $\omega_1 \models e$ or $\omega_2 \models \neg e$.
2. GENERAL CASE: Given nonempty prior ($\sqsubseteq \neq \sqsubseteq_\bot$) and satisfiable evidence($[e] \neq \emptyset$),
   (a) If $\omega \models e$ and $\omega' \models e$ then $\omega \sqsubseteq_e^\circledast \omega'$ iff $\omega \sqsubseteq \omega'$
   (b) If $\omega \models \neg e$ and $\omega' \models \neg e$ then $\omega \sqsubseteq_e^\circledast \omega'$ iff $\omega \sqsubseteq \omega'$
   (c) If $\omega \models e$ and $\omega' \models \neg e$ then $\omega \sqsubset_e^\circledast \omega'$

4.2 Lexicographic Contraction

As in the case of lexicographic revision, we take it that the pre-contraction plausibility ordering (belief state) $\sqsubseteq$, after suspension of belief x, evolves to the new ordering $\sqsubseteq_x^\ominus$ via the use of a state contraction operator $\ominus$ as follows.

Definition 3 **(Lexicographic State Contraction $\ominus$)**

1. SPECIAL CASES:
   (a) $(\sqsubseteq_\top)_x^\ominus = \sqsubseteq_\top$ *for all* x.
   (b) $(\sqsubseteq_\bot)_x^\ominus = \sqsubseteq_\top$ *for all* x.

2. GENERAL CASE: *Given non-trivial pre-contraction state* $\sqsubseteq$, $\sqsubseteq_\bot \neq \sqsubseteq \neq \sqsubseteq_\top$,
   (a) If $\omega \models x$ *and* $\omega' \models x$ *then* $\omega \sqsubseteq_x^\ominus \omega'$ *iff* $\omega \sqsubseteq \omega'$
   (b) If $\omega \models \neg x$ *and* $\omega' \models \neg x$ *then* $\omega \sqsubseteq_x^\ominus \omega'$ *iff* $\omega \sqsubseteq \omega'$
   (c) Let $\chi \in \{x, \neg x\}$ *and* $\overline{\chi}$ *be its complement. If* $\omega \models \chi$ *and* $\omega' \models \overline{\chi}$*, then* $\omega \sqsubseteq_x^\ominus \omega'$ *iff for every chain* $\omega_0 \sqsubset \omega_1 \sqsubset \cdots \sqsubset \omega$ *in* $[\chi]$ *of length* n*, there exists a chain* $\omega_0' \sqsubset \omega_1' \sqsubset \cdots \sqsubset \omega'$ *of length* n *in* $[\overline{\chi}]$.

Corresponding conditions 2(a) and 2(b) in LEXICOGRAPHIC REVISION and LEXICOGRAPHIC CONTRACTION are more or less identical: irrespective of whether x is a piece of evidence to be incorporated into the current knowledge corpus (or a belief to be suspended) the prior ordering between any two worlds, ω and ω' that both satisfy x (or both invalidate x), should not be disturbed in the process. These conditions are also accepted in the Darwiche-Pearl approach to belief revision. Condition 2(c) in LEXICOGRAPHIC CONTRACTION states the effect of a contraction by x on two worlds, one of which satisfies x and the other does not: the ranking of all worlds in $[\neg x]$ are simultaneously shifted so that the best (minimal) worlds in $[\neg x]$ attain the same rank as the $\sqsubseteq$-minimal worlds (i.e., the worlds in $[K]$). This is reminiscent of Spohn's account of belief contraction via Ordinal Conditional Functions [Spohn, 1988], and we take it that this is the relational analogue of Spohn's account of iterated belief contraction. This definition of $\ominus$ captures in a simple way a generalised version of the Harper Identity provided in [Nayak *et al.*, 2006]. Figure 1 provides a simple illustration of lexicographic state contraction. The "state space" Ω is partitioned by the pre-contraction plausibility preorder $\sqsubseteq$ into eight "cells", $0, 1, \ldots 7$, the worlds in i being preferable to those in $i+1$, for $0 \leq i \leq 6$, while the worlds in the same cell are equally preferable. These cells are further partitioned by $[x]$ and $[\neg x]$: the cells $X_0, \ldots, X_7$ constituting $[x]$, and the cells $Y_0, \ldots, Y_4$ constituting $[\neg x]$. After the contraction by x, the "best" worlds in $[x]$ as well as the "best" worlds in$[\neg x]$ are accorded the most preferred status in the post-contraction state; the second most preferred worlds consisting of the pre-contraction "next best" worlds from $[x]$ and $[\neg x]$, and so on.

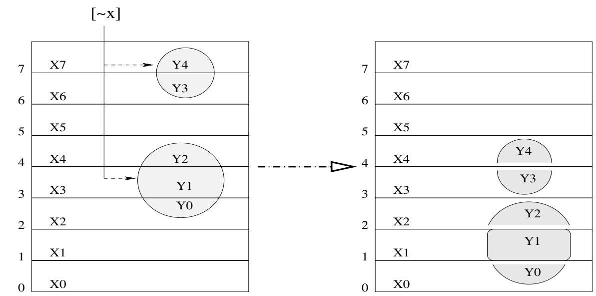

Figure 1: State before and after a contraction by x

Earlier we indicated that there is a close connection between Lexicographic Contraction and Principled Iterated Contraction. For this purpose we need to generate a contraction operation from a given plausibility preorder in order to make this connection. We use a slight variation of the standard transformation.

Definition 4 **($\sqsubseteq$ to $-$)**

$$y \in K_x^{-\sqsubseteq} \textit{ iff } \begin{cases} \Omega \subseteq [y] & \textit{if } \sqsubseteq \, = \, \sqsubseteq_\bot \\ min_\sqsubseteq \Omega \cup min_\sqsubseteq[\neg x] \subseteq [y] & \textit{otherwise.} \end{cases}$$

Our last result (Theorem 2) below shows that if the two sequential contractions in $(K_x^-)_y^-$ are constructed from the belief states $\sqsubseteq$ and $\sqsubseteq_x^\ominus$ respectively, where K is the belief set $K_\sqsubseteq$ associated with the belief state $\sqsubseteq$, then PRINCIPLED FACTORED INSERTION is satisfied.

Theorem 2 *Let* $\sqsubseteq$ *be the original belief state,* K *its associated belief set* $K_\sqsubseteq$*, and* $\ominus$ *the lexicographic state contraction operation. If we define, for every sentence* α *and* β*,*

1. K_α^- *using the contraction operation* $-_\sqsubseteq$ *generated from* $\sqsubseteq$ *via Definition 4, and*

2. $(K_\alpha^-)_\beta^-$*, the result of removing* β *from* K_α^-*, using the contraction operation* $-_{(\sqsubseteq \ominus \alpha)}$ *generated from* $\sqsubseteq_\alpha^\ominus$ *via Definition 4*

then the (dynamic) contraction operation $-$ *thus defined is a* PRINCIPLED ITERATED CONTRACTION OPERATION.

Proof Sketch: Note that, given a preorder $\sqsubseteq$ and sentence x, we have $\sqsubseteq_x^{\ominus}$, denoted $\sqsubseteq'$ below for readability, is a total preorder. Hence the operation $-_{(\sqsubseteq \ominus x)}$, defined as required, is an AGM-rational belief contraction operation. A point to be noted is that the two contraction operations in an expression of the form $(K_x^-)_y^-$ are technically different operations, the first corresponding to $- := -_{\sqsubseteq}$ and the second to $-' := -_{\sqsubseteq'}$, and so it should more accurately be represented as $(K_x^-)_y^{-'}$.[8]

We need to show that the operation pair $(-, -')$ satisfies the property PRINCIPLED FACTORED INSERTION. It is easily shown that in presence of the condition $y \in K_x^-$, the conditions

1. $x \rightarrow y \in (K_x^-)_y^{-'}$ amounts to $min_{\sqsubseteq'}[\neg y] \subseteq [\neg x]$
2. $x \vee y \in (K_x^-)_y^{-'}$ amounts to $min_{\sqsubseteq'}[\neg y] \subseteq [x]$, and
3. the two conditions $x \rightarrow y \notin (K_x^-)_y^{-'}$ and $x \vee y \notin (K_x^-)_y^{-'}$ jointly amount to $\exists\, \omega, \omega' \in min_{\sqsubseteq'}[\neg y]$ s.t. $\omega \in [x]$ and $\omega' \in [\neg x]$.

Note that these three cases are mutually exclusive, and it suffices in these cases to show respectively that

A $min_{\sqsubseteq'}[\neg y] = min_{\sqsubseteq}([\neg x] \cap [\neg y])$

B $min_{\sqsubseteq'}[\neg y] = min_{\sqsubseteq}([x] \cap [\neg y])$

C $min_{\sqsubseteq'}[\neg y] = min_{\sqsubseteq}([\neg x] \cap [\neg y]) \cup min_{\sqsubseteq}([x] \cap [\neg y])$

The information contained in (1), (2) and (3), together with the definition of the operation $\ominus$, in particular, that the relative ordering of worlds inside $[x]$ (respectively $[\neg x]$) is not disturbed, help to carry out the proofs of (A), (B) and (C), correspondingly. For instance, the fact from (1), that $min_{\sqsubseteq'}[\neg y] \subseteq [\neg x]$ and the condition that for $\omega, \omega' \in [\neg x]$, $\omega \sqsubseteq' \omega'$ iff $\omega \sqsubseteq \omega'$ drive the proof of (A).

A point to note that the condition (2c) in Definition 3 is not used in the proof. ■

A natural question to raise at this point is whether the property PRINCIPLED FACTORED INSERTION, together with the AGM postulates of contraction, fully characterises Lexicographic Contraction. It is easily verified that it does not. Given $y \in K_x^-$, if we equate $(K_x^-)_y^-$ with $K_x^- \cap K_{x \vee y}^-$ when $\not\vdash x \vee y$, and with $K_x^- \cap K_{x \vee y}^-$ when $\not\vdash x \rightarrow y$ otherwise, it will satisfy PRINCIPLED FACTORED INSERTION, and yet a *non-Lexicographic* state contraction operation, PRIORITY,[9] will lead to such contraction behaviour.

5 Conclusion

Our aim was to examine iterated belief contraction in a principled way. We started with Rott's QUALIFIED INSERTION that, upon examination, was found to be wanting. We then showed that a judicious combination of QUALIFIED INSERTION with FACTORING leads to what we consider to be a pivotal principle of iterated belief change. We have argued that while Lexicographic State Contraction satisfies it, this principle is too weak to completely characterise a Lexicographic Contraction. We hope to completely characterise Lexicographic Contraction in our future work.

[8] See the dynamics of revision operations in [Nayak *et al.*, 2003] for motivation.

[9] See [Nayak *et al.*, 2006]. This can be pictured as a variation of Figure 1 where in the post-contraction state, $Y_0, \ldots, Y_4$ get the ranks $0, \ldots, 4$; X_0 retains rank 0, and $X_1, \ldots, X_7$ get ranks $5, \ldots, 11$. That is, every non-best world in $[x]$ is *dis*-preferred to every world in $[\neg x]$, subject to other relevant conditions.

References

[Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

[Bochman, 2001] A. Bochman. Contraction of epistemic states: A general theory. In M.-A. Williams and H. Rott, editors, *Frontiers in Belief Revision*, pages 195–220. Kluwer, 2001.

[Boutilier, 1996] C. Boutilier. Iterated revision and minimal revision of conditional beliefs. *Journal of Philosophical Logic*, 25:262–304, 1996.

[Chopra *et al.*, 2002] Samir Chopra, Aditya Ghose, and Thomas Meyer. Iterated revision and the axiom of recovery: A unified treatment via epistemic states. In *Proceedings of the ECAI 2002: 15th European Conference on Artificial Intelligence, ECAI*, pages 541–545, 2002.

[Darwiche and Pearl, 1997] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artifical Intelligence*, 89:1–29, 1997.

[Gärdenfors, 1988] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge Massachusetts, 1988.

[Grove, 1988] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

[Nayak *et al.*, 2003] A. C. Nayak, M. Pagnucco, and P. Peppas. Dynamic Belief Revision Operators. *Artifical Intelligence*, pages 193–228, 2003.

[Nayak *et al.*, 2006] A. C. Nayak, R. Goebel, M. A. Orgun, and T. Pham. Taking Levi Identity Seriously: A plea for iterated belief contraction. In *Proc. of the 1st Int. Conf. on Knowledge Science, Engineering and Management (KSEM 06)*, pages 305–317. LNAI, Springer, 2006.

[Nayak, 1994] A. C. Nayak. Iterated Belief Change based on Epistemic Entrenchment. *Erkenntnis*, 41:353–390, 1994.

[Rott, 2001] Hans Rott. *Change, Choice and Inference: A study of belief revision and nonmonotonic reasoning*. Oxford Science Publications, Clarendon Press, 2001.

[Rott, 2006] Hans Rott. Shifting priorities: Simple representations for twenty-seven iterated theory change operators. *Manuscript*, 2006.

[Spohn, 1988] W. Spohn. Ordinal conditional functions. a dynamic theory of epistemic states. In W. L. Harper and B. Skirms, editors, *Causation in Decisions, Belief Change and Statistics, vol. II*, pages 105–134. Kluwer, 1988.

[Williams, 1994] M-A. Williams. Transmutations of knowledge systems. In *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning*, pages 619–29, 1994.