

Technical Report TR-ARP-2-95

Automated Reasoning Project  
Research School of Information Sciences and Engineering  
and Centre for Information Science Research  
Australian National University

January 11, 1995

## REALISTIC BELIEF REVISION

Greg Restall and John Slaney

**Abstract** In this paper we consider the implications for belief revision of weakening the logic under which belief sets are taken to be closed. A widely held view is that the usual belief revision functions are highly classical, especially in being driven by consistency. We show that, on the contrary, the standard representation theorems still hold for paraconsistent belief revision. Then we give conditions under which consistency is preserved by revisions, and we show that this modelling allows for the *gradual* revision of inconsistency.

# 1 Realistic Logics

Belief Revision is a rich and diverse field. The unit of study is most often a belief set — a set  $K$  of sentences (propositions, whatever) closed under a consequence relation. In Gärdenfors' canonical text [6], and in nearly all other studies of belief revision, the notion of consequence is taken to be *superclassical*. Consequence at least includes classical propositional consequence. This is a theoretical simplification. No-one believes that belief is closed under that sort of consequence. If it were, we would believe all tautologies, and furthermore, we would only have inconsistent beliefs when believing everything. There is something to be said for abandoning the assumption that belief sets are closed under classical consequence, for a more restricted notion of logical closure. One way to do this is to abandon the notion of closure altogether, and to merely work with belief bases. But this is going too far, because it is useful to close beliefs under some kind of consequence. If you believe  $A$  and you believe  $B$ , then you (at least implicitly) believe their conjunction,  $A \wedge B$ . You may retract your belief in  $A \wedge B$  once it becomes explicit, but you still have the belief until you make the contraction. Similarly, if you believe  $A$ , you at least implicitly believe the (inclusive) disjunction  $A \vee B$  of  $A$  with any other proposition  $B$ . This disjunction may not occur to you, but in a good sense it still counts as one of your beliefs. A similar case can be made that you believe  $A$  if and only if you believe  $\sim\sim A$ , that belief in  $\sim(A \vee B)$  amounts to belief in  $\sim A \wedge \sim B$ , and that belief in  $\sim(A \wedge B)$  amounts to belief in  $\sim A \vee \sim B$ .<sup>1</sup> But you cannot make the case as strongly that belief in  $A \wedge \sim A$  entails belief in  $B$  for any  $B$  whatsoever, or that any belief entails belief in  $A \vee \sim A$  for any  $A$  whatsoever. So, we need a notion of logical consequence which will enable us to keep the 'sensible' inferences without their harmful cousins. Thankfully, the notion of *first degree entailment*, first formulated by Smiley around 1960, given a paraconsistent interpretation by Dunn [4], and then recommended for use in this context by Belnap [2] and Slaney [9] among many others, meets our needs. It is a simple modification of classical consequence in which formulae take as truth values *subsets* of  $\{\mathsf{T}, \mathsf{F}\}$ , instead of simply either of  $\mathsf{T}$  or  $\mathsf{F}$  alone. So, in our language  $\mathcal{L}$  built up from atomic formulae in a set  $\text{Atoms}$ , with connectives among  $\wedge$ ,  $\vee$  and  $\sim$ , formulae can be both true and false, or neither true nor false. To be precise, a valuation is a map  $\mathcal{V} : \text{Atoms} \rightarrow \mathcal{P}(\{\mathsf{T}, \mathsf{F}\})$ , extended to the whole language  $\mathcal{L}$  as follows:

- $\mathsf{T} \in \mathcal{V}(A \wedge B)$  iff  $\mathsf{T} \in \mathcal{V}(A)$  and  $\mathsf{T} \in \mathcal{V}(B)$ .
- $\mathsf{F} \in \mathcal{V}(A \wedge B)$  iff  $\mathsf{F} \in \mathcal{V}(A)$  or  $\mathsf{F} \in \mathcal{V}(B)$ .

---

<sup>1</sup>We acknowledge that constructivists will not be happy with all of these inferential moves, and looking at what one can do without these kinds of negation postulates is worthwhile. However, we must leave it for another time.

- $\top \in \mathcal{V}(A \vee B)$  iff  $\top \in \mathcal{V}(A)$  or  $\top \in \mathcal{V}(B)$ .
- $\text{F} \in \mathcal{V}(A \vee B)$  iff  $\text{F} \in \mathcal{V}(A)$  and  $\text{F} \in \mathcal{V}(B)$ .
- $\top \in \mathcal{V}(\sim A)$  iff  $\text{F} \in \mathcal{V}(A)$ .
- $\text{F} \in \mathcal{V}(\sim A)$  iff  $\top \in \mathcal{V}(A)$ .

It is useful to have propositional constants which denote the ‘false only’ truth value  $\{\text{F}\}$  and the ‘true only’ value  $\{\top\}$ .

- $\mathcal{V}(\perp) = \{\text{F}\}$ , and  $\mathcal{V}(\top) = \{\top\}$ .

Once we have our notion of a valuation, we can define first degree entailment. A formula  $A$  is a consequence of a set  $X$  of formulas just when any valuation  $\mathcal{V}$  making every element of  $X$  (at least) true, also makes  $A$  (at least) true. We write this as  $X \vdash A$ , or  $A \in \text{Cn}(X)$ , where  $\text{Cn}(X)$  is the set of all formulae entailed by  $X$ . Note that we have all of the usual properties of conjunction and disjunction (including distribution:  $A \wedge (B \vee C) \vdash (A \wedge B) \vee (A \wedge C)$ ) and the de Morgan laws all hold as entailments. As well, we have  $A \dashv \vdash \sim \sim A$ .<sup>2</sup> However, we do not have  $A \wedge \sim A \vdash B$  (an evaluation could make  $A$  both true and false without making  $B$  true) and neither do we have  $B \vdash A \vee \sim A$  (no matter what  $B$  is,  $A$  could be assigned  $\{\}$ , and so  $A \vee \sim A$  can fail). Note too that *theorems* (formulae  $A$  such that  $\top \vdash A$ ) in our vocabulary must contain either  $\top$  or  $\perp$ .<sup>3</sup>

Of course, first degree entailment is by no means the complete story. We need to consider extensions to deal with conditionals, together with quantifiers, and other deductive machinery. There are indications of how we would like see things developed [8, 9] but here we rest content with the  $\wedge, \vee, \sim, \top, \perp$  fragment of the logic. This is a well-worn resting place, because the logic of first-degree entailment has a proud place in contexts of knowledge representation. It appears in Barwise and Perry’s situation semantics [1]. Our valuations correspond neatly to their *abstract situations*. For them, a situation is a piece of the world which decides (either for or against) some issues, but not necessarily all of them. Abstract situations (used to represent belief states) may also have contradictory information about some issues, reflecting self-contradictory beliefs.

There are also connections with the burgeoning literature on *partial logic*.<sup>4</sup> According to this research community, partial evaluations are useful for representing knowledge states. A partial evaluation is just a partial function from atoms to  $\{\top, \text{F}\}$  instead of a total function, and it is extended to (partially) map formulae to  $\{\top, \text{F}\}$  in the usual way. But this is no different

---

<sup>2</sup>This is shorthand for  $\{A\} \vdash \sim \sim A$  and  $\{\sim \sim A\} \vdash A$ .

<sup>3</sup>To verify this, suppose that  $A$  doesn’t contain  $\top$  or  $\perp$ . Let  $\mathcal{V}$  be an evaluation which sets every atom  $\{\}$ . Then  $A$  must also be evaluated as  $\{\}$ , so it is not a theorem.

<sup>4</sup>For introductory works, consult [3] and [7].

to first degree entailment, if we restrict our evaluations to those in which atomic formulae are consistent. We have a traditional partial evaluation (instead of  $\mathcal{V}(A) = \{\}$ , we say that  $\mathcal{V}(A)$  is undefined). Even granting that knowledge brooks no contradiction (after all, what is known is true) the same cannot be said for belief. So, once we allow that beliefs can be inconsistent, we have a simple generalisation of the machinery of partial logic to first degree entailment.

So, first degree entailment sits happily in the field of knowledge and belief representation. In what follows we will be using this account of entailment to define belief sets, and we will show how the standard representations of belief revision fare in this new wider context, and that the new account of belief sets enable us to do things with belief sets that were heretofore impossible.<sup>5</sup>

## 2 Realistic Contraction and Revision

From now on we assume that we have at hand a notion  $Cn$  of consequence which at least includes first degree entailment. Given that notion of consequence, we define a belief set to be a set  $K$  of formulae such that  $K = Cn(K)$ . Note that any belief set includes  $\top$ , and any other formula  $A$  where  $\top \vdash A$ , but that belief sets need not include  $A \vee \sim A$ , and they can include  $B \wedge \sim B$  without including  $\perp$ . The only belief set which includes  $\perp$  is the trivial belief set  $K_{\perp}$ , which is the set of all propositions in our language.

The simplest operation on belief sets is that of adding another proposition, and closing under logical consequence. We define  $K_A^+$ , the result of adding  $A$  to  $K$  to be  $Cn(K \cup \{A\})$ .

The contraction and revision operations are more interesting. We will start with contraction, the operation of removing a proposition  $A$  from a belief set. Gärdenfors' eight original postulates for a contraction of a belief set  $K$  are as follows.

- (K<sup>-</sup>1)  $K_A^-$  is closed.
- (K<sup>-</sup>2)  $K_A^- \subseteq K$ .
- (K<sup>-</sup>3) If  $A \notin K$  then  $K_A^- = K$ .
- (K<sup>-</sup>4) If  $\top \not\vdash A$  then  $A \notin K_A^-$ .
- (K<sup>-</sup>5) If  $A \in K$  then  $K \subseteq (K_A^-)_A^+$ .
- (K<sup>-</sup>6) If  $A \dashv\vdash B$  then  $K_A^- = K_B^-$ .

---

<sup>5</sup>Fuhrmann [5] seems to be the only instance in the rather large belief revision literature in which the proposal to allow non-trivial inconsistent belief sets is taken seriously and developed a little way. He does not consider the obvious dual problem of belief sets including all logical theorems.

$$(K^-7) K_A^- \cap K_B^- \subseteq K_{A \wedge B}^-.$$

$$(K^-8) \text{ If } A \notin K_{A \wedge B}^- \text{ then } K_{A \wedge B}^- \subseteq K_A^-.$$

The most problematic of these eight requirements is the condition (K<sup>-</sup>5) of *recovery*. The idea behind recovery is that by removing  $A$  from  $K$  you only make a minimal change to  $K$ , So minimal that by adding  $A$  to the result, and closing under consequence you get all of  $K$  back. This postulate has been widely criticised, because many intuitive operations of contraction simply do not validate it. This concurs with our approach, because the representations considered in the next section will not validate recovery. For our purposes, a ‘contraction operator’ will be a function  $K^-$  satisfying the conditions  $K^- \{1, 2, 3, 4, 6, 7, 8\}$ .

Given a contraction operator  $K^-$ , we can define a revision operator  $K^*$  in the usual way. To revise your belief set  $K$  by  $A$ , you first retract your original belief that  $\sim A$  (if you have one) and then add  $A$ , closing under consequence. So, you can define  $K_A^*$  to be  $(K_{\sim A}^-)_A^+$ . This is the Levi Identity. Then we would get a revision operator  $K^*$  satisfying the following postulates.

$$(K^*1) K_A^* \text{ is closed.}$$

$$(K^*2) A \in K_A^*.$$

$$(K^*3) K_A^* \subseteq K_A^+.$$

$$(K^*4) \text{ If } \sim A \notin K \text{ then } K_A^+ \subseteq K_A^*.$$

$$(K^*5) K_A^* = K_{\perp} \text{ if and only if } A \vdash \perp$$

$$(K^*6) \text{ If } A \dashv \vdash B \text{ then } K_A^* = K_B^*.$$

$$(K^*7) K_{A \wedge B}^* \subseteq (K_A^*)_B^+.$$

$$(K^*8) \text{ If } \sim B \notin K_A^* \text{ then } (K_A^*)_B^+ \subseteq K_{A \wedge B}^*.$$

We leave the verification of these postulates to the reader. The reasoning is not any more difficult than the classical case.

Traditionally, we can also define contraction functions from revision functions, by way of the Harper Identity:  $K_A^- = K \cap K_{\sim A}^*$ . However, in our non-classical environment, this definition does not always define a contraction function, for the following reason. We need not have  $A \notin K_A^-$ , because  $A$  might be both in  $K$ , and in  $K_{\sim A}^*$ . The new theory  $K_{\sim A}^*$  may be inconsistent about  $A$ . This does seem odd, because we would not expect  $K_{\sim A}^*$  to be inconsistent about  $A$ , because we have asked it to revise with respect to  $\sim A$ . It ought remove  $A$  and then add  $\sim A$ . However, it could well be that adding  $\sim A$  might bring with it  $A$ .  $\sim A$  might entail its own negation (as it would if it were of the form  $B \wedge \sim B$ ). In that case, adding  $\sim A$  would bring with it  $A$ , no matter how much we would like to avoid this. Classically, this

means that  $\sim A \vdash \perp$ , and hence,  $\top \vdash A$ , so  $A \in K_A^-$  is no problem, since you cannot contract away theorems. However, in our context we need not have  $\sim A \vdash \perp$ . In other words,  $A$  could entail its own negation without  $A$  being *trivialising*. As a result, the Harper identity fails. This will become important when we discuss Grove’s system of spheres.

### 3 Representations

That is enough of what does not work in the nonclassical environment. There is a lot of good news about what *does* work. In this section we will see that each of the standard representation results, of epistemic entrenchment, transitively relational partial meet contraction, and Grove’s systems of spheres each generalise to our setting, and they do not commit us to the counterintuitive recovery postulate.

#### 3.1 Epistemic Entrenchment

First we can consider entrenchment. Here the guiding idea is that we rank our beliefs according to some notion of ‘entrenchment’. The more deeply held a belief, less likely we are to give it up. So, we assume we have at hand a notion  $\sqsubseteq$  of epistemic entrenchment. (We  $A \sqsubseteq B$  to mean  $A$  is more deeply entrenched than  $B$ , in contrast to Gärdenfors [6], and others, who write it as  $B \leq A$ . We prefer this notation, because it makes it clear that  $A$  is *deeper* than  $B$ .) We follow Gärdenfors in positing five postulates of entrenchment.

(EE1)  $\sqsubseteq$  is transitive.

(EE2) If  $A \vdash B$  then  $B \sqsubseteq A$ .

(EE3) Either  $A \wedge B \sqsubseteq A$  or  $A \wedge B \sqsubseteq B$ . (As a corollary, it follows that  $\sqsubseteq$  is a total order.)

(EE4) If  $K \neq K_\perp$  then  $B \sqsubseteq A$  for all  $B$  if and only if  $A \notin K$ . (In other words, all elements not in  $K$  are equally at the top of the entrenchment order.)

(EE5) If  $A \sqsubseteq B$  for all  $B$  then  $\top \vdash A$ . (So, theorems alone are the most deeply entrenched propositions.)

Note that the postulates are more plausible when interpreted in our context than classically, for now the most deeply entrenched propositions are only  $\top$  and what it entails, instead of all of the theorems of classical logic, of arbitrary complexity. Similarly, the condition (EE2) only relates entrenchments to entailments in our base logic, which can be a lot weaker than classical logic. Better still, we have the following results.

- If  $\sqsubseteq$  is an entrenchment ordering, then defining

$$K_A^- = \begin{cases} K \cap \{B : A \vee B \sqsubset A\} & \text{if } \top \not\vdash A \\ K & \text{otherwise} \end{cases}$$

makes  $K^-$  a contraction function. (Where we define  $B \sqsubset A$  to mean  $A \not\sqsubseteq B$ ). Given  $\sqsubseteq$ , the resulting  $K^-$  is said to be the contraction function determined by  $\sqsubseteq$ .

- Conversely, if  $K^-$  satisfies each condition, then we can define an entrenchment ordering  $\sqsubseteq$  by setting  $B \sqsubseteq A$  if and only if either  $A \notin K_{A \wedge B}^-$  or  $\vdash A \wedge B$ . Given  $K^-$ , the resulting  $\sqsubseteq$  is said to be the entrenchment ordering determined by  $K^-$ .

For the first result, retracting  $A$  from  $K$  is given by taking away those  $B$ s where  $A \sqsubseteq A \vee B$ . That is, the  $B$ s are not as deeply entrenched in  $K$  as  $A$  is. Conversely, we take  $B$  to be as deeply entrenched as  $A$  if either  $A$  and  $B$  are both theorems, or when you contract  $A \wedge B$  from  $K$ , you take  $A$  out of the result. Clearly,  $A$  could not be more deeply entrenched than  $B$  in this context, for then you would have removed  $B$  from  $K$  and kept  $A$ .

As before, the proofs of these results are no different to the classical proofs, so we leave them to the reader.

It is illuminating to consider the fate of the recovery postulate, given this account of contraction. Classically, we can reason as follows. Suppose  $B \notin K_A^-$ , but  $B \in K$ . We wish to show that  $B \in (K_A^-)_A^+$ . If we have  $B \notin K_A^-$  and  $B \in K$ , then we must have  $A \vee B \not\sqsubseteq A$ . In the classical context, we have  $B \in (K_A^-)_A^+$  only when  $A \supset B \in K_A^-$ . That is,  $\sim A \vee B \in K_A^-$ . But this happens when  $\sim A \vee B \in K$  (and that's taken care of, since  $B \in K$ ) and  $A \vee (\sim A \vee B) \sqsubset A$ . But this is immediate, because, in the classical context,  $A \vee (\sim A \vee B)$  is equivalent to  $\top$ , and  $\top \sqsubset A$ , because  $A$  is less deeply entrenched than the theorems, since it was successfully retracted from  $K$ .

This line of reasoning is *very* sensitive to the choice of logic to be used. Firstly, we need  $A \supset B \in K_A^-$  to give  $B \in (K_A^-)_A^+$ . In other words, we need  $A, A \supset B \vdash B$ , or equivalently,  $A, \sim A \vee B \vdash B$ . Once we admit inconsistent states, this is not plausible, because a state could be inconsistent about  $A$ , without  $B$  following. So, if inconsistency is possible, the deduction stops here. If partiality is possible, the deduction breaks down as well. In that case, we cannot say that  $A \vee (\sim A \vee B)$  is equivalent to  $\top$ , because the excluded middle  $A \vee \sim A$  might fail, for lack of information about  $A$ . In either case, the deduction from an entrenchment ordering to the recovery postulate fails.

It is worth pausing for a moment to consider what happens in the presence of a stronger conditional operator, like the  $\rightarrow$  of a relevant logic, or of the logic **BN** [9]. In this case we cannot make the move from  $B \in K$  to  $A \rightarrow B \in K$ , because in neither calculus do you have  $B \vdash A \rightarrow B$ . Neither

will you have  $A \vee (A \rightarrow B)$  as a theorem, so the consequence of recovery need not follow, even in the presence of a stronger conditional.

It is also helpful to consider what contraction operators can look like. Here is a particularly simple example. Take  $K = Cn\{p, q\}$ , and define  $\sqsubseteq$  as follows.  $A \sqsubseteq B$  iff either  $\top \vdash A$ , or  $A \in K$  and  $B \notin K$ . This is the coarsest entrenchment ordering, which takes theorems to be most deeply entrenched, other elements of  $K$  next, and non-elements of  $K$  last. It is simple to check that it is an entrenchment ordering. Consider  $K_p^-$ . It is  $K \cap \{B : p \vee B \sqsubseteq p\}$ . But here,  $p \vee B \sqsubseteq p$  if and only if  $\top \vdash p \vee B$ , and this obtains only when  $\top \vdash B$ . As a result  $K_p^- = Cn\{\top\}$ , and  $(K_p^-)_p^+ = Cn\{p\} \neq K$ . Our entrenchment ordering can take  $p$  and  $q$  to be of equal entrenchment, so that if one goes, so does the other — without leaving behind anything that would connect the two. As a result, the recovery postulate fails in general. On the other hand, given a conditional ( $A \supset B$  will do if we are avoiding inconsistency, otherwise we need a more vertebrate conditional) we can leave behind a ‘connection’ of the form  $p \rightarrow q$ , so any future addition of  $p$  will bring with it  $q$ . Given a non-classical notion of logical closure, we can keep the association between entrenchment of propositions and revision, without the untoward consequence of the recovery postulate. The recovery postulate dictates that the ‘connections’ must always be there. In our context we have the freedom to postulate connections between propositions or not, as the evidence warrants.

### 3.2 Partial Meet Contractions

Another popular representation of contraction functions is in terms of ‘partial meets’. In the classical context, we define  $K \perp A$  to be the class of all *maximal belief sets* which are subsets of  $K$  and which do not contain  $A$ . Let  $K \perp = \bigcup_A K \perp A$  collect all of these maximal belief sets together. A ‘maxichoice’ contraction function takes  $K_A^-$  to be a member of  $K \perp A$ . But the resulting belief set is, in general, ‘too big’. So, we can try taking  $K_A^-$  to be  $\bigcap \gamma(K \perp A)$ , the intersection (meet) of an appropriately chosen collection of elements of  $K \perp A$ . So, we take some favoured elements of  $K \perp A$ , and let  $K_A^-$  be what is contained in all of them. The result will be a contraction function which satisfies conditions (K<sup>-</sup>1) to (K<sup>-</sup>6), assuming that the underlying logic is at least classical. In our context, it is trivial to show that the contraction function will satisfy each of these postulates apart from recovery, (K<sup>-</sup>5).

To get the seventh and eighth postulates, we need to do a little more work. We need to assume that the choice function  $\gamma$  is not totally arbitrary, but rather, given by a global preference relation on  $K \perp$ . Given a relation  $\leq$  on  $K \perp$ , we can define  $K_A^-$  to be the intersection of all of the most preferred elements of  $K \perp A$  (if there are any) and  $K$  otherwise. In other words, we take  $K_A^-$  to be  $\bigcap \{K' : K' \in K \perp A \text{ and } K'' \leq K' \text{ for all } K'' \in$



$K \perp A$  if  $\top \not\vdash A$ , and  $K$  otherwise. It is simple to show that the operation  $K^-$ , so defined satisfies (K<sup>-</sup>7). If we take  $\leq$  to be a transitive relation, we also have (K<sup>-</sup>8). All this is standard, and the usual proofs go through in our context as well.

However, given the absence of recovery, and given the fact that theories need not contain all instances of excluded middle, not every contraction function appears in this way. Consider our simple example where  $K = Cn\{p, q\}$ . There is a contraction function for which  $K_p^- = Cn\{\top\}$ . Consider  $K \perp p$ . Every element of  $K \perp p$  contains  $p \vee \sim p$ , simply because if a theory doesn't, it is not a *maximal* subtheory of  $K$  not entailing  $p$ . For we can easily extend it by making  $p$  false. So, every partial meet contraction function defined in this way gives us  $p \vee \sim p \in K_p^-$ . But not every contraction function need be like this. We may be completely ignorant about  $p$  upon retracting it from  $K$ .

There is a simple fix to this problem, by allowing  $K \perp A$  to include *all* belief sets weaker than  $K$  which do not entail  $A$ . The resulting partial meet contraction function defined in the same way will satisfy each of our postulates, and furthermore, given a contraction function  $K^-$ , we may define a rank ordering on  $K \perp$ , as follows. Set  $K'' \leq K'$  if and only if  $K'' = K' = K$ , or  $K'' \in K \perp A$  for some  $A \in K$ ,  $K' \in K \perp A$  and  $K_A^- \subseteq K'$  for some  $A \in K$  and finally, for all  $A$ , if  $K', K'' \in K \perp A$  and  $K_A^- \subseteq K''$  then  $K_A^- \subseteq K'$  too.

It is not difficult to show that this gives a rank ordering on belief sets which generates the original contraction function. The checking is tedious, and we leave the details to interested readers.

### 3.3 Spheres

Finally, we will sketch the way that Grove's system of spheres can be adapted to our setting. We assume acquaintance with the classical context. Recall that in the traditional setting, a belief set  $K$  is represented by the set  $[K]$  of possible worlds at which it is true. Contractions, are generated by a 'system of spheres', a series of sets of possible worlds, each containing those before it in the series, starting with  $[K]$ , and which taken together cover every possible world. These sets give a measure of closeness to  $[K]$ . A world  $w_1$  is closer than  $w_2$  to  $[K]$  when it is in some set  $S$  in the series which does not contain  $w_2$ . Then, to find  $K_A^-$ , consider the set  $S$  of closest worlds to  $[K]$  at which  $\sim A$  is true.  $K_A^-$  is then the set of all propositions true at all worlds in  $[K] \cup S$ . So,  $K_A^-$  will not make  $A$  true,<sup>6</sup> as  $\sim A$  is true in the worlds in  $S$ . Similarly  $K_A^- \subseteq K$ , as every ' $K$ -world' is in the set  $[K] \cup S$ . It is not difficult to show that this construction gives us a contraction function.<sup>7</sup>

<sup>6</sup>Unless  $A$  is true in all worlds, and so,  $S$  is empty.

<sup>7</sup>There are some fine points we are stepping over lightly here. It takes some work to ensure that there is a closest set of worlds, as the definition requires.

The same kind of construction works with weaker logics too. However, we need make two important changes. The first is the definition of ‘worlds’. Instead of complete consistent sets of propositions, we must take worlds to be *prime theories*. That is, we consider those sets  $P$  of propositions, closed under  $Cn$ , such that if  $A \vee B \in P$ , then  $A \in P$  or  $B \in P$ . These correspond to our valuations  $\mathcal{V}$ . For any valuation  $\mathcal{V}$ , the set of all propositions made (at least) true by  $\mathcal{V}$  consists of a prime theory. These ‘worlds’ allow inconsistencies and incompleteness.

Another change is required before we have a contraction function from a system of spheres on the set of prime theories. The classical definition requires us to find  $K_A^-$  by adding to  $[K]$  the closest worlds at which  $\sim A$  is true. But this is no good in our context, for two reasons. Firstly, it could be that the closest worlds at which  $\sim A$  is true also feature  $A$ . As a result,  $K_A^-$  as defined could still contain  $A$ . Another problem is the fact that if  $A$  is in  $K$ , then so is  $A \vee \sim A$ . Similarly,  $A \vee \sim A$  is true at the worlds closest to  $K$  at which  $\sim A$  is true. So,  $A \vee \sim A \in K_A^-$  as we have defined it. This is not in general true for contraction functions as we have defined them, so, this definition is not as broad as it could be. Note that these problems are just like those that beset the Harper identity, which we saw earlier.

There is a simple fix for these problems, and that is to alter the definitions as follows. Define  $K_A^-$  by adding to  $[K]$  the closest worlds at which  $A$  fails. Then if there is any world at all at which  $A$  fails (that is, if  $\top \not\vdash A$ ) then  $A \notin K_A^-$ . And the closest worlds at which  $A$  fails could well also reject  $A \vee \sim A$ , so we have no requirement that these excluded middles remain. As a result, we have a definition in terms of spheres of worlds which does justice to the contraction postulates (as is easily checked).

## 4 Consistency and Inconsistency

We have seen that expanding our horizons by considering belief sets which are not closed under classical consequence does not take away any of the traditional representation results. Instead, it opens up new possibilities and it eliminates some counterintuitive properties of the traditional accounts of belief revision. In this final section we will consider how belief sets can remain consistent, or gradually extract themselves from inconsistency, which is much more like what we do in practice. This fulfils Fuhmann’s desire when he writes:

In the face of inconsistent theories we should want two things:

- (a) *localise inconsistencies* — an inconsistent theory should not be rendered totally corrupt just because some inconsistency has crept into the theory; and

- (b) *locally restore consistency* — we should be able to resolve one inconsistency at a time by contracting an inconsistent theory such that other inconsistencies, which we cannot yet resolve, may be carried over into the contracted theory. [5] pages 186 and 187

Desideratum (a) is satisfied because the background logic is paraconsistent. Our belief sets can contain inconsistencies without being trivial. In the rest of the paper we shall consider desideratum (b). To show that we can keep parts of a theory consistent while restoring consistency to the inconsistent part, we will use the notion of a *vocabulary*. A vocabulary  $V$  is a set of atoms in our language. The restriction  $K \upharpoonright V$  is the set of formulae in  $K$  built up from atoms in  $V$ ,  $\top$  and  $\perp$ . It is simple to show that if no atom  $p$  in  $V$  is inconsistent in  $K$  (that is,  $p \wedge \sim p \in K$ ) then  $K \upharpoonright V$  is consistent. This is how we will keep track of a part of a belief set  $K$  being consistent — it will have some vocabulary in which it is consistent. The inconsistency is restricted to a different ‘subject matter’. Now suppose  $K$  is inconsistent, but that  $K \upharpoonright V$  is consistent, for some suitable vocabulary  $V$ . Then we wish to revise  $K$ , to remove an inconsistency, keeping  $K \upharpoonright V$  consistent. But this is simple. Consider contractions from the perspective of epistemic entrenchment orderings. We say that an entrenchment ordering  $\subseteq$  *keeps  $V$  fixed* if the  $\subseteq$ -deepest formulae are the theorems, and the elements of  $Cn(K \upharpoonright V)$  are deeper than any other formulae (other than theorems) in  $K$ . Then, if  $A \in K$ , but  $A \notin K \upharpoonright V$ , then  $K \upharpoonright V \subseteq K \upharpoonright V$ , as is simple to check. This keeps the beliefs in the  $V$ -vocabulary untouched, while retracting the inconsistencies.

It follows that we can contract away contradictions step by step, while keeping our ‘safe’ information immune from change. Note too that there is no compulsion that  $K \upharpoonright V$  be consistent throughout this process. We could have a some vocabulary such that beliefs in it are to be kept fixed during a revision — this may include inconsistencies which are just too difficult to deal with at the current time. This too is allowed on our account of revision.

A simple example should suffice to illustrate the process. Suppose we have a theory like  $Cn\{p, q, r \wedge \sim r, s \wedge \sim s\}$ . Given an appropriate entrenchment ordering, we can contract on  $s \wedge \sim s$ , say, perhaps keeping neither  $s$  nor  $\sim s$ , nor  $s \vee \sim s$ , but keeping everything else, because we have kept the beliefs in the vocabulary  $\{p, q, r\}$  more deeply entrenched (even  $r \wedge \sim r$ , which may be surprising — we may hold to  $r \wedge \sim r$  more strongly than to  $s \vee \sim s$ ). As a result, our belief set is now  $Cn\{p, q, r \wedge \sim r\}$ . Then we can contract on  $r \wedge \sim r$ , keeping, say  $r$ , because it is more deeply entrenched than  $\sim r$ . During this contraction we can keep the vocabulary  $\{p, q\}$  constant, so we end up with  $Cn\{p, q, r\}$ , a consistent belief set.

## 5 Conclusion

We have seen that first-degree entailment gives us the techniques to deal with inconsistent belief sets without requiring that all inconsistency be dealt with at once. This has not been at the expense of the formal representation results of belief revision. The theory of belief revision which allows for inconsistency has all of the formal beauty, and the intuitive grounding of the model theory of the classical account of belief revision. The only things it lacks are the oversimplifications and questionable consequences of the classical theory. And that is not a bad thing at all.

## References

- [1] Jon Barwise and John Perry. *Situations and Attitudes*. MIT Press, Bradford Books, 1983.
- [2] Nuel D. Belnap. A useful four-valued logic. In J. Michael Dunn and George Epstein, editors, *Modern Uses of Multiple-Valued Logics*, pages 8–37. D. Reidel, Dordrecht, 1977.
- [3] S. Blamey. Partial logic. In Gabbay and Günthner, editors, *Handbook of Philosophical Logic*, volume 3. Reidel, Dordrecht, 1986.
- [4] J. Michael Dunn. Intuitive semantics for first-degree entailments and “coupled trees”. *Philosophical Studies*, 29:149–168, 1976.
- [5] André Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20:175–203, 1991.
- [6] Peter Gärdenfors. *Knowledge in Flux*. MIT Press, Bradford Books, 1988.
- [7] J. Halpern and Y. Moses. Towards a theory of knowledge and ignorance. In Kr. Apt, editor, *Logics and Models of Concurrent Systems*. Springer-Verlag, Berlin, 1985.
- [8] Greg Restall. Four valued semantics for relevant logics (and some of their rivals). *Journal of Philosophical Logic*, 1995. to appear.
- [9] John K. Slaney. The implications of paraconsistency. In *Proceedings of IJCAI-91*, pages 1052–1057. 1991.