Representing Defaults as Sentences with Reduced Priority*

Mark Ryan
Department of Computing
Imperial College
London SW7 2BZ, UK
E-mail: mdr@doc.ic.ac.uk

Abstract

We distinguish between two ways of thinking about defaults. The first way, in which defaults augment known premises by 'strengthening' the underlying logic, is the traditional approach taken by most existing formalisms. In the second way, defaults are represented in the set of premises, but obtain their default status by having a reduced priority relative to the known premises. In this paper we:

1. Compare and contrast the approaches. We argue that the second approach makes for simpler representation of defaults and their interactions.

2. Describe a syntax and semantics for the second, less well-known approach; we introduce the notion of *ordered theory presentation* (OTP) to represent theories with defaults.

3. Show how ordered theory presentations can represent familiar examples of interacting defaults in an intuitively clear and simple way; we give the Tweety example and the Yale Shooting example. We also show that the OTP framework is particularly well suited to inheritance examples.

4. Show formal properties of OTPs, in particular cumulativity, and suggest connections with circumscription.

5. Show how OTPs may be used to model *belief revision* and compare the result with the standard theory.

1 Introduction

Most systems for reasoning with defaults treat them as a way of *strengthening the underlying logic*. For example, in circumscription defaults are represented by the policy of minimising certain predicates. Models of the circumscribed theory are those models of the original theory which have minimal extensions of those predicates. In particular, anything that can be proved without the defaults (i.e., without the minimisation) can also be proved with them; thus, the process of minimisation strengthens the deductive power of the logic. The same is true in, for example, negation as failure viewed as a default system; the default mechanism (in which the defaults are the negations of atomic formulas) allows us to derive more from a set of clauses than is classically derivable.

There is another view of defaults which is less widely known, although it has been described before [Bib85, Bre89, Poo88]. Whereas on the first view we had too few consequences of a theory, and used the default mechanism to add to them, on the second view we have too many consequences and the default mechanism reduces their number. In the second view, defaults are represented as *sentences in the theory* instead of as a means of augmenting the logic. The set of facts together with the set of defaults is in general contradictory. But the defaults are assigned a lower status, or reduced priority, than the other more certain sentences in the theory; this avoids contradictory conclusions. Much of this paper will flesh out both the syntax and the semantics of this 'reduced priority'.

We consider that the second view of defaults is preferable. Firstly, it provides a clearer way of specifying the default information. The fact that defaults are expressed as ordinary sentences using the full range of logical operators obviates the need for coding tricks which are often necessary in, for example, circumscription. Secondly, it treats defaults as part of the knowledge being represented, instead of as part of the logic. This gives improved knowledge representation.

In this paper we describe a system for representing defaults which falls into the second view of the two described above. In that system, a theory is presented as a partially-ordered set of sentences. (The exact definition is given in §2.) All of the sentences which we wish to represent are included in this set. That some of them are defaults and some are not is represented by their position in the ordering. The lower a sentence is

*To appear in: **Proceedings of the Third International Conference on Knowledge Representation and Reasoning** (KR'92), Morgan Kaufmann, 1992.

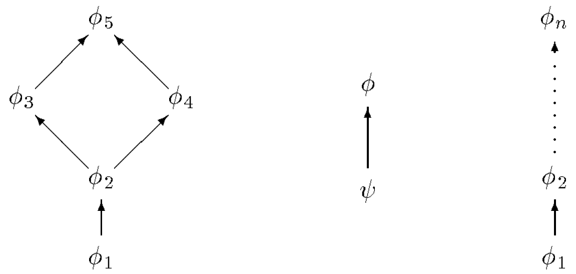

Figure 1: Three OTPs

in the ordering, the less of a default it is, and the minimal sentences are those which are not defaults at all. In the system we describe, we can have several levels of defaults—those below in the ordering override those above, if there is a conflict. Sentences can be a default relative to one sentence but not relative to another. We can specify which sentences are to override exactly which others.

This way of presenting a theory we call an *ordered theory presentation*, or OTP for short. Consider the first example of an OTP in figure 1 to make the above discussion of priority a little more concrete. In that ordered presentation, there is one 'fact', namely ϕ_1. It has a greater priority than the other sentences. The others are defaults, but still, some are stronger than others. For example, ϕ_3 is stronger than ϕ_5, weaker than ϕ_2 and incomparable in strength with ϕ_4. Thus, the arrow is read as 'is stronger than' or 'dominates'. This information is part of the knowledge being represented.

If a sentence dominates another, that means that it can override it if the two conflict. The meaning of the second OTP in figure 1 is $\phi \wedge \psi$ if ϕ and ψ are mutually consistent; otherwise it is ψ with as much of ϕ as is consistent with ψ. Thus, if they are inconsistent, ψ overrides ϕ. But this overriding, when it happens, is in general only partial. ψ doesn't override all of ϕ, just those bits which conflict with it. The machinery needed for this is described later in the paper.

OTPs were first described in [Rya91], where they were called 'structured theories'. This paper is self contained, but some technical details and many proofs have been omitted here to leave space for new results. The most complete account of OTPs to date is [Rya92a], copies of which are available from the author.

The remainder of the paper is organised as follows. In §2 we give examples of ordered theory presentations and define their semantics. In §3 we examine some standard examples of default reasoning using OTPs, and in §4 we show the relation with other default systems. Finally, in §5 we show how to use OTPs for belief revision.

2 Ordered theory presentations

An ordered presentation of a theory is a *partially ordered multi-set of sentences*. Sentences lower in the ordering take priority over those above. Earlier we simplified by saying that it was a partially ordered *set*, but we have to consider multi-sets because the same sentence may occur twice, in different places in the order. An informal syntax of graphs for OTPs was used in §1, which is used for much of the paper; a more formal notation is introduced in §2.2.

2.1 Examples and motivation

This section is intended to illustrate by example the intended behaviour of OTPs. The reader can check the examples against his or her intuitions. All of them work out successfully in the formalism described in the paper. While reading these examples, it is important to keep the following points in mind:

1. In an OTP, sentences lower in the ordering take precedence over those above.
2. When a sentence lower in the ordering contradicts a sentence above it in the ordering, the lower sentence overrides the higher one. But in general, this overriding is only partial. The lower sentence does not cancel the effect of the higher one completely.
3. In evaluating an OTP (that is, in working out the theory it presents), the idea is to use as much of the available information as possible but to avoid contradictions.

We take the underlying logic (classical propositional logic or classical predicate logic) as given.

Example 2.1 Here are two OTPs and the ordinary sentences to which they are equivalent.

$$\begin{array}{c} p \wedge q \\ \uparrow \\ \neg p \end{array} \quad \equiv \quad \neg p \wedge q \qquad\qquad \begin{array}{c} p \wedge q \\ \uparrow \\ \neg p \vee \neg q \end{array} \quad \equiv \quad p \leftrightarrow \neg q.$$

In the first case, the OTP consists of the sentences $\neg p$ and $p \wedge q$, but with the former overriding the latter. Thus, $p \wedge q$ is a default relative to $\neg p$. The OTP means that we want $\neg p$ first and foremost, and subject to that, as much of $p \wedge q$ as possible. But $p \wedge q$ conflicts with $\neg p$, so we can't have it all; we can only have the q component. Therefore we get $\neg p$ and q.

In the second case, the default $(p \wedge q)$ is the negation of the given sentence $(\neg p \vee \neg q)$. The overall effect of the OTP is to give us the certain sentence (the $\neg p \vee \neg q$), and then as much of the default as is consistent. Of $p \wedge q$, we can have either p or q but not both. That is why we end up with $p \leftrightarrow \neg q$.

Example 2.2

$$\begin{array}{c} p \\ \uparrow \\ q \\ \uparrow \\ \neg p \vee \neg q \end{array} \qquad \equiv \qquad \neg p \wedge q$$

This is like the second case of example 2.1, except now there is a priority expressed between p and q. This priority is expressed by their location in the ordering. The bottom sentence (the most important) says that we want one of p and q to fail; but subject to that we want q. This gives us $\neg p \wedge q$, since they are consistent. Then, subject to all *that*, we want p. But we've ruled that out by now, so we end up with $\neg p \wedge q$.

Example 2.3 This example will turn out to have importance in §5.

$$\begin{array}{c} p \wedge q \wedge r \\ \uparrow \\ \neg p \vee \neg q \vee \neg r \\ \uparrow \\ (p \leftrightarrow q) \vee \neg r \end{array} \qquad \equiv \qquad (p \leftrightarrow q) \wedge (p \leftrightarrow \neg r)$$

To see this is correct, separate the cases of r and $\neg r$. If r, then we must have $p \leftrightarrow q$ in order to satisfy the most important sentence (the bottom one). To satisfy the next sentence, we must have $\neg p$ or $\neg q$. Since we already have $p \leftrightarrow q$, this means we have $\neg p \wedge \neg q$. Now we have determined the value of all three atoms, for we have $\neg p \wedge \neg q \wedge r$. On the other hand, if $\neg r$ then both the bottom sentence and the middle one are satisfied. We want as much of the top one as possible, which is $p \wedge q$. Therefore, we get $p \wedge q \wedge \neg r$. The presentation is thus equivalent to $(\neg p \wedge \neg q \wedge r) \vee (p \wedge q \wedge \neg r)$, which is elementarily equivalent to $(p \leftrightarrow q) \wedge (p \leftrightarrow \neg r)$.

Example 2.4

$$\begin{array}{c} \forall x.\, p(x) \\ \uparrow \\ \exists x.\, \neg p(x) \end{array} \qquad \equiv \qquad \exists x.\, (\neg p(x) \wedge \forall y.\, (x \neq y \rightarrow p(y)))$$

The more important sentence (the bottom one) says that there is at least one individual which has not got the property p. But, subject to satisfying that, we want to satisfy as much of the upper sentence as possible; it says that all individuals have the property p. We conclude therefore, that precisely one individual fails p; all the others satisfy it. This is stated by the theory on the right.

The examples illustrate the intended behaviour of ordered presentations. Our aim in the next section is to define their semantics formally. We do so in a logically clean way, so that our definitions do not interfere with the mechanism of the underlying logic.

2.2 The semantics of ordered theory presentations

We will define the models of ordered theory presentations. Since the sentences of an OTP are in general inconsistent, we cannot expect its models to satisfy all the sentences. Instead, they should satisfy the lower ones, and then as much of the higher ones as possible. To achieve this we define for each OTP an ordering of the interpretations of the language. This ordering ranks interpretations according to how well they satisfy the sentences of the OTP; and this ranking respects the ordering of the sentences in the OTP. Then models of the OTP are taken to be the maximal interpretations. This strategy of ordering models is well-known in the default reasoning literature [Bes88, McC80, Sho88]

First, it is necessary to have a more formal notation for OTPs than the graphs of the last section. We have seen that an ordered theory presentation is a collection of sentences equipped with a partial order. But to cover the case that the same sentence occurs several times in different places in the presentation, it is necessary to posit a 'carrier set' on which the order is defined and whose points are labelled by sentences.

Definition 2.5 An ordered theory presentation Γ is a tuple $\langle X, \leqslant, F \rangle$ where X is a finite set (the carrier set), $\leqslant$ is a partial order on X, and F is a function mapping X to sentences.

The intuitive meaning of the ordering is: if $x < y$ then the sentence $F(x)$ has greater priority (or more influence) than $F(y)$. This information is used when $F(x)$ and $F(y)$ conflict. We will assume that we are working with a fixed language L over propositional logic or predicate logic with equality[1], with interpretations $\mathcal{M}$ and a satisfaction relation $\Vdash \subseteq \mathcal{M} \times L$ between interpretations of the language and sentences.

As already stated, to define the models of an ordered theory presentations Γ we define an ordering $\sqsubseteq^{\Gamma}$ on interpretations in $\mathcal{M}$ which measures how well an interpretation satisfies Γ. $M \sqsubseteq^{\Gamma} N$ shall mean that N is as good (or better) than M at satisfying Γ. Models of Γ are then taken to be maximal interpretations in this ordering[2]. The definition of the ordering relies on orderings $\sqsubseteq_{\phi}$, one for each sentence ϕ of the language.

[1] In fact, the definitions and results presented here work with other logics, including modal and intuitionistic logics. But in this paper we restrict ourselves to classical logic.

[2] The technique of ordering interpretations which is used in in this paper is well-established in the literature. It originates in McCarthy's first circumscription paper [McC80], and has been generalised in various ways [Sho88, Bes88, KLM90, Vel91, etc.]. In all of those papers, the ordering works in the opposite way to the one we have used for OTPs, that is, $M < N$ means M is better than N; and therefore, one is interested in minimal models. The reader may wonder why we chose to fly in the face of

The relation $\sqsubseteq_\phi$ grades interpretations according to how well they satisfy ϕ. To define $\sqsubseteq_\phi$, it is necessary to define a notion which we call 'natural entailment', written $\stackrel{\simeq}{\models}$. This definition in turn relies on the notion of the *monotonicities* of a sentence. We start therefore with the definition of monotonicities. Then we proceed to the definition of $\stackrel{\simeq}{\models}$, then $\sqsubseteq_\phi$, then $\sqsubseteq^\Gamma$.

The positive monotonicities of a sentence are the predicates whose extension can be increased in any model of the sentence. The negative monotonicities are the predicates which can be decreased in the model. We define this formally as follows:

Definition 2.6 1. In predicate logic, the *extension* of a predicate symbol p in a model is the set of tuples of which p is true in the model. In propositional logic, the extension of a proposition p in a model is a singleton $\{*\}$ if p is true in the model; if p is false, it is $\varnothing$.

2. Extensions are naturally ordered by inclusion. We define $M \leqslant^p N$ if M and N are exactly alike except that the p-extension of M is included in that of N.

3. If $\phi \neq \bot$ then ϕ is *monotonic in p* (written $p \in \phi^+$) if $M \leqslant^p N$ and $M \Vdash \phi$ imply $N \Vdash \phi$. Similarly ϕ is *anti-monotonic in p* ($p \in \phi^-$) if $N \leqslant^p M$ and $M \Vdash \phi$ imply $N \Vdash \phi$. The case that $\phi = \bot$ is handled separately, for technical reasons which will become clear; we define $\bot^+ = \bot^- = \varnothing$.

That is to say, $p \in \phi^{+(-)}$ if increasing (decreasing) the extension of p in any model of ϕ results in another model of ϕ.

Example 2.7 Let $(L, \mathcal{M})$ be classical propositional logic over $\{p, q\}$. For several examples of ϕ, the sets ϕ^+ and ϕ^- are shown in the following table.

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

Example 2.8 Let $(L, \mathcal{M})$ be classical predicate logic over p (unary) and q (binary).

ϕ	ϕ^+	ϕ^-
$\forall x.\, p(x)$	$\{p,q\}$	$\{q\}$
$\exists x.\, p(x)$	$\{p,q\}$	$\{q\}$
$\forall x.\, \exists y.\, q(x,y)$	$\{p,q\}$	$\{p\}$
$\forall x.\, (p(x) \rightarrow \exists y.\, q(x,y))$	$\{q\}$	$\{p\}$
$\forall x.\, \forall y.\, (q(x,y) \rightarrow q(y,z))$	$\{p\}$	$\{p\}$

In classical logic we may characterise $\phi^\pm$ more syntactically, by means of positive and negative occurrences. Recall that *p occurs positively in* ϕ if it occurs in ϕ within the scope of an even number of negation operators, after the operators $\rightarrow$ and $\leftrightarrow$ have been unpacked in terms of their standard definitions. Similarly, if in such circumstances it appears in the scope of an odd number of negation signs then it *occurs negatively*.

Proposition 2.9 $p \in \phi^{+(-)}$ iff ϕ can be written with only positive (negative) occurrences of p.

Notice that the definition is *semantic* in the sense that it is not sensitive to the way ϕ is written. That is, writing $\phi \dashv\models \psi$ if $\phi \models \psi$ and $\psi \models \phi$, we have that $\phi \dashv\models \psi$ implies $\phi^\pm = \psi^\pm$.

Having defined monotonicities, we turn to the definition of natural entailment. Let ϕ and ψ be sentences of L.

Definition 2.10 ϕ *naturally entails* ψ, written $\phi \stackrel{\simeq}{\models} \psi$, if $\phi \models \psi$, and $\phi^+ \subseteq \psi^+$, 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.

Remark 2.11 1. $\stackrel{\simeq}{\models}$ is a reflexive and transitive relation.

2. We have that $p \wedge q \stackrel{\simeq}{\models} p$ and $p \wedge q \stackrel{\simeq}{\models} p \vee q$, but $p \wedge q \not\stackrel{\simeq}{\models} p \leftrightarrow q$ and $p \not\stackrel{\simeq}{\models} p \vee q$. Moreover, $\bot \stackrel{\simeq}{\models} \phi$ for all ϕ. (That was the reason for requiring $\bot^\pm = \varnothing$.) The full picture for natural and ordinary entailment for propositional logic with the predicates p, q is given in figure 2.

this well-established convention, in choosing to order interpretations in the opposite sense and therefore to seek $\sqsubseteq^\Gamma$-*maximal* interpretations. There are two reasons. The first is that the fact that other workers order models in the opposite way is for the historic reason that in circumscription one wants to minimise abnormality predicates; this reason does not apply in the more abstract setting of OTPs. On the contrary, it is more intuitive to move *upwards* in an ordering when one is moving to better and better models. The second reason is that one typically looks at *ascending chains* and *maximal elements* in domain theory and information systems theory, where we see links with our work. Cf. proposition 2.19.

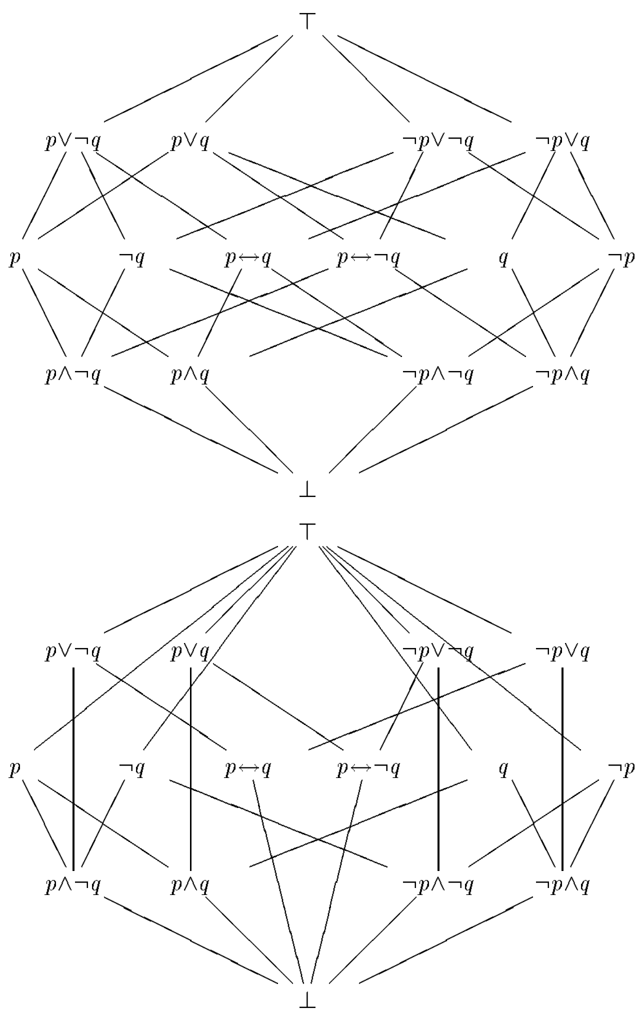

Figure 2: The *ordinary* and *natural* consequence relations over $\{p, q\}$

3. Also, $\forall x.\, p(x) \mathrel{\dot{\models}} \psi$ implies ψ can be written with no negative occurrences of p and no occurrences of any other predicate.

Natural entailment is something like 'relevant entailment'; it stops us adding irrelevant disjuncts in our conclusions. 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

$$\phi \mathrel{=\!\!|\!\!\models} \psi \text{ iff } \phi \mathrel{\dot{=\!\!|}\!\dot{\models}} \psi.$$

Our interest in natural entailment is in order to achieve the definition of $\sqsubseteq_\phi$, to which we now turn. 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 2.12 M satisfies ϕ no worse than N, written $M \sqsubseteq_\phi N$, if for each ψ such that $\phi \mathrel{\dot{\models}} \psi$, $M \Vdash \psi$ implies $N \Vdash \psi$.

Examples will be given shortly. It is easy to verify that

Proposition 2.13 1. $\sqsubseteq_\phi$ is a pre-order, that is to say, it is reflexive and transitive.

2. If $\phi \not\equiv \bot$, the maximal elements of $\sqsubseteq_\phi$ (which are in fact maxi*mum*) are just the models of ϕ.

3. If $\phi \mathrel{=\!\!|\!\!\models} \psi$ then $\sqsubseteq_\phi = \sqsubseteq_\psi$.

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 incomparable extent than some N which perhaps also fails fully to satisfy ϕ. We now define $\sqsubseteq^\Gamma$ in terms of $\sqsubseteq_\phi$ as follows.

Definition 2.14 $M \sqsubseteq^\Gamma N$ if for each $x \in X$, $M \not\sqsubseteq_{F(x)} N$ implies there exists $y \leqslant x$ such that $M \sqsubset_{F(y)} N$.

One can read this as saying: N is as good as M overall $[M \sqsubseteq^\Gamma N]$ if whenever it appears not to be so at a point x $[M \not\sqsubseteq_x N]$ then there is a more important point y $[y \leqslant x]$ where N is doing better than M $[M \sqsubset_y N]$. Informally, the definition says: if things appear to go wrong at a particular x, then they go well at some y in a more important position than x.

Remark 2.15 The definition of $\sqsubseteq^\Gamma$ is more perspicuous if Γ is linear. Let Γ be the third OTP of figure 1. Then $\sqsubseteq^\Gamma$ is the lexicographic combination of the $\sqsubseteq_{\phi_i}$s, i.e.

$$\begin{aligned} M \sqsubseteq^\Gamma N \quad \text{iff} \quad & M \sqsubset_{\phi_1} N \\ & \text{or } (M \sqsubseteq_{\phi_1} N \text{ and } M \sqsubset_{\phi_2} N) \\ & \text{or } (M \sqsubseteq_{\phi_1,\phi_2} N \text{ and } M \sqsubset_{\phi_3} N) \\ & \text{or } \ldots \text{ or} \\ & (M \sqsubseteq_{\phi_1 \cdots \phi_{n-1}} N \text{ and } M \sqsubset_{\phi_n} N). \end{aligned}$$

Proposition 2.16 $\sqsubseteq^\Gamma$ is also a pre-order.

Finally, we define the models of Γ. They are simply the interpretations which are rated maximally by $\sqsubseteq^\Gamma$.

Definition 2.17 $M \Vdash \Gamma$ if M is $\sqsubseteq^{\Gamma}$-maximal.

These definitions represent the guts of the system we propose. Before turning to examples, we present some results.

Proposition 2.18 If ϕ and ψ are mutually consistent then

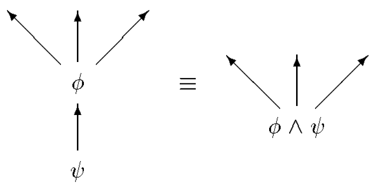

Proposition 2.19 Let Γ be an OTP and $M \in \mathcal{M}$. There exists $N \in \mathcal{M}$ such that $M \sqsubseteq^{\Gamma} N$ and N is $\sqsubseteq^{\Gamma}$-maximal.

Proposition 2.20 If $\Gamma \models \phi$ then $\phi \neq \bot$.

The last of these says that no contradictions may be concluded from any OTP. This may seem surprising, but is really quite rational!

The proofs of these propositions may be found in [Rya92a]. They rely on the compactness of the underlying logic (which in this paper we have assumed is classical propositional or predicate logic).

We now turn to some examples, applying the definitions given so far.

Example 2.21 The working for example 2.3 is given in figure 3. For each sentence ϕ in the OTP, the ordering $\sqsubseteq_{\phi}$ is shown. Then these are combined in the manner of remark 2.15 to yield the final model ordering, whose maximal elements are 001 and 110. The formula with precisely these models is $(p \leftrightarrow q) \wedge (p \leftrightarrow \neg r)$.

In figure 3, we show the ordering on interpretations by means of similar diagrams to the ordering of sentences in OTPs. It is hoped that this is not confusing. If such a diagram has sentences at its nodes, it is an OTP. If it has interpretations at its nodes, it is the diagram corresponding to an ordering $\sqsubseteq_{\phi}$ or $\sqsubseteq^{\Gamma}$ for some sentence ϕ or OTP Γ.

Further examples of $\sqsubseteq_{\phi}$ and $\sqsubseteq^{\Gamma}$ for propositional and predicate logic are given in [Rya92a].

3 Representing defaults in OTPs

We will concentrate on two classic examples, one about inheritance and one about temporal reasoning. To the reader acquainted with default systems they will be very familiar. Although hackneyed, they are excellent examples for showing the key differences between formalisms.

Inheritance example. We will consider the well-known example concerning birds and penguins and whether they can fly. The class of penguins is a subclass of the class of birds. But the property of being able to fly, which holds of birds by default, is not inherited by penguins. In the usual formulation of this example, we have the *factual* premise 'Penguins are birds', together with the *defaults* 'Birds can fly' and 'Penguins cannot fly'. Using predicates to represent the obvious classes, we have $\forall x.\,\big(p(x) \rightarrow b(x)\big)$, $\forall x.\,\big(b(x) \rightarrow f(x)\big)$, and $\forall x.\,\big(p(x) \rightarrow \neg f(x)\big)$.

We want the following results:

1. If Fred is stated to be a bird (whether he is also a penguin or not is not stated), we want to conclude that he can fly.
2. But if it is stated that he is a penguin, we want to conclude that he cannot fly.

The reason this example is interesting is that there are two defaults which compete in certain circumstances. It is easy to get result 1 correctly, but it is in the case of result 2 that the defaults conflict. Our intuition that the second of the two defaults should have priority and block the application of the first is based on the *specificity principle*, which states that *defaults about a specific class of objects take priority over defaults about a more general class.* We use this principle to order the sentences, obtaining for case 1:

$$\forall x.\,\big(b(x) \rightarrow f(x)\big)$$
$$\uparrow$$
$$\forall x.\,\big(p(x) \rightarrow \neg f(x)\big)$$
$$\uparrow$$
$$\forall x.\,\big(p(x) \rightarrow b(x)\big) \wedge b(\text{Fred})$$

This OTP proves $f(\text{Fred})$ as required. The OTP for case 2 of the example has $p(\text{Fred})$ instead of $b(\text{Fred})$, and proves $\neg f(\text{Fred})$.

This shows the fundamental difference between the two approaches to default representation discussed in the Introduction. We have here a set of sentences which we wish to represent, but like much of what is taken to be common knowledge, they conflict with each other. To handle this, we regard some as being weaker than others. The stronger sentences may partially override the weaker ones. These are the basic principles of this second way of representing defaults.

Multiple inheritance The framework of ordered theory presentation is much better suited to inheritance examples than the above analysis indicates. Instead of expressing the fact that penguins are birds by the formula $\forall x.\,(p(x) \rightarrow b(x))$, we construct a specification for birds, and then construct a specification for penguins by stating that they inherit the properties of birds.

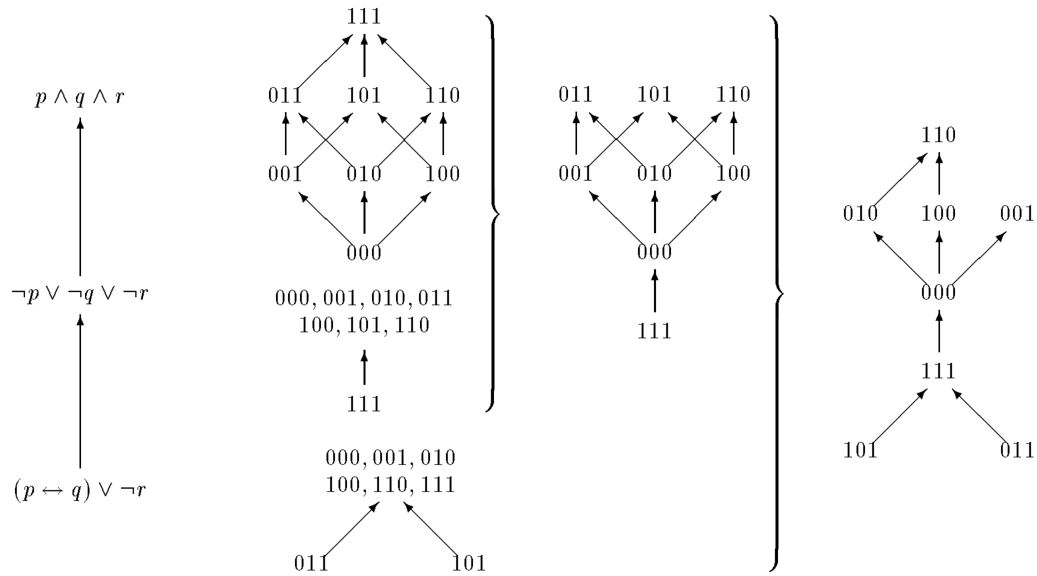

Figure 3: The working for example 2.21

Our specification for birds is

$$f \wedge e$$

meaning that they *f*ly and lay *e*ggs. Now we construct the specification of penguins by inheriting the properties of birds and overriding as necessary:

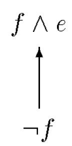

The ordering for the OTP comes straight from the inheritance ordering.

Although our examples of OTPs in this paper have been linear, the definitions allow the sentence ordering to be partial. (Examples of partially ordered theory presentations are given in [Rya92a].) In that way we can handle multiple inheritance. For example, a rendering of the familiar inheritance situation for Nixon is given in the first of the following diagrams (in which the arrows mean 'is-a'), while the second shows the corresponding OTP, in which we have propositions expressing the properties *b*iped, *p*acifist, believer in *G*od, supporter of *B*ush, and being named *N*ixon.

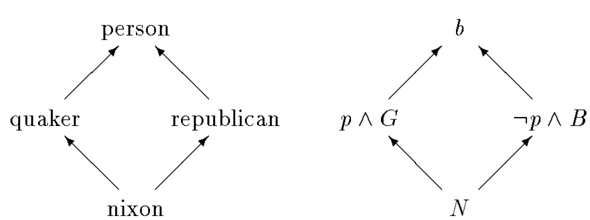

Temporal example. We will describe our solution to the Yale Shooting problem. We assume familiarity with the description of the problem and with the well-known pitfalls in representing the facts and defaults in a way which yields the correct prediction [HM86, Kau86, Sho88, Bak91, Sha92].

At the basis of the problem there are two competing defaults, one expressing the persistence of the loadedness of the gun and the other the persistence of the man's aliveness. The essence of a large class of solutions focusses on the idea, due to Y. Shoham [Sho88], that defaults relating to earlier states of the system should take priority over defaults relating to later states. Thus, we may stipulate a principle for persistence defaults[3], analogous to the specificity principle for inheritance defaults. The *chronology principle* states that: *defaults about an earlier state take priority over defaults about a later state.*

We can use this principle directly in the context of ordered theory presentations. To illustrate this, we shall dramatically simplify the problem by coding it in propositional logic with three states, which we represent by indices on the propositions[4].

[3] It is important to note that this principle is appropriate only when using defaults to predict the outcome of action sequences, i.e. for 'prediction problems'. It is not appropriate for other examples of uses of persistence defaults, such as 'explanation problems' where it is desired to account for a known outcome, in which this principle manifestly gets the wrong answer. An example of this is H. Kautz' 'stolen car problem' [Kau86].

[4] This is not the coding of the example given in Reiter's logic by Hanks and McDermott in the usual paper.

Let the three states be represented by the set $\{1, 2, 3\}$, in which 1 is the result of the loading action, 2 is the result of waiting and 3 results from the shoot action. Let ℓ_i and a_i mean respectively that the gun is loaded and the man is alive in state i. We have the facts

$$\ell_1 \qquad a_1 \qquad \ell_2 \rightarrow \neg a_3$$

and we have the defaults

$$a_i \rightarrow a_{i+1} \qquad \ell_i \rightarrow \ell_{i+1} \qquad (i \in \{1, 2\})$$

Notice that we have not represented the fact that being loaded in a state is an exception to the persistence of alive in the state which follows a shoot action—as was done in the original coding of [HM86]. We do not need to do this, because that fact is represented by the chronology principle, which says that the persistence of earlier fluents shall have priority over the persistence of later ones. We use this to arrive at the ordered theory presentation:

$$\begin{array}{c} a_2 \rightarrow a_3 \\ \uparrow \\ (\ell_1 \rightarrow \ell_2) \wedge \\ (a_1 \rightarrow a_2) \\ \uparrow \\ \ell_1 \wedge a_1 \wedge \\ (\ell_2 \rightarrow \neg a_3) \end{array}$$

This OTP proves $\neg a_3$ as required.

We do not intend to conclude from this analysis that the logic of ordered theory presentations is superior to all the other default systems because it obtains the correct answer to the Yale Shooting Problem. Such a conclusion would be naïve for many reasons. For example, our solution is a crude application of the chronology principle, but, as H. Kautz' stolen car example shows [Kau86], this is not appropriate for all examples of reasoning about actions. We have illustrated that the theory of OTP given in this paper does correctly implement prioritisation of defaults in a natural way which allows for clear knowledge representation. We also hope that we have shown that the representation of defaults, and interacting defaults in particular, is clearer in the theory of OTPs than in many of its rivals.

We have simplified rather dramatically by using a propositional language and making explicit the identities of the states. This simplification is justified since the same problem occurs in this simpler setting as occurred in Hanks and McDermott's, but the simpler setting is rather easier to understand. However, it is true that the simpler setting may not do justice to some of the subtler solutions to the problem which have appeared in the literature. As these are not the main interest of this paper, I feel this is not a significant loss.

4 Relation with other default systems

4.1 Circumscription

PRELIMINARY REPORT

The aim of this topic, which has not yet been achieved, is to provide theorems which show how to translate between OTPs and circumscriptive theories. In this section we give the story so far by means of results, examples and conjectures.

We assume familiarity with the ideas of circumscription [McC80, Lif85], prioritised circumscription [Lif87], and also with *propositional* circumscription. Circumscribing a proposition in a theory means trying to make it false, just as circumscribing a predicate means trying to make its extension as small as possible. It is easy to show that the circumscription of a set of propositions (allowing another set to vary) in a propositional theory is again a propositional theory.

We also adopt the following notation:

- $\mathsf{Circ}_{p;z}(\phi)$ is the circumscription of p in ϕ, allowing z to vary and keeping everything else constant. p and z may be tuples of propositions or predicates.
- $\mathsf{Circ}_p^z(\phi)$ is also the circumscription of p in ϕ, but keeping z fixed and allowing everything else to vary.

It turns out that OTPs translate into circumscriptive theories in which everything is allowed to vary, so we will often be interested in the special case $\mathsf{Circ}_p^{\varnothing}(\phi)$, which we abbreviate to $\mathsf{Circ}_p(\phi)$.

From circumscription to OTPs. The simplest case is $\mathsf{Circ}_p(\phi)$ where p is a single proposition or predicate. The corresponding OTPs are respectively

$$\begin{array}{c} \neg p \\ \uparrow \\ \phi \end{array} \qquad\qquad \begin{array}{c} \forall \underline{x}.\neg p(\underline{x}) \\ \uparrow \\ \phi \end{array}$$

This follows from the facts that:

- If p is a proposition, $M \sqsubseteq_{\neg p} N$ iff $N \Vdash p$ implies $M \Vdash p$.
- If p is a predicate, $M \sqsubseteq_{\forall \underline{x}.\neg p(\underline{x})} N$ iff M, N are isomorphic structures in terms of the functions of the language and (modulo that isomorphism) the p-extension of N is included in that of M.

The parallel circumscription of several propositions or predicates $\mathsf{Circ}_{p_1 \ldots p_n}(\phi)$ is respectively

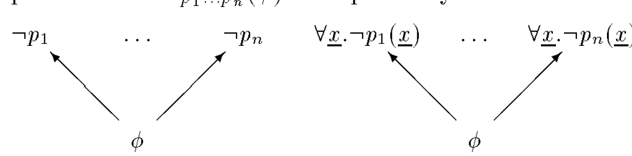

and the prioritised circumscription $\mathsf{Circ}_{p_1 > \ldots > p_n}(\phi)$ becomes

$$\begin{array}{ccc} \neg p_n & \qquad & \forall \underline{x}. \neg p_n(\underline{x}) \\ \uparrow & & \uparrow \\ \vdots & & \vdots \\ \neg p_1 & & \forall \underline{x}. \neg p_1(\underline{x}) \\ \uparrow & & \uparrow \\ \phi & & \phi \end{array}$$

since, as implied by [Lif87, eq. 9], in the context of the given priorities we have

$$\begin{array}{lll} p \sqsubset p' & \text{iff} & p_1 \sqsubset p'_1 \\ & & \text{or } (p_1 \sqsubseteq p'_1 \text{ and } p_2 \sqsubset p'_2) \\ & & \text{or } (p_{1,2} \sqsubseteq p'_{1,2} \text{ and } p_3 \sqsubset p'_3) \\ & & \text{or } \ldots \text{ or} \\ & & (p_{1\ldots n-1} \sqsubseteq p'_{1\ldots n-1} \text{ and } p_n \sqsubset p'_n). \end{array}$$

Compare remark 2.15. Here, $p_2 < p_1$ means that p_1 is circumscribed with a greater priority than p_2, while $p \sqsubseteq p'$ ($p \sqsubset p'$) means that the extension of p is (strictly) included in that of p'.

A corollary of this analysis is the obvious definition of 'partially prioritised circumscription', in which predicates are circumscribed with priorities specified by a partial order.

Definition 4.1 Let $<$ be a partial order on the tuple p of predicates $p_1, \ldots, p_n$. Then $p \sqsubseteq p'$ if for each i with $p_i \not\sqsubseteq p'_i$ there exists j with $p_i \leqslant p_j$ such that $p_j \sqsubset p'_j$.

Compare definition 2.14.

From OTPs to Circumscription. The situation in this direction is rather more complicated. We will consider only the simplest case of

$$\begin{array}{c} \forall x. \phi(x) \\ \uparrow \\ \psi \end{array}$$

It will not be hard to generalise. Let Γ be this OTP. The basic idea is to split $\phi(x)$ into 'components' $\phi_1(x) \wedge \ldots \wedge \phi_n(x)$ and consider the circumscription

$$\mathsf{Circ}_{\mathsf{ab}_1 \ldots \mathsf{ab}_n}(\psi \wedge \bigwedge_i (\neg \phi_i(x) \rightarrow \mathsf{ab}_i(x)))$$

But what constraints should there be on the way ϕ is split into components? As a minimum one would expect to require *conjunctive normal form*, but this is not good enough, as we now show.

Example 4.2 Suppose $\phi = p \wedge q$. There are several normal forms, two such being $p \wedge q$ itself and $p \wedge (\neg p \vee q)$. We have

$$\Gamma \equiv \mathsf{Circ}_{\mathsf{ab}_1, \mathsf{ab}_2}(\psi \wedge (p \rightarrow \mathsf{ab}_1) \wedge (q \rightarrow \mathsf{ab}_2))$$

whatever ψ may be, but we also have that

$$\Gamma \not\equiv \mathsf{Circ}_{\mathsf{ab}_1, \mathsf{ab}_2}(\psi \wedge (p \rightarrow \mathsf{ab}_1) \wedge (\neg p \vee q \rightarrow \mathsf{ab}_2)).$$

For example, set $\psi = \neg p$. Then $\Gamma \equiv \neg p \wedge q$, but the right-hand side is simply $\neg p$.

The example shows that we probably also require that the split of ϕ must not contain unnecessary occurrences of propositions or predicates of the 'wrong' polarity. It is a consequence of Lyndon's Theorem [CK90] that:

Proposition 4.3 For any formula ϕ there is an equivalent formula ϕ' such that every predicate occurring positively (negatively) in ϕ' occurs positively (negatively) in every formula equivalent to ϕ.

In other words, we can find for every ϕ an equivalent ϕ' with no eliminable occurrences of predicates. We conjecture that this is the normal form we seek.

4.2 Formal properties of OTPs

The study of default systems has been transformed by a new concern, namely the formal properties of the generated consequence relation. The first default systems introduced in the 1980 special issue of *Artificial Intelligence* [AIJ80] did not even have well-defined consequence relations. Gabbay and Clark [Gab91, CG88] first observed that, instead of focussing on the *negative* properties of such consequence relations, that is, their *non*-monotonicity, one should instead ask what positive properties they have. They gave the name 'cautious monotonicity' to the property

$$\frac{\Phi \mid\!\sim \phi \quad \Phi \mid\!\sim \psi}{\Phi, \phi \mid\!\sim \psi}.$$

This property, which is weaker than full monotonicity, has become widely accepted as a desirable property for default systems.

The story of the properties of default consequence relations has been pursued in the work of Kraus/Lehmann/Magidor [KLM90, Leh89] and also by Makinson [Mak88, Mak92]. Makinson's [Mak92] is, in my opinion, the most authoritative and systematic study to date. He describes and motivates a set of conditions on a default consequence relation and analyses existing systems according to whether they have the conditions. In this section we outline his principal conditions and check the theory of OTPs of this paper against them.

In Makinson's work, the expression $\Phi \mid\!\sim \psi$ should be read as: ψ follows from Φ *in the context of an understood set of defaults*. It is unfortunate (and detracts slightly from Makinson's systematic study) that these defaults are nowhere made explicit. Consequently, the behaviour of the consequence relation under variations of the defaults—and for that matter, questions of default representation—are not examined at all in his work.

Makinson's conditions also refer to classical consequence, written $\Vdash$. $\Phi \Vdash \psi$ is to be read as ψ follows from Φ without using the defaults. The understood set of defaults can be thought of as augmenting classical consequence to default consequence. Therefore, the first property we may expect is

Supraclassicality:

$$\frac{\Phi \Vdash \psi}{\Phi \mid\!\sim \psi}.$$

It says that anything which can be derived without the defaults can also be derived with them.

The next three conditions are together called '**cumulativity**'. The first is simply

Inclusion: if $\psi \in \Phi$ then $\Phi \mid\!\sim \psi$.

The next two are weak forms of the standard Tarski conditions of cut and monotonicity:

Cautious monotonicity:

$$\frac{\Phi \mid\!\sim \phi, \text{ for all } \phi \in \Psi \quad \Phi \mid\!\sim \psi}{\Phi, \Psi \mid\!\sim \psi}$$

Weak cut:

$$\frac{\Phi \mid\!\sim \phi, \text{ for all } \phi \in \Psi \quad \Phi, \Psi \mid\!\sim \psi}{\Phi \mid\!\sim \psi}$$

For the justification of these principles in intuitive terms, we cannot do better than quote Makinson. "Cut may be seen as expressing a determination not to allow the length, intricacy or manner of a derivation of a conclusion to reduce the freedom with which it is used in further inference. There is no 'diminution of usability' with respect to distance from origins. Once inferred, a proposition may be called upon in conjunction with the original information, unless genuinely new (i.e. uninferable) information is also added. Cautious monotonicity, on the other hand, may be seen as expressing a certain irreversibility in the drawing of conclusions. Once inferred, a proposition may be retained irrespective of what other inferred propositions are added to the stock of usable information. We need never go back unless, once more, genuinely new information is brought in" [Mak92].

The next condition we will consider is

Distributivity: If Φ and Ψ are $\Vdash$-closed sets of sentences (that is, $\Phi \Vdash \phi$ implies $\phi \in \Phi$, and similarly for Ψ) then

$$\frac{\Phi \mid\!\sim \phi \quad \Psi \mid\!\sim \phi}{\Phi \cap \Psi \mid\!\sim \phi}.$$

Makinson considers other conditions, but these are the principal ones.

We have already noted that Makinson's conditions make no reference to the set of defaults which are implicit in the relation $\mid\!\sim$. On the other hand, one of the attractive features of the framework of Ordered Theory Presentations as a default system is that there is *no difference* between defaults on one hand and 'sure facts' or facts on the other, except the priority they are given in the ordering. We view this as a desirable feature since we believe that, philosophically, the so-called sure facts and the defaults have the same provenance. They should all form part of the theory from which we make deductions. A sentence does not have the status of a default in isolation, but only in relation to other sentences; to be precise, it is a default relative to those sentences which can override it.

Nevertheless, we can go quite some way in examining Makinson's conditions in the context of ordered theory presentations over classical logic. In order to emulate variation of the facts with a fixed set of defaults, we can consider the consequences of the following ordered presentation with Δ fixed and Φ varying:

$$\begin{array}{c}(\Delta)\\ \uparrow\\ \Phi\end{array}$$

This is the OTP Δ with Φ appended at the bottom, which we will write as $\Delta * \Phi$ until the end of this section. We can think of this OTP as a way of representing that which in other default formalisms might be called 'the theory Φ with defaults Δ'. Notice that Δ is itself an OTP; that is, we are still allowing defaults with different priorities. Using this idea we can define a consequence relation $\mid\!\sim$ which embodies the defaults, as in Makinson's work. The obvious thing to do is to let $\Phi \mid\!\sim \psi$ mean $\Delta * \Phi \models \psi$. However, we know from proposition 2.20 that $\perp$ does not have its classical behaviour in the context of OTPs. We can get improved results by setting:

Definition 4.4 $\Phi \mid\!\sim \psi$ if $\bigwedge \Phi = \perp$ or $\Delta * \Phi \models \psi$.

That is to say, if Φ is contradictory then it entails everything; otherwise, it entails just what the illustrated OTP entails.

Recall that the technique of model ordering which originates in McCarthy's first circumscription paper [McC80] has been generalised in various ways [Sho88, Bes88, KLM90, Vel91, etc.]. It is further generalised by Makinson in [Mak92], where he proves that preferential model structures which satisfy a condition which he calls 'stopperedness' generate inference relations which satisfy each of the conditions on inference relations defined above. We therefore need simply show that the relation $\mid\!\sim$ of definition 4.4 *is* such a relation to prove that

Proposition 4.5 $\mid\!\sim$ satisfies supraclassicality, inclusion, cumulativity, and distributivity.

This is indeed the case, the condition of stopperedness following from our proposition 2.19. As before, the full proofs are spelled out in [Rya92a]. We thus have shown that OTPs over classical logic can yield a default inference relation in the sense of Makinson, with good formal properties.

5 Belief Revision

Ordered theory presentations have significant application in belief revision, whose basic question is: how should new information be incorporated into a belief state to result in a belief state which contains the new information and as much of the original belief state as is consistentΓ The best-known work on this subject is called the AGM theory (after its originators, C. Alchourrón, P. Gärdenfors and D. Makinson) [Gär88].

A full account of the AGM theory and the belief revision functions obtained from ordered theory presentations is given in [Rya92b] and [Rya92a]. We summarise our main findings below, but the interested reader should consult the more expansive references.

The AGM theory represents belief states as deductively-closed sets of sentences. Let K be such a belief state and ϕ a sentence. The revision of K by ϕ is written $K*\phi$. The AGM theory sets out eight postulates which a belief revision function $*$ must satisfy, known as K1–K8. (These may be found in any of the standard references; we repeat them in a generalised form below.)

We argue, however, that the eight axioms are neither *sound* not *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 is to give 'counterexamples' to K4 and K8, which are given later in the paper. My argument against completeness is the following proposition, which shows that K1–8 admit revision functions which fail to preserve any of the original belief state in many cases.

Proposition 5.1 The revision function

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

satisfies axioms K1–8.

In addition to this undesirable property of the AGM system, there is the further fact of that system that one cannot perform revision more than once. Repeated or iterated revision is not constrained by the axioms, and none of the models proposed for the AGM axioms (like revision by selection functions and epistemic entrenchment [Gär88]) define it[5].

Before considering how belief revision works in the context of OTPs, we have to generalise the AGM axioms. As things stand, they 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.

K1 $\mathcal{K}*\phi$ is a belief state;
K2 $\phi\in|\mathcal{K}*\phi|$;
K3 $|\mathcal{K}*\phi|\subseteq|\mathcal{K}|+\phi$;
K4 If $\neg\phi\notin|\mathcal{K}|$ then $|\mathcal{K}|+\phi\subseteq|\mathcal{K}*\phi|$;
K5 $\mathcal{K}*\phi$ is contradictory implies $\phi=\bot$;
K6 If $\models\phi\leftrightarrow\psi$ then $|\mathcal{K}*\phi|=|\mathcal{K}*\psi|$;
K7 $|\mathcal{K}*(\phi\wedge\psi)|\subseteq|\mathcal{K}*\phi|+\psi$;
K8 If $\neg\psi\notin|\mathcal{K}*\phi|$ then $|\mathcal{K}*\phi|+\psi\subseteq|\mathcal{K}*(\phi\wedge\psi)|$.

We now turn to the belief revision theory offered by the OTP framework. We define

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

As belief revision gives rise only to *linear* OTPs we can write them with a more succinct notation. The OTP of example 2.2 will be written $[p,q,\neg p\vee\neg q]$.

Revision on these belief states is defined as follows:

$$\Gamma*\phi = \begin{cases} \bot & \text{if } \phi=\bot \\ [\phi] & \text{if } \phi\neq\bot \text{ and } \Gamma=\bot \\ \Gamma \text{ appended with } \phi & \text{otherwise} \end{cases}$$

The general case, therefore, is that we simply append the revising sentence. In other words, belief states are (usually) just revision histories.

[5] For the expert reader, we remark that 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 [Rot, Sch91]. 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 [Han91].

We argue that this function performs intuitively correct belief revision. As well as allowing repeated revision, it has the 'persistence' requirement mentioned above. However, this belief revision function, while satisfying $\mathcal{K}1$, $\mathcal{K}2$, $\mathcal{K}3$, $\mathcal{K}5$, $\mathcal{K}6$ and $\mathcal{K}7$, fails to satisfy $\mathcal{K}4$ and $\mathcal{K}8$. The counterexample to these two is given by example 2.3 and figure 3 of this paper, by setting $\mathcal{K} = [p \wedge q \wedge r, \neg p \vee \neg q \vee \neg r]$ and $\phi = (p \leftrightarrow q) \vee \neg r$ for K4; and $\mathcal{K} = [p \wedge q \wedge r]$, $\phi = \neg p \vee \neg q \vee \neg r$ and $\psi = (p \leftrightarrow q) \vee \neg r$ for K4. An explanation in both technical and intuitive terms of this counterexample may be found in the references already cited.

Acknowledgements

I am very grateful to Murray Shanahan for useful discussions, particularly concerning the relationships with Circumscription.

References

[AIJ80] Artificial Intelligence. Special Issue on Non-Monotonic Logic, volume 13, 1980.

[Bak91] A. B. Baker. Nonmonotonic reasoning in the framework of situation calculus. *Artificial Intelligence*, 49:5–23, 1991.

[Bes88] P. Besnard. The preferential-models approach to non-monotonic logics. In P. Smets, A. Mamdani, D. Dubois, and H. Prade, editors, *Non-standard Logics for Automated Reasoning.* Academic Press, 1988.

[Bib85] W. Bibel. Methods of automated reasoning. In J. Bibel, editor, *Fundamentals in Artificial Intelligence.* Lecture Notes in Computer Science 232, Springer Verlag, 1985.

[Bre89] G. Brewka. Preferred subtheories: An extended logical framework for default reasoning. In *Proc. International Joint Conf. on Artificial Intelligence (IJCAI)*, pages 1043–1048. Morgan Kaufmann, 1989.

[CG88] M. R. B. Clarke and D. M. Gabbay. 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.

[CK90] C. C. Chang and H. K. Keisler. *Model Theory.* North-Holland, third edition, 1990.

[Gab91] D. M. Gabbay. Theoretical foundations for non-monotonic reasoning. part 2: Structured non-monotonic theories. In *Proc. Third Scandanavian Conference on Artificial Intelligence (SCAI'91)*, 1991.

[Gär88] P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States.* MIT Press, 1988.

[Han91] S. O. Hansson. *Belief Base Dynamics.* PhD thesis, Department of Philosophy, Uppsala University, 1991.

[HM86] S. Hanks and D. McDermott. Default reasoning, non-monotonic logics and the frame problem. In *Proc. Fifth National Conference on Artificial Intelligence (AAAI)*, pages 328–333, 1986.

[Kau86] H. Kautz. The logic of persistence. In *Proc. Fifth National Conference on Artificial Intelligence*, pages 401–405, 1986.

[KLM90] S. Kraus, D. Lehmann, and M. Magidor. Non-monotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.

[Leh89] D. Lehmann. What does a conditional knowledge base entail? In *Proc. First International Conference on Principles of Knowledge Representation and Reasoning (KR'89).* Morgan Kaufmann, 1989.

[Lif85] V. Lifschitz. Computing circumscription. In *Ninth International Joint Conference on Artificial Intelligence*, pages 121–127, 1985.

[Lif87] V. Lifschitz. Pointwise circumscription. In M. L. Ginsberg, editor, *Readings in Non-monotonic Logic.* Morgan Kaufmann, 1987.

[Mak88] D. Makinson. General theory of cumulative inference. In M. Reinfrank, J. de Kleer, and M. L. Ginsberg, editors, *Non-monotonic Reasoning.* Lecture Notes in Artificial Intelligence 346, Springer-Verlag, 1988.

[Mak92] 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.

[McC80] J. McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.

[Poo88] D. Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36:27–47, 1988.

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

[Rya91] M. D. Ryan. Defaults and revision in structured theories. In *Proc. Sixth IEEE Symposium on Logic in Computer Science (LICS)*, pages 362–373, 1991.

[Rya92a] M. Ryan. *Ordered Presentations of Theories: Default Reasoning and Belief Revision.* PhD thesis, Department of Computing, Imperial College, 1992. Copies available from author.

[Rya92b] M. D. Ryan. Belief revision and ordered theory presentations. In P. Dekker and M. Stokhof, editors, *Proc. Eighth Amsterdam Colloquium on Logic*, 1992.

[Sch91] K. Schlechta. Some results on theory revision. In A. Fuhrmann and M. Morreau, editors, *The Logic of Theory Change.* Lecture Notes in Artificial Intelligence 465, Springer Verlag, 1991.

[Sha92] M. Shanahan. A circumscriptive calculus of events. Technical report, Imperial College Department of Computing, 1992.

[Sho88] Y. Shoham. *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence.* MIT Press, 1988.

[Vel91] F. Veltman. Defaults in update semantics. Technical Report LP-91-02, Institute for Language, Logic and Information, Amsterdam, 1991.