Belief revision and ordered theory presentations*

Mark Ryan†

Abstract

The standard theory of belief revision—known as the AGM theory after its authors, C. Alchourrón, P. Gärdenfors and D. Makinson [3]—suffers from at least two disadvantages. One is that it represents belief states as *infinite* objects, namely deductively closed sets of sentences. Another is that existing belief revision models make *too strong assumptions* about what information is available to guide revisions. A consequence of this is that repeated revisions are impossible.

In this paper, we show how to use *ordered theory presentations* to represent belief states[1]. Ordered theory presentations are theory presentations equipped with a partial order. In this paper we are only concerned with a special case, in which the sets are finite and the order is total, or linear. We define a revision operator, which is shown to satisfy some, but not all, of the AGM postulates. No information other than that encoded in the OTP is needed to effect the revision; this makes repeated revision easy.

1 Introduction

The central question in belief revision is the following: given a *belief state* and a *sentence*, how should one obtain a new belief state in which the sentence is *true*, but which preserves as much of the old belief state as possible? In other words, one wants a function

$$* : \text{belief states} \times \text{sentences} \rightarrow \text{belief states}$$

such that

1. ϕ is true in $? * \phi$; that is, the revision has been *effective*; and
2. given this constraint, $? * \phi$ contains 'as much' of ? as is consistent; that is, old beliefs persist through revisions if they can.

The case of interest, of course, is that in which $\neg\phi$ is true in ?, so that the revision is more than just *refinement*, or the addition of compatible information. We also hope that

3. $? * \phi$ does not contain any extraneous information which was present in neither ? nor ϕ.

In the above requirements, some things are easy to formulate and some are not. We assume that any satisfactory representation of belief states comes with a function $|\cdot|$ which takes a belief state and returns the set of sentences true in it. $|?|$ is called the *extension* of ?. But formalising the 'as much' requirement and the requirement of no extraneous information (numbers 2 and 3) is not so easy, and is the subject of this paper.

*To appear in the proceedings of the Eighth Amsterdam Colloquium, 1991

†Dept. of Computing, Imperial College, 180 Queens Gate, London SW7 2BZ, Great Britain. Phone: +44 71 589 5111 ext. 5074. Email: `mdr@doc.ic.ac.uk`.

[1] Ordered theory presentations have already been described in a previous paper [12]. In that paper, they were called 'structured theories', but since they are in fact a way of using an ordering on a set of sentences to present a theory, their name has been changed to *ordered theory presentations*, or *ordered presentations*, or even OTPs, for short. Familiarity with [12] is not required in order to understand this paper.

Belief revision has obvious applications in artificial intelligence (eg. robotics), computer science (eg. deductive databases), the philosophy of science, social theory and so on. It also has applications beyond the idea of revising 'beliefs'. For example, in specification theory and in AI, there is the well-known frame problem to do with the semantics of actions. Given (the representation of) a state of a system and the post-condition of an action performed in the state, what is the state which results from performing the action? The same requirements on the revision function apply here too: the post-condition should be true in the resulting state, which (given this constraint) should preserve as much of the original state as possible.

We begin the next section by describing the standard theory of belief revision, known as the AGM theory. The AGM theory suffers from several disadvantages. One is that it represents belief states as *infinite* objects, namely deductively closed sets of sentences. Another is that existing belief revision models make *too strong assumptions* about what information is available to guide revisions. A consequence of this is that repeated revisions are impossible.

We propose representing belief states by linear ordered theory presentations instead of by theories. Definitions and examples will be given. The result is a finite representation of belief states. The revision operator is shown to satisfy some, but not all, of the AGM postulates. The counterexamples for the AGM postulates which fail are motivated. The important point is that no information other than that encoded in the OTP is needed to effect the revision; this makes repeated revision easy.

The paper is organised as follows. The AGM theory is described in the next section. Section 3 describes the properties which a theory of belief revision should have. The notion of *ordered theory presentations* (OTPs) is introcuced in section 4, which is the main section. Sections 5 and 6 discuss the properties of the revision operator on OTPs; and further examples are given in section 7.

2 The AGM theory

The standard theory of belief revision—known as the AGM theory after its authors, C. Alchourrón, P. Gärdenfors and D. Makinson [3]—models belief states as *deductively-closed sets of sentences.* More recent developments of the AGM theory are described in S. O. Hansson's thesis [4]. It describes a small set of postulates which any belief revision operator should satisfy (see below). If K is a belief state and ϕ a formula, then $K * \phi$ is a belief state, the result of revising K with ϕ. As already noted, the case of interest is when $\neg\phi \in K$, that is, when the revising sentence conflicts with the current belief state. We suppose we are utilising classical logic with the usual connectives, and the usual entailment relation $\models$.

Notation 1

1. Let Φ be a set of sentences. $\mathsf{Cn}(\Phi) = \{\phi \mid \Phi \models \phi\}$.
2. L is the set of all sentences in the language.
3. $K + \phi = \mathsf{Cn}(K \cup \{\phi\})$.

The AGM postulates are the following:

K1 $K * \phi$ is a deductively-closed theory;

K2 $\phi \in K * \phi$;

K3 $K * \phi \subseteq K + \phi$;

K4 If $\neg\phi \notin K$ then $K + \phi \subseteq K * \phi$;

K5 $K * \phi = L$ implies $\phi = \bot$;

K6 If $\models \phi \leftrightarrow \psi$ then $K * \phi = K * \psi$;

K7 $K * (\phi \wedge \psi) \subseteq K * \phi + \psi$;

K8 If $\neg\psi \notin K * \phi$ then $K * \phi + \psi \subseteq K * (\phi \wedge \psi)$.

K1 says that $K * \phi$ should be a belief state. K2 says that the revision should be successful, i.e. the resulting theory should at least contain ϕ. The third axiom says that $K * \phi$ should have no more than what we would get by just adding ϕ set-theoretically and closing under entailment. Of course, if ϕ is inconsistent with K then adding it in that way would yield the whole of L (the theory with every sentence in it). K4 asserts that if ϕ *is* consistent with K then we get precisely the result of adding it set-theoretically. We should point out that this is one of the (two) axioms with which we take issue in §6. K5 says that the revision yields the contradictory theory L only if ϕ is inconsistent. This is not just that ϕ is inconsistent *with* K, but is inconsistent on its own. The converse, that revising with an inconsistent sentence yields the inconsistent theory, is guaranteed by K2. K6 says simply that revising with logical equivalents yields the same theory.

K7 and K8 are more complicated, approximating what happens with repeated revisions. They are analogues of K3 and K4.

Note that K7 and K8 do not contain expressions like $K * \phi * \psi$, and therefore do not constrain repeated revision in any explicit way. The only constraints on repeated revision are those inherited from the more general case of revision which K1–8 describe.

We submit that the AGM axioms to be neither *sound* nor *complete* with respect to intuitively rational belief revision. Of course such a statement is necessarily imprecise, because 'intuitively rational' belief revision is not amenable to mathematical description. The argument to show lack of soundness consists of 'counterexamples' to K4 and K8 later in the paper. (Again the scare quotes show that these are not counterexamples to any fully spelled-out conjecture.) The argument against completeness is the following proposition, which shows that K1–8 admit revision functions which have no element of the 'persistence' requirement (number 2 above).

Proposition 2 The revision function

$$K * \phi = \begin{cases} K + \phi & \text{if } \neg\phi \notin K \\ \mathsf{Cn}\{\phi\} & \text{otherwise} \end{cases}$$

satisfies axioms K1–8.

Proof K1–4 and K6 are immediate.

K5 Suppose $K * \phi = L$. By K3, $K + \phi = L$, so $\neg\phi \in K$. Therefore, $K * \phi = \mathsf{Cn}\{\phi\}$, so $\phi = \bot$.

K7 Suppose $\neg(\phi\wedge\psi) \in K$. Then $K*(\phi\wedge\psi) = \mathsf{Cn}(\phi\wedge\psi)$. If $\neg\phi \in K$ then $(K*\phi)+\psi = \mathsf{Cn}(\phi\wedge\psi)$; otherwise it is $(K + \phi) + \psi$, which contains $\mathsf{Cn}(\phi \wedge \psi)$.

Otherwise $\neg(\phi\wedge\psi) \notin K$, and so $\neg\phi \notin K$. $K*(\phi\wedge\psi) = K+(\phi\wedge\psi) = (K+\phi)+\psi = (K*\phi)+\psi$.

K8 Suppose $\neg\psi \notin K * \phi$. Suppose also that $\neg\phi \notin K$. Then $K * \phi = K + \phi$. Therefore, $K*\phi+\psi = K+\phi+\psi$. Also, since $\neg\psi \notin K*\phi$ and $K*\phi = K+\phi$, we have that $\neg\psi \notin K+\phi$ and therefore $\neg(\phi \wedge \psi) \notin K$. Therefore, $K * (\phi \wedge \psi) = K + (\phi \wedge \psi) = K + \phi + \psi$, as required.

Now suppose $\neg\phi \in K$. Then $\neg(\phi\wedge\psi) \in K$, so $K*\phi+\psi = K+\phi+\psi = K+(\phi\wedge\psi) = K*(\phi\wedge\psi)$. ◇

Of course there are more interesting functions satisfying the axioms. The following two are the most important in the AGM literature: *partial meet* revision; and revision by *epistemic entrenchment*.

2.1 Selection functions

Suppose K is a belief state and ϕ is a sentence other than $\bot$. Let

$$K|_\phi = \text{ the } \subseteq\text{-maximal elements of } \{K' \subseteq K \mid \neg\phi \notin K\},$$

that is, the set of maximal subsets of K which are consistent with ϕ. $K|_\phi$ may be pronounced 'K without ϕ'. The operation of partial meet revision assumes a selection function S_K which selects *some* of these subsets. Then revision is defined by

$$K * \phi = \begin{cases} \mathsf{Cn}\left(\bigcap S_K(K|_\phi, \phi) \cup \phi\right) & \text{if } \phi \neq \bot; \\ L & \text{otherwise.} \end{cases}$$

That is to say, if $\phi \neq \bot$ it is the intersection of those ϕ-consistent maximal subsets chosen by S_K with ϕ added set-theoretically. If $\phi = \bot$ it is simply L, the inconsistent theory (the set of all sentences).

It should be clear that this is unsatisfactory, since the whole problem of how to make a revision has just been packaged up in the existence of a selection function, and has not been solved at all. Obviously, the selection function must depend on K. Therefore, we need not bother with the information K alone provides us, since everything we need might just as well be given by this magical S! The drawback of coding everything in S is that repeated revisions are then impossible.

There is a limiting case of partial meet revision, in which $S_K(K|_\phi, \phi)$ is always a singleton. This case is known as maxichoice contraction. There is another limiting case in which $S_K(K|_\phi, \phi) = \bigcap K|_\phi$, the intersection of all the candidate theories, which is known as full meet revision. The first of these is unsatisfactory for the same reason as the general case, namely that the selection function remains to be defined. (It has other, worse, problems too, detailed in Gärdenfors' book.) The second limiting case does not have this problem, and is worth spelling out in full, since it fully specifies how to carry out a revision without the need for extra information. According to it,

$$K * \phi = \begin{cases} \mathsf{Cn}\left(\bigcap K|_\phi \cup \phi\right) & \text{if } \phi \neq \bot; \\ L & \text{otherwise.} \end{cases}$$

It is straightforward to check that this definition satisfies the postulates K1-K8. But there are problems. Consider, for example, how to revise $\mathsf{Cn}(\{p, q\})$ with $\neg p \vee \neg q$. Intuitively, there are at least three plausible answers: $\mathsf{Cn}(\{p\})$, $\mathsf{Cn}(\{q\})$ and $\mathsf{Cn}(\{p \leftrightarrow \neg q\})$. Full meet contraction gives us the last of these, because no information is available to chose whether to give up p or to give up q. But, in practice there may be criteria for choosing to give up one rather than the other. This is what leads to consideration of selection functions, since they could encode the extra information required. But then, as already remarked, repeated revision is impossible. The moral *we* draw from this situation is different. It is that *deductively closed theories are inadequate as representations of belief states*. We return to this point later, after considering the other main way of providing the information necessary to guide revisions, namely epistemic entrenchment orderings.

2.2 Epistemic entrenchment

Revision by epistemic entrenchment is effected as follows. First we require an epistemic entrenchment ordering on the current belief state. This is a linear pre-order on the sentences in the state, which represents the degree to which they are believed. Those less entrenched according the ordering are dispensed with more readily in the case of a revision which conflicts with the current state. An epistemic entrenchment ordering for a belief state K must satisfy the following axioms:

- **E1** If $\phi \leqslant_K \psi$ and $\psi \leqslant_K \chi$ then $\phi \leqslant_K \chi$ (*transitivity*);
- **E2** If $\phi \models \psi$ then $\phi \leqslant_K \psi$ (*dominance*);
- **E3** Either $\phi \leqslant_K \phi \wedge \psi$ or $\psi \leqslant_K \phi \wedge \psi$ (*conjunctiveness*);
- **E4** If K is consistent then $\phi \leqslant_K \psi$ for all ψ iff $\phi \notin K$ (*minimality*);
- **E5** If $\phi \leqslant_K \psi$ for all ϕ, then $\models \psi$ (*maximality*).

As in the case of the K postulates, these axioms are intended to encode rationality constraints on what an epistemic entrenchment ordering might be. For example, E2 says that it is always better to give up logically weaker sentences during the course of a revision; therefore, these should be less entrenched. E3 says that giving up a conjunction is at least as hard as giving up either of the conjuncts. Taken together, axioms E1–E3 imply that $\leqslant_K$ is a linear order, that is, either $\phi \leqslant_K \psi$ or $\psi \leqslant_K \phi$ (or both). E4 says that a sentence is minimally entrenched in K iff it is not in K. E5 says that just the tautologies are maximally entrenched.

Given a belief state K, an epistemic entrenchment ordering $\leqslant_K$ on K, and a sentence ϕ, the revision of K by ϕ is given by

$$K * \phi = \begin{cases} \mathsf{Cn}\big(\{\psi \in K \mid \neg\phi <_K \neg\phi \vee \psi\} \cup \{\phi\}\big) & \text{if } \phi \neq \bot; \\ L & \text{otherwise.} \end{cases}$$

($<$ is the usual strict counterpart of $\leqslant$, defined by: $\phi < \psi$ if $\phi \leqslant \psi$ and $\psi \not\leqslant \phi$.)

Full motivation for the K and EE axioms, as well as for the definition of $*$ in terms of $\leqslant_K$, can be found in Gärdenfors' book [3].

We now summarise the main weaknesses we have described of the AGM theory. Belief states are represented as deductively closed theories. This means that they are (in general) impossible to write down fully, or to store on a computer. Moreover, as noted, they are incapable of representing the necessary information required to chose between alternative revisions. Therefore, extra information in the form of a selection function or an EE ordering is required. This information is not deemed part of the belief state, and is lost during the revision, making further revision impossible. It is worth pointing out that this means that the real *intention* of axiom K1 is not satisfied by these revision functions. Its intention is that after a revision we should end up with an object of the same type as the one with which we started. Obviously, both partial meet revision and revision by epistemic entrenchment fail this requirement. In those cases we start off with a pair, respectively of type $\langle K, S_K\rangle$ and $\langle K, \leqslant_K\rangle$, and end up with something of type K.

There are some proposals for modifying the AGM theory to solve some of these problems. For example, some work has been done on theory base revision to address the problem of the infinite nature of deductively closed sets of sentences. In that work, belief states are represented as finite sets of sentences (theory *bases* or theory *presentations*) [1, 5, 9]. But each of these authors assume the existence of something like a selection function or an EE ordering, so are subject to objections on those grounds. There are proposals of non-deterministic revision [6], which alleviate the need for a selection function, but they rely on infinite belief state representations.

There are proposals to allow repeated revision using EE orderings, either by keeping a single EE ordering for all belief states or assuming the existence of a function which, for every belief state, gives an EE ordering [10, 13]. But as neither the single ordering nor this function is itself revised in the course of belief revisions, it is easy to find examples which are in contradiction with intuitions about iterated belief change [4].

Another modification of the AGM theory which allows EE orderings to be revised is given by H. Rott [11]. He defines revision of EE orderings as follows.

$$\psi \leqslant_{K*\phi} \chi \text{ if } \phi \rightarrow \psi \leqslant_K \phi \rightarrow \chi.$$

However, as he points out, this fails to capture much of the intuition of repeated revision because any further revision of $K * \phi$ always includes ϕ.

3 Criteria for belief revision

In this section we enumerate what we claim are the criteria by which to judge a theory of belief revision.

1. Finite representation of belief states.
2. Persistence.
3. Iteration: what you put in is what you get out.
4. The "intentions" behind the K axioms of AGM.

The criterion of finite representation means that all belief states can be explicitly written down or represented on a computer. The advantages of this should be easy to see; one in particular is that one can give examples of belief revision in action! (See section 7.)

Persistence means that as much of the former belief state should survive a revision as possible. We rule out revisions like the one of proposition 2.

The iteration criterion says that you should get out of a revision an object of the same type as you put in. As mentioned, this is violated by AGM, since you put in either an EE ordering, or a theory coupled with a selection function; but, all you get out is a theory. We call this *iteration* since, if it obtains, it guarantees that revisions may be repeated. Its absence is a serious problem in AGM.

The last criterion, concerning the K axioms of AGM, is deliberately expressed in a vague way. Obviously, if belief states are not represented as deductively closed sets of sentences then it is impossible to test them literally. Also, as we have noted, they do not specify what should happen under repeated revision, in terms of expressions of the form $K * \phi * \psi$. This is presumably because the AGM models do not support repeated revision. Moreover, for reasons which we will discuss in section 6, we dispute two of the AGM axioms. In view of these reasons, we can only say that something like the intention of the AGM axioms is desirable.

The AGM axioms K1–8 rely on a particular representation of belief states (namely, deductively closed sets of sentences). Therefore, direct comparison with theories of belief revision which use other representations of belief states is impossible. To overcome this we can we rewrite the axioms in a more general way, which assumes only the following:

1. A set of belief states, together with a subset of 'contradictory' belief states.
2. A function $*$ (revision) which takes a belief state and a sentence to a belief state;
3. A function $|\cdot|$ (extension) which takes a belief state and returns the set of sentences true in it.

Here are the axioms rewritten in this way. We will write $\mathcal{K}$ for a typical 'abstract' belief state.

- $\mathcal{K}1$ $\quad \mathcal{K} * \phi$ is a belief state;
- $\mathcal{K}2$ $\quad \phi \in |\mathcal{K} * \phi|$;
- $\mathcal{K}3$ $\quad |\mathcal{K} * \phi| \subseteq |\mathcal{K}| + \phi$;
- $\mathcal{K}4$ $\quad$ If $\neg\phi \notin |\mathcal{K}|$ then $|\mathcal{K}| + \phi \subseteq |\mathcal{K} * \phi|$;
- $\mathcal{K}5$ $\quad \mathcal{K} * \phi$ is contradictory implies $\phi = \bot$;
- $\mathcal{K}6$ $\quad$ If $\models \phi \leftrightarrow \psi$ then $|\mathcal{K} * \phi| = |\mathcal{K} * \psi|$;
- $\mathcal{K}7$ $\quad |\mathcal{K} * (\phi \wedge \psi)| \subseteq |\mathcal{K} * \phi| + \psi$;
- $\mathcal{K}8$ $\quad$ If $\neg\psi \notin |\mathcal{K} * \phi|$ then $|\mathcal{K} * \phi| + \psi \subseteq |\mathcal{K} * (\phi \wedge \psi)|$.

4 Ordered theory presentations

Here we present a system for belief revision which satisfies each of the criteria described above. Belief states are represented by *ordered theory presentations*, which were introduced in [12]. The reader unfamiliar with that paper is at no disadvantage, but should ignore the remainder of this subsection, skipping to section 4.1 below. The following explanations will be of use to readers familiar with [12].

In that paper, a 'structured theory' was defined to be a finite, *partially* ordered set of points each of which is labelled by a sentence in the language. We have since changed the name of these objects to 'ordered theory presentations', because what they really are is a way of using an ordering on a set of sentences to present a theory. Thus, they are theory presentations[2] equipped with a partial order.

[2] Theory presentations are sometimes also called theory bases.

Here, we are interested only in the linear case, that is to say, in the case that the order is *total*. Thus, we can simplify the definition, letting *linear* ordered theory presentations be simply finite sequences (or lists) of formulas.

4.1 Intuitions

A (linear) ordered theory presentation is a finite list of formulas; ? $= [\phi_1, \phi_2, \ldots \phi_n]$. Here, n is said to be the *length* of ?. The *extension* of ? is the deductively-closed theory which ? presents; that is, it is the set of sentences entailed by ?, after taking account of the various conflicts in ?. This is defined precisely later on, but the intuition is the following: ? presents the theory which first of all has ϕ_n, and then has as much of ϕ_{n-1} as possible while retaining consistency, and then ... up to ϕ_1. Put another way, we start with ϕ_1. Then we 'force in' ϕ_2, overriding as necessary. Then ... and so on until ϕ_n.

The following are examples of belief states.

1. $[p \wedge q]$ 2. $[p, q]$ 3. $[p \wedge q, \neg p]$

Their lengths are 1, 2 and 2 respectively. OTPs 1 and 2 above both have the extension $\mathsf{Cn}(\{p \wedge q\})$. But in 2, p is less entrenched than q, and will disappear if a revision which demands that one of p and q goes. Thus, we stipulate:

Sentences later in the sequence are more entrenched than those earlier.

Belief state 3 has the extension $\mathsf{Cn}(\{\neg p \wedge q\})$. This is because $\neg p$, which is more entrenched than $p \wedge q$, overrides the p component of $p \wedge q$. But the q component is not overridden. Thus,

Sentences later in the sequence have the effect of overriding those earlier, in the case of conflict.

It should be noted that we have not yet given a definition of extension; this comes in section 4.2. The principles above are to motivate the right intuitions.

It should now come as no surprise to find that

Revision of OTPs is effected by appending the revising sentence to the end of the sequence.

Thus, the three belief states mentioned above can be revised by $\neg p \vee \neg q$, yielding

$1'$. $[p \wedge q, \neg p \vee \neg q]$ $2'$. $[p, q, \neg p \vee \neg q]$ $3'$. $[p \wedge q, \neg p, \neg p \vee \neg q]$.

The extension of presentation $1'$ is $\mathsf{Cn}(\{p \leftrightarrow \neg q\})$, which was the outcome of the corresponding example for full meet revision described above. But presentation $2'$ has as extension $\mathsf{Cn}(\{q\})$. Since 1 and 2 had the same extension and $1'$ and $2'$ do not, it should be clear that *there is more to an OTP than its extension.*

Ordered theory presentation $3'$ has the extension $\mathsf{Cn}(\{\neg p \wedge q\})$, which is the same as it had before the revision. This is because the revising sentence was consistent with the belief state it revised.

It remains to show how to formalise the computation of extension for ordered theory presentations, which is the task of the next section. First, let us observe some important facts about OTP revision.

1. OTPs have *memory.* If ? is an OTP, then the extension of ? $* p * q * \neg(p \wedge q)$ includes $q \wedge \neg p$. This is because the theory was more recently revised with q than with p, so q is more entrenched. Older information is discarded more readily than newer.
2. But, information is never wantonly discarded.
3. The more you revise an ordered presentation, the more complicated (= longer) it gets. That is because ordered presentations are nothing more than *revision histories.*

4.2 Extension

In classical logic, a theory presentation denotes a set of models, containing those which satisfy each of the sentences[3]. Two presentations are equivalent if they denote the same sets. In this section we will define models of ordered presentations, extending the notion of equivalence to OTPs. Thus, the extension of a ordered presentation is the theory of the set of models it denotes. It turns out that

1. Every OTP has a non-trivial extension, and
2. In the case of classical propositional logic at least, the set of models of any OTP is denoted by a finite theory presentation.

Definition 3 An *ordered theory presentation* of length n is a finite list of formulas of length n. Notation: ? $= [\phi_1, \phi_2, \ldots, \phi_n]$.

As stated, our aim is to define the models of such ordered theory presentations. We will assume that we are working with a fixed propositional language L, and that $\Vdash$ is the satisfaction relation between interpretations (assignments of truth values to propositional symbols) and sentences of the language. L is ambiguously the language and the set of sentences of the language. Everything here is easily re-worked for predicate logic, or indeed for any logic defined in terms of interpretations and satisfaction. (See [12] for details; occasionally we will add parenthesised notes to this effect.)

Let $\mathcal{M}$ be the set of interpretations of L, and let $\Vdash$ be the satisfaction relation between interpretations and sentences of the language: $\Vdash \subseteq \mathcal{M} \times L$.

To define the models of an ordered theory presentation ? we need to define an ordering $\sqsubseteq^\Gamma$ on interpretations in $\mathcal{M}$ which measures how well an interpretation satisfies ?. This relies on orderings $\sqsubseteq_\phi$, one for each sentence ϕ of the language. To define $\sqsubseteq_\phi$, it is necessary to define a notion which we call 'natural entailment', written $\models\!\!\!\!^{\smile}$. This definition in turn relies on the notion of the *monotonicities* of a sentence. Lest the reader be daunted by these nestings of definitions, we repeat the list here. We define, in order,

1. *Monotonicities* of a sentence ϕ, written $\langle \phi^+, \phi^- \rangle$.
2. *Natural entailment*, a relation $\models\!\!\!\!^{\smile}$ between sentences, being a sub-relation of ordinary entailment $\models$.
3. For each sentence ϕ, a reflexive and transitive order $\sqsubseteq_\phi$ on the interpretations of the language $\mathcal{M}$; as will be seen, this order grades interpretations according to how nearly they satisfy ϕ.
4. For each ordered presentation ?, a reflexive and transitive order $\sqsubseteq^\Gamma$ on $\mathcal{M}$; this order grades interpretations according to how well they satisfy ?.
5. The models of ?, being a set $[\![?]\!] \subseteq \mathcal{M}$.

For the definition of the monotonicities of ϕ, we need the following notation. If M is an interpretation of L and p is a propositional symbol in L, then $M^{[p \mapsto t]}$ is an interpretation identical with M except possibly that it assigns **true** to p. (If M already assigns **true** to p then $M^{[p \mapsto t]}$ is simply M.) $M^{[p \mapsto f]}$ is defined analogously.

Definition 4 Let ϕ be a sentence *other than* $\bot$ and p any propositional symbol.

1. ϕ is *monotonic in p* if $M \Vdash \phi$ implies that $M^{[p \mapsto t]} \Vdash \phi$.
2. ϕ is *anti-monotonic in p* if $M \Vdash \phi$ implies that $M^{[p \mapsto f]} \Vdash \phi$.
3. ϕ^+ and ϕ^- are the sets of symbols in which ϕ is monotonic and anti-monotonic respectively.

[3] Of course, the same goes for many non-classical logics too, including intuitionistic logic, modal logics, temporal logics etc. Readers familiar with [12] will recall that ordered theory presentations are defined over any logic defined in terms of models and satisfaction.

The case that $\phi = \bot$ is handled separately; we define $\bot^+ = \bot^- = \emptyset$.

Thus, ϕ is monotonic in p if "increasing" the truth value of p in a model of ϕ preserves satisfaction of ϕ. Similarly, ϕ is anti-monotonic in p if descreasing the truth value so preserves satisfaction. (In the case that L is a predicate language, one must define the *extension* of a predicate symbol P in a model to be the set of tuples which satisfy P in the model. ϕ is (anti-)monotonic in the predicate symbol P if increasing (decreasing) the extension of P in a model of ϕ also results in a model of ϕ.)

Example 5 Suppose L has just the propositional symbols p and q.

ϕ	ϕ^+	ϕ^-
$\top$	$\{p, q\}$	$\{p, q\}$
p	$\{p, q\}$	$\{q\}$
q	$\{p, q\}$	$\{p\}$
$p \wedge q,\ p \vee q$	$\{p, q\}$	$\emptyset$
$p \rightarrow q$	$\{q\}$	$\{p\}$
$p \leftrightarrow q$	$\emptyset$	$\emptyset$
$\bot$	$\emptyset$	$\emptyset$

Having defined monotonicities, we turn to point 2 of the five-point plan mentioned above, i.e. the definition of natural entailment. Let ϕ and ψ be sentences of L.

Definition 6 ϕ *naturally entails* ψ, written $\phi \stackrel{\smile}{\models} \psi$, if

1. $\phi \models \psi$, and
2. $\phi^+ \subseteq \psi^+$ and $\phi^- \subseteq \psi^-$

Natural entailment is a sub-relation of ordinary entailment; in addition to ordinary entailment we require that the monotonicities of the premise be preserved by the conclusion.

Proposition 7 $\stackrel{\smile}{\models}$ is reflexive and transitive[4]. $\diamondsuit$

Example 8 The relations $\models$ and $\stackrel{\smile}{\models}$ on the set of sentences formed from the language containing the propositions $\{p, q\}$ are shown in figure 1 for comparison. Thus: $p \wedge q \stackrel{\smile}{\models} p$ and $p \wedge q \stackrel{\smile}{\models} p \vee q$, but $p \wedge q \not\stackrel{\smile}{\models} p \leftrightarrow q$ and $p \not\stackrel{\smile}{\models} p \vee q$. Moreover, $\bot \stackrel{\smile}{\models} \phi$ for all ϕ.

The definition of natural entailment is perhaps not very satisfying, because (one might ask), what is so special about preserving monotonicities? One way to answer this is purely pragmatic: as we will see, it is essential for the next definition, which does have a satisfying feel. But first, we justify the term *natural entailment* by showing examples of how much more natural this entailment really is.

Natural entailment is something like relevant entailment; it stops us adding irrelevant disjuncts in our conclusions. (This is not the same notion of relevance as Anderson/Belnap, for there one is interested in stopping irrelevant conjuncts in the premisses.) The following entailments, which are ordinarily valid, are not naturally valid:

$$\begin{array}{lll} p \models p \vee q & \quad p \models q \rightarrow p & \quad \\ p \models p \vee \neg q & \quad \neg p \models p \rightarrow q & \quad p \wedge q \models p \leftrightarrow q \end{array}$$

Regarding the first pair, the premise p tells us nothing about q, and therefore it is suspect to introduce q or $\neg q$ as a disjunct. The second pair are the standard inelegancies of material implication, which everyone will be glad to see the back of. Finally, we dislike $p \wedge q \models p \leftrightarrow q$ because

[4] the proofs of this proposition and others in this section which are given without proofs may be found in [12].

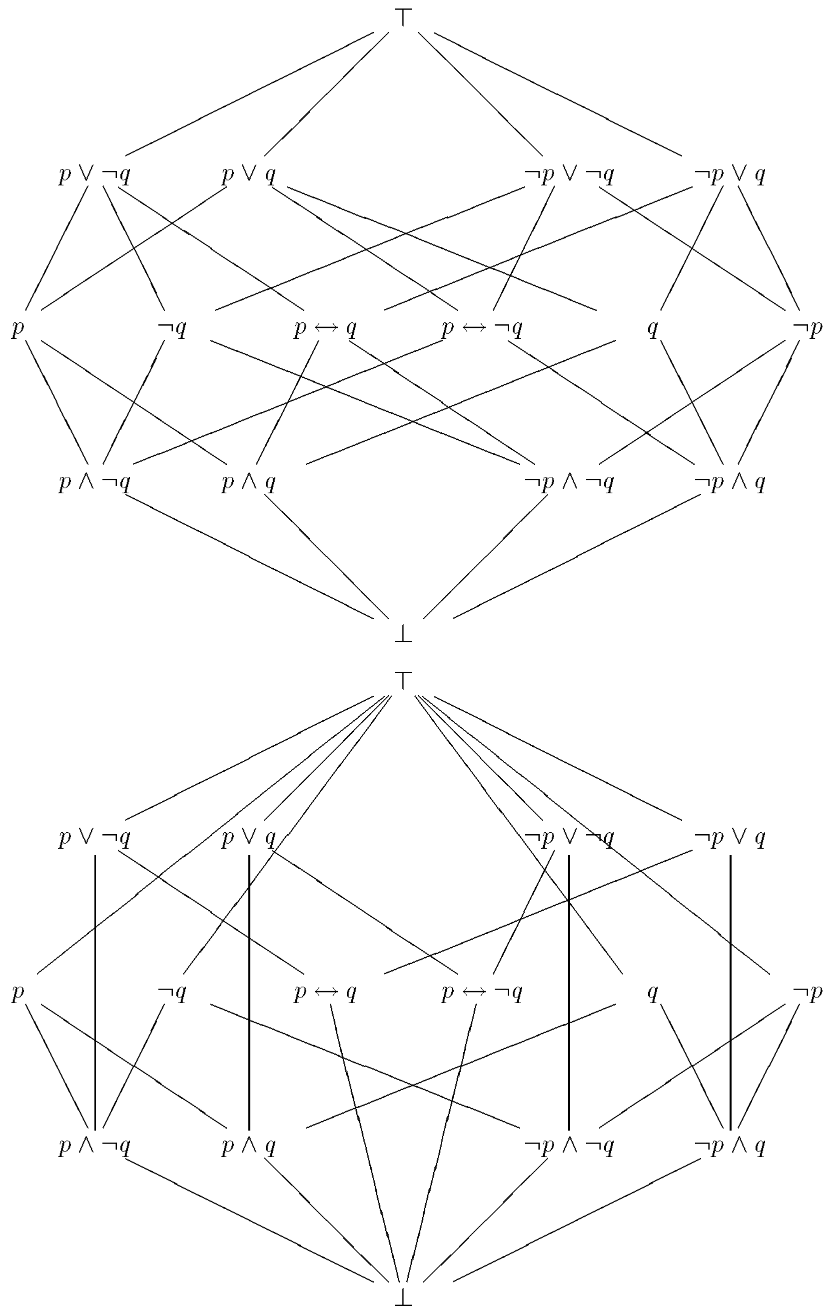

Figure 1: The *ordinary* and *natural* entailment relations over $\{p, q\}$

the right-hand side suggests that p and q are in some way bound together, whereas the left-hand side only says that they are both true.

On the other hand, the simplicity of the definition and the fact that it is based on satisfaction by models ensures that there is nothing untoward going on. In particular, if ϕ and ψ are classically equivalent then they are naturally equivalent; indeed

$$(\phi \models \psi \text{ and } \psi \models \phi) \text{ iff } (\phi \mathrel{\dot{\models}} \psi \text{ and } \psi \mathrel{\dot{\models}} \phi)$$

I really do think there is something wonderful about natural entailment, even though I only stumbled upon it as a means of achieving the following definition of $\sqsubseteq_\phi$. More in another paper!

How, then, to define $\sqsubseteq_\phi$? As stated, $M \sqsubseteq_\phi N$ means that N is as good at satisfying ϕ as M is. It is not just that N satisfies ϕ and M does not; perhaps neither satisfy ϕ, but N more nearly does. For example, let M be a propositional interpretation which assigns false to both p and q; and let N assign true and false to p and q respectively. Then $M \sqsubseteq_{p\wedge q} N$, while $N \not\sqsubseteq_{p\wedge q} M$. Neither satisfy $p \wedge q$, but at least N satisfies p; M doesn't satisfy either of p and q.

This example shows that one has to look at which *consequences* of ϕ are satisfied by M and N. However, defining $M \sqsubseteq_\phi N$ to mean that N satisfies all the consequences of ϕ which M does gives us precisely the bipartite ordering rejected in the preceding paragraph. This is because ϕ has too many irrelevant consequences; we should just look at the *natural* ones.

Definition 9 $M \sqsubseteq_\phi N$, if for each ψ,

$$\phi \mathrel{\dot{\models}} \psi \Rightarrow (M \Vdash \psi \Rightarrow N \Vdash \psi)$$

We can show that $\sqsubseteq_\phi$ has precisely the mathematical behaviour we want.

Proposition 10

1. $\sqsubseteq_\phi$ is a pre-order, that is to say, it is reflexive and transitive.
2. If $\phi \neq \bot$, the maximal elements of $\sqsubseteq_\phi$ (which are in fact maxi*mum*) are just the models of ϕ.
3. The orderings $\sqsubseteq_\bot$ and $\sqsubseteq_\top$ are the indiscrete ordering; that is, $\sqsubseteq_\bot = \sqsubseteq_\top = \mathcal{M} \times \mathcal{M}$. $\diamondsuit$

The proofs of these assertions, together with many examples of $\sqsubseteq_\phi$ for various sentences ϕ, can be found in [12].

We have defined, for each sentence ϕ, an ordering on interpretations $\sqsubseteq_\phi$ which measures the extent to which interpretations satisfy ϕ. If M satisfies ϕ to the fullest extent (that is, if it simply satisfies it) then M is $\sqsubseteq_\phi$-maximum. If M does not fully satisfy ϕ then it may satisfy it to a greater, lesser, equal or incomporable extent than some N which perhaps also fails fully to satisfy ϕ. As stated, examples of this ordering on models for various sentences ϕ can be found in [12].

We continue with the five-point plan at the beginning of this section. We define $\sqsubseteq^\Gamma$ by induction on ?. If ? is the empty list [ ], then $M \sqsubseteq^\Gamma N$ for all M, N. Otherwise, if ? is $[\phi_1, \ldots, \phi_n]$ then $M \sqsubseteq^\Gamma N$ if either $M \sqsubset_{\phi_1} N$, or $M \sqsubseteq_{\phi_1} N$ and also $M \sqsubseteq^{\Gamma'} N$, where ?′ is $[\phi_2, \ldots, \phi_n]$. In other words, to determine whether N is as good at satisfying ? as M is, one looks at the most important sentence. If N is strictly better than M as far as that sentence is as concerned, then N is certainly as good overall. Otherwise, it is at least necessary for N to be as good on that sentence, and the question of whether it as good overall defers to the remaining sentences which are treated in a similar way. We summarise:

Definition 11

1. $M \sqsubseteq^{[\,]} N$ always; and
2. $M \sqsubseteq^{\Gamma * \phi} N$ if $M \sqsubset_\phi N$ or ($M \sqsubseteq_\phi N$ and $M \sqsubseteq^\Gamma N$).

Remark 12 $M \sqsubset^{\Gamma * \phi} N$ if $M \sqsubset_\phi N$ or ($M \sqsubseteq_\phi N$ and $M \sqsubset^\Gamma N$).

$? * \phi$ is ? with ϕ appended. Finally, it is easy to define the models of ?. They are simply the interpretations which are rated maximally by $\sqsubseteq^{\Gamma}$.

Definition 13 $M \Vdash ?$ if M is $\sqsubseteq^{\Gamma}$-maximal.

Naturally we expect that the highest priority sentence is satisfied by models of the theory:

Proposition 14 Let $? = [\phi_1, \phi_2, \ldots, \phi_n]$ be an OTP and let $M \Vdash ?$. If $\phi_n \neq \bot$ then $M \Vdash \phi_n$. ♢

Definition 15 Let ? be an ordered theory presentation.

1. The *extension* of ?, written $|?|$, is the theory of the set of models of ?:

$$|?| = \{\phi \mid M \Vdash ? \text{ implies } M \Vdash \phi\}$$

2. The *consequences* of ? are the sentences in its extension:

$$? \models \phi \text{ if } \phi \in |?|.$$

4.3 Summary of definitions

In this section we summarise the position so far. We assume we are working in classical logic. Ordered theory presentations consist of finite lists of sentences in the language (def. 3). To define the models of an ordered theory presentation, we first define the monotonicities $\langle \phi^+, \phi^- \rangle$ of each sentence ϕ (def. 4). This is a pair of sets of atomic sentences. Then we define the relation of natural entailment between sentences (def. 6). We claim, in passing, that this has intuitive properties which ordinary entailment fails, but our main purpose is to use it to define the degree of satisfaction between an interpretation and a sentence. We do this by ordering interpretations according to how well they satisfy a particular sentence ϕ, in definition 9. These ordering are used to define an ordering for the whole OTP which measures how well interpretations do at satisfying it. Finally, its models are the interpretations maximal in this ordering. The extension of an OTP and its consequences are then straightforward to define (def. 15).

5 Properties of OTPs

Proposition 16 Every ordered presentation has a model.

Proof This is a complicated proof, using Zorn's lemma to find maximal elements in the order $\sqsubseteq^{\Gamma}$. For details, see [12]. ♢

A consequence of this result is that contradictions can never be derived from an ordered presentation, not even one with the contradictory sentence in it!

Proposition 17 If $? \models \phi$ then $\phi \neq \bot$.

Proof Let $M \Vdash ?$. Since $M \Vdash \phi$, $\phi \neq \bot$. ♢

Definition 18 Let ? and Δ be OTPs.

1. ? and Δ are *statically equivalent*, written $? \equiv \Delta$, if they have the same extension:

$$? \equiv \Delta \text{ if } |?| = |\Delta|.$$

2. ? and Δ are *dynamically equivalent*, if, for all ϕ, $? * \phi \equiv \Delta * \phi$.

Dymamic equivalence implies static equivalence, but the converse is not so as the following example shows.

Example 19 $[p, q] \equiv [p \wedge q]$, since both have the models $\{11\}$ in the obvious notation. But $[p, q, \neg p \vee \neg q] \not\equiv [p \wedge q, \neg p \vee \neg q]$, since the model sets are respectively $\{01\}$ and $\{01, 10\}$.

However, if $? \models \phi$ we would not expect that revising ? by ϕ should change the set of models.

Proposition 20 If $? \models \phi$ then $? \equiv ? * \phi$. $\diamondsuit$

We also obtain weak analogues of the usual structural properties:

Proposition 21

1. Weak inclusion: if $\phi \neq \bot$ then $? * \phi \models \phi$
2. Weak monotonicity: $\dfrac{? \models \phi \quad ? \models \psi}{? * \phi \models \psi}$
3. Weak cut: $\dfrac{? * \phi \models \psi \quad ? \models \phi}{? \models \psi}$

These principles are accepted as being requirements which a system for belief revision or for defaults should have (see for example [2, 7, 8]).

6 Belief revision: the AGM postulates

As stated, we intend to use these ordered theory presentations as representations of belief states in order to model belief revision. The obvious way to do this is to let

$$\text{belief states} = \text{ordered theory presentations}$$

and define $? * \phi$ to be ? with ϕ appended; of course we have been implicitely assuming this definition so far in the paper. Note that under this arrangement there are no contradictory theories (proposition 17).

In this setting, we can investigate the truth or falsity of the abstract K axioms given in section 3. We obtain the following.

$\mathcal{K}$**1** $? * \phi$ is a belief state.

This is true. If ? is an OTP then so is $? * \phi$.

$\mathcal{K}$**2** $\phi \in |? * \phi|$.

This is false. For example, $\bot \notin |[\,] * \bot|$; for, as one can check, $|[\,] * \bot| = \mathsf{Cn}(\emptyset)$. However, $\mathcal{K}2$ is true if $\phi \neq \bot$, by proposition 14.

$\mathcal{K}$**3** $|? * \phi| \subseteq |?| + \phi$.

True. We need to show that $M \Vdash ?$ and $M \Vdash \phi$ imply $M \Vdash ? * \phi$. Suppose not, i.e. suppose $M \sqsubset^{\Gamma * \phi} N$ for some N. By lemma 12, either $M \sqsubset_\phi N$, which contradicts $M \Vdash \phi$ (proposition 10) or $M \sqsubset^\Gamma N$, which contradicts $M \Vdash ?$ (definition 13).

$\mathcal{K}$**4** If $\neg\phi \notin |?|$ then $|?| + \phi \subseteq |? * \phi|$

This is false. Let $\phi_1 = p \wedge q \wedge r$, $\phi_2 = \neg p \vee \neg q \vee \neg r$ and $\phi_3 = (p \leftrightarrow q) \vee \neg r$. The counterexample is obtained by setting: $? = [\phi_1, \phi_2]$ and $\phi = \phi_3$. To see this, we should first examine the orderings for each of ϕ_1, ϕ_2 and ϕ_3. They are shown in the top half of figure 2. Applying definition 11, the orderings $\sqsubseteq^\Gamma$ and $\sqsubseteq^{\Gamma * \phi}$ (i.e. $\sqsubseteq^{[\phi_1, \phi_2]}$ and $\sqsubseteq^{[\phi_1, \phi_2, \phi_3]}$ respectively) are as shown in the bottom half of the figure. We can check the following:

- $\neg\phi_3 \notin |[\phi_1, \phi_2]|$, that is to say, there is a model M such that M is $\sqsubseteq^{[\phi_1, \phi_2]}$-maximal and $M \not\Vdash \neg\phi_3$. Such an M is 110. Thus, the antecedent of K4 holds.

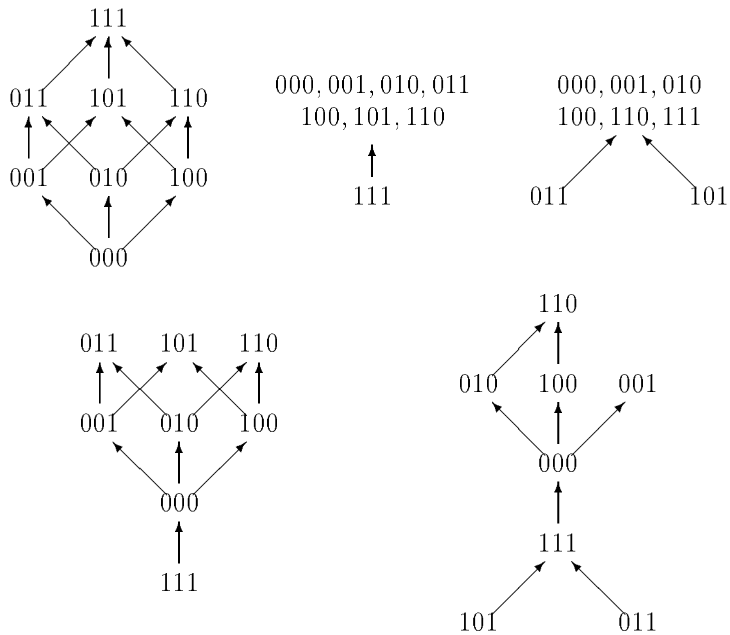

Figure 2: The counterexample to $\mathcal{K}4$ (see text)

– But the consequent is false. For if $|?|+\phi \subseteq |?*\phi|$ then $M \Vdash ?*\phi$ implies $M \Vdash |?|+\phi$, i.e. $M \Vdash ?$ and $M \Vdash \phi$. But by inspecting the diagrams we can find M such that $M \Vdash ?*\phi$ but $M \nVdash ?$, namely $M = 001$.

$\mathcal{K}$5 $?*\phi$ is contradictory implies $\phi = \bot$.

This is vacuously true since there are no contradictory belief states.

$\mathcal{K}$6 If $\models \phi \leftrightarrow \psi$ then $|?*\phi| = |?*\psi|$.

True. Suppose $\models \phi \leftrightarrow \psi$. It is sufficient to prove $\sqsubseteq_\phi = \sqsubseteq_\psi$. Suppose $M \sqsubseteq_\phi N$, and $\psi \mathrel{\models\!\!\!\!\!\equiv} \chi$ and $M \Vdash \chi$. By reflexivity, $\phi \mathrel{\models\!\!\!\!\!\equiv} \psi$, so by transitivity $\phi \mathrel{\models\!\!\!\!\!\equiv} \chi$. Therefore, $N \Vdash \chi$, so $M \sqsubseteq_\psi N$. The converse is proved similarly.

$\mathcal{K}$7 $|?*(\phi \wedge \psi)| \subseteq |?*\phi| + \psi$.

True. We need to show that if $M \Vdash ?*\phi$ and $M \Vdash \psi$ then $M \Vdash ?*(\phi \wedge \psi)$. If $\phi = \bot$ then $?*\phi \equiv ?*(\phi \wedge \psi)$, and we are done. So suppose $\phi \neq \bot$, and $M \Vdash ?*\phi$ and $M \Vdash \psi$, but $M \sqsubset^{\Gamma*(\phi\wedge\psi)} N$ for some N. Since $M \Vdash ?*\phi$ and $\phi \neq \bot$, we have $M \Vdash \phi$ by proposition 14. Therefore, $M \Vdash \phi \wedge \psi$. By lemma 12, either $M \sqsubset_{\phi\wedge\psi} N$, which contradicts $M \Vdash \phi \wedge \psi$, or $M \sqsubset^{\Gamma} N$. But this also leads to a contradiction, for then, since $M \sqsubseteq_\phi N$, we obtain $M \sqsubset^{\Gamma*\phi} N$ by lemma 12, contradicting $M \Vdash ?*\phi$.

$\mathcal{K}$8 If $\neg\psi \notin |?*\phi|$ then $|?*\phi| + \psi \subseteq |?*(\phi \wedge \psi)|$

False. The counterexample given for $\mathcal{K}4$ holds here too. Set $? = [p \wedge q \wedge r]$, $\phi = \neg p \vee \neg q \vee \neg r$ and $\psi = (p \leftrightarrow q) \vee \neg r$.

On this way of using OTPs as belief states, we have shown that $\mathcal{K}1$, $\mathcal{K}3$, $\mathcal{K}5$, $\mathcal{K}6$ and $\mathcal{K}7$ are valid; that $\mathcal{K}2$ is valid under the proviso that $\phi \neq \bot$; and that $\mathcal{K}4$ and $\mathcal{K}8$ are not valid.

It is worth pointing out that the lack of contradictory belief states and the partial failure of $\mathcal{K}2$ are easily solved, by adding a new belief state to represent the contradictory belief state and modifying the definition of revision. Thus,

$$\text{belief states} = \text{ordered theory presentations} \cup \{\bot\}.$$

Revision on these belief states is defined as follows:

$$? \breve{*} \phi = \begin{cases} \bot & \text{if } \phi = \bot \\ [\phi] & \text{if } \phi \neq \bot \text{ and } ? = \bot \\ ? * \phi & \text{otherwise} \end{cases}$$

This emulates what the AGM axioms intend for $\bot$, in that

1. There is a unique contradictory belief state.
2. Revising any state with the contradictory sentence results in the contradictory state (K2).
3. The contradictory state can only be obtained in this way (K5), so in particular
4. Revising the contradictory state with a non-contradictory sentence will *not* result in the contradictory state.

For the psychological plausibility of these stipulations, or otherwise, see [3]. Especially the first one is debatable! Our point is simply that if we take this definition of $? \breve{*} \phi$ on board, we obtain that $\mathcal{K}1$, $\mathcal{K}2$, $\mathcal{K}3$, $\mathcal{K}5$, $\mathcal{K}6$, and $\mathcal{K}7$ are satisfied, and $\mathcal{K}5$ is satisfied in a more satisfying manner. $\mathcal{K}4$ and $\mathcal{K}8$ are still false for the same reasons.

6.1 The AGM axioms $\mathcal{K}4$ and $\mathcal{K}8$

$\mathcal{K}4$ and $\mathcal{K}8$ are serious violations of the AGM axioms, and there is no easy way of making them satisfied in the framework of OTPs. One must face the question: are they desirable axioms for belief revision? We believe the answer is no.

Consider the diagrams given in figure 2. As far as our counterexample is concerned, the question of the validity of $\mathcal{K}4$ hinges on whether $001 \sqsubset_{\phi_1} 110$ or not. If this was so, then we would also have $001 \sqsubset^{[\phi_1,\phi_2,\phi_3]} 110$ and $[\phi_1, \phi_2, \phi_3]$ would have only the model 110. Therefore, $\mathcal{K}4$ (and $\mathcal{K}8$) would hold.

Should $001 \sqsubset_{p \wedge q \wedge r} 110$ be the case? At first sight it seems clear that 110 is better at satisfying $p \wedge q \wedge r$ than 001 is, for 110 satisfies two of the atomic propositions while 001 satisfies only one. But this kind of cardinality argument is flawed. Why is it better to satisfy $p \wedge q$ rather than r? Perhaps r itself expresses a conjunction of facts. Are two oranges better than one apple?

The AGM book does not provide any argument in favour of $\mathcal{K}4$ and $\mathcal{K}8$. Consider the following story. I am expecting a friend called John to arrive. He can come by car, bike, or train. I am doubtful about whether he will arrive or not, however, because I believe that his car and bike are both at the repairers; and also, the trains are not working today (for a change). Let:

p	mean that his car is unavailable for use
q	his bike is unavailable
r	the trains are unavailable

Initially I believe

$$p \wedge q \wedge r.$$

Now suppose John actually arrives. I have no reason to doubt that he came by one of the usual means of transport (for example, he didn't ask me for money for a taxi). Therefore I revise my beliefs by

$$\neg p \vee \neg q \vee \neg r.$$

In the course of conversation it turns out that the repairer phoned him this morning to say that both his car and his bike were available for collection. I reason as follows. If the trains are still not working, he may have asked Richard for a lift to the repairer. His bike fits in the back of Richard's car, so then they could have collected both items. But, Richard may have been unavailable or unwilling. Either way, he will have collected both items or neither, so I revise with:

$$r \rightarrow (p \leftrightarrow q)$$

If the trains are working ($\neg r$) I cannot draw the conclusion $p \leftrightarrow q$, since he may have gone by train to pick up either the car or the bike, or neither, or he may still have asked Richard and got both.

The question now is: have I got enough information to conclude which means of transport were available for John to use?

We believe the answer is no. Suppose r, that is, the trains are still not working. I have already reasoned that this implies $p \leftrightarrow q$, and since John is actually here (so $\neg p \vee \neg q \vee \neg r$), it must be that $\neg p \wedge \neg q$. Therefore, $\neg p \wedge \neg q \wedge r$. On the other hand, suppose $\neg r$, i.e. that the trains *are* working. This tells me nothing about p and q. But since I started with the belief that $p \wedge q$ and John's arrival (by train, presumably) is consistent with these, I retain them. Therefore, $p \wedge q \wedge \neg r$. So I conclude $(\neg p \wedge \neg q \wedge r) \vee (p \wedge q \wedge \neg r)$, or, equivalently, $(p \leftrightarrow q) \wedge (p \leftrightarrow \neg r)$.

We have argued that it is not rational to conclude $\neg r$ in this case. We have also noted that the theory of belief revision outlined in this paper does not conclude $\neg r$. Indeed, we have argued that it concludes precisely what it is rational to conclude. It should be pointed out in fairness to the AGM theory that it does not insist on $\neg r$ either. To see this, consider what happens if the revision function specified in proposition 2 is applied to the revision history in question. We get

$$\begin{aligned} &\mathsf{Cn}\{p, q, r\} * (\neg p \vee \neg q \vee \neg r) * (r \rightarrow (p \leftrightarrow q)) \\ &\quad = \mathsf{Cn}\{\neg p \vee \neg q \vee \neg r\} * (r \rightarrow (p \leftrightarrow q)) \\ &\quad = \mathsf{Cn}\{(p \wedge q) \vee \neg r\} \end{aligned}$$

$\neg r$ is not derivable from this theory.

What we have shown is that if we augment the system of OTPs for belief revision so as to obtain $\mathcal{K}4$ and $\mathcal{K}8$, then we would have a system which concluded $\neg r$ in this case, which is undesirable.

7 Examples

Here we list some facts about linear OTPs, together with some references to examples in the literature to which the facts seem relevant.

$$\begin{aligned} |[p]| &= \mathsf{Cn}(\{p\}) \\ |[p, q]| &= \mathsf{Cn}(\{p, q\}) \\ |[p, q, \neg q]| &= \mathsf{Cn}(\{p, \neg q\}) \\ |[p, q, \neg p]| &= \mathsf{Cn}(\{\neg p, q\}) \\ |[p \wedge q, \neg p]| &= \mathsf{Cn}(\{\neg p, q\}) \\ |[p \wedge q, \neg p \vee \neg q]| &= \mathsf{Cn}(\{p \leftrightarrow \neg q\}) \\ |[p, q, \neg p \vee \neg q]| &= \mathsf{Cn}(\{\neg p, q\}) \\ |[p \vee q, \neg q]| &= \mathsf{Cn}(\{p, \neg q\}) \end{aligned}$$

We also have that

$$s \rightarrow p \in |[s, s \rightarrow p, s \rightarrow q, \neg q, \neg p]|$$

(cf. Hansson [4, page 7:12]), and, for example,

$$\begin{array}{lcl} p \leftrightarrow q \in |[p, q]|, & \text{but} & p \leftrightarrow q \notin |[p, q, \neg p]| \\ p \leftrightarrow q \in |[p, p \leftrightarrow q]| & \text{and} & p \leftrightarrow q \in |[p, p \leftrightarrow q, \neg p]| \end{array}$$

(cf. [4, page 4:3]).

References

[1] A. Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20:175–203, 1991.

[2] D. M. Gabbay and M. R. B. Clarke. An intuitionistic basis for non-monotonic logic. In P. Smets, A. Mamdani, D. Dubois, and H. Prade, editors, *Non-standard Logics for Automated Reasoning*. Academic Press, 1988.

[3] P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States.* MIT Press, 1988.

[4] S. O. Hansson. *Belief Base Dynamics.* PhD thesis, Department of Philosophy, Uppsala University, 1991.

[5] S. O. Hansson. From logical atoms to basic beliefs. Submitted for publication, 1992.

[6] S. Lindström and W. Rabinowicz. Epistemic entrenchment with incomparabilities and relational belief revision. In A. Fuhrmann and M. Morreau, editors, *The Logic of Theory Change*, volume 465 of *LNAI*. Springer Verlag, 1991.

[7] D. Makinson. General patterns in non-monotonic reasoning. In D. Gabbay, C. Hogger, and J. Robinson, editors, *Handbook of Logic in Artificial Intelligence.* Oxford University Press, 1992. Forthcoming.

[8] D. Makinson and P. Gärdenfors. Relations between the logic of theory change and non-monotonic logic. To appear.

[9] B. Nebel. *Reasoning and Revision in Hybrid Representation Systems.* Springer Verlag LNAI 422, 1990.

[10] H. Rott. Preferential belief change using generalised epistemic entrenchment. Konstanzer Berichte zur Logik und Wissenschaftstheorie 15.

[11] H. Rott. Two methods of constructing contractions and revisions of knowledge systems. *Journal of Philosophical Logic*, 20:149–173, 1991.

[12] M. D. Ryan. Defaults and revision in structured theories. In *IEEE Symposium on Logic in Computer Science (LICS)*, pages 362–373, 1991.

[13] K. Schlechta. Some results on theory revision. In A. Fuhrmann and M. Morreau, editors, *The Logic of Theory Change*, volume 465 of *LNAI*. Springer Verlag, 1991.