

Allowed Arguments

Michael Morreau

Philosophy Department & UMIACS
University of Maryland at College Park
College Park, MD 20742-7615
U.S.A.

Abstract

Starting with a first-order modal conditional logic which allows unlimited nesting of defaults and embeddings into any context analyzable in possible-worlds theory, I introduce two simple notions of default reasoning: the syntactic notion of (prioritized) Allowed Consequence and the semantic notion of (prioritized) Allowed Entailment. I prove that the one is sound and complete relative to the other.

1 Introduction

The better known accounts of defeasible reasoning are somewhat limited in the expressive power of their object languages. This is apparent for example when it comes to formalizing Raz's [1990] *prima facie* reasons to act, which underlie the planning and justification of acts. These suggest an analysis in a formalism for expressing defeasible knowledge, such as Reiter's [1980] Default Logic: that p is a *prima facie* reason to do a can be thought of as a 'normal' default rule with the reason p as its prerequisite and as its consequent a sentence expressing that a ought to be done. Such an analysis runs into a difficulty though. In some cases not discussed by Raz the very fact that there is a reason to act is itself, in turn, a reason to act. Thus Jones' lateness is a reason for his boss to fire him; and this very fact is in turn a reason for Jones to act: it is a reason for him to arrive on time in the future (assuming of course he doesn't want to be fired). This second claim cannot be represented in way I've suggested though, since a default rule which represented it would have, as its prerequisite, a default rule expressing that Jones' lateness is a *prima facie* reason for him to be fired. Such *embedded* defaults are not within the scope of Reiter's formalism since its default rules are not a part of its object language. They are metalinguistic rules of inference.

McCarthy's [1980] Circumscription can in principle be used to represent embedded defaults, since it represents defaults in the object language. But in order to represent the

notion that an act *ought* to be done, which *prima facie* reasons seem to require, this formalism will have to be generalized from classical logic to some more expressive underlying language.

It is unlikely that Default Logic can be extended to cover embedded defaults; Circumscription, on the other hand, can certainly be extended to more expressive languages (Thomason [1990] extends it to Montague's Intensional Logic). But you wonder whether it is not better to start off with a more expressive language, formalizing defaults as *ceteris paribus* conditionals in modal logic, since modal conditional logic allows unlimited embeddings of conditionals; also, propositional attitudes, deontic contexts, counterfactuals and other constructions are already quite well understood in that setting.

This idea was first pursued by Delgrande [1988]. Others including Asher and Morreau [1991] and Boutilier [1992] later took it up, but their results so far have been disappointing. Delgrande had a difficulty which Circumscription and Default Logic do not face with what he called "irrelevant information." Asher & Morreau [1991] solved this particular problem, but their nonmonotonic consequence notion, Commonsense Entailment, is conceptually exotic and technically formidable. Boutilier [1992] treated conditionals expressing defaults as necessarily true, if true at all; this squanders the resources of modal conditional logic for representing nested defaults and embeddings of defaults into belief contexts and counterfactuals. Finally, these approaches have not led to attractive syntactic characterizations of nonmonotonic reasoning. Not even characterizations of the sound-though-incomplete kind which McCarthy [1980] offers.

In this paper I will present an account of nonmonotonic reasoning based in modal conditional logic which does not have these shortcomings. The goal is the same as that of Asher and Morreau [1991]: to provide an underpinning for analyses requiring the resources of possible-worlds theory. In regard to the inferences it allows, the notion of nonmonotonic consequence which I will present here is very similar to their Commonsense Entailment (though section 7 shows how it allows defeasible inferences including *strengthening of conditional antecedents* and the *hypothetical syllogism* which have as far as I know not

Research supported in part by the NSF and ARL.

been dealt with before). One important advantage of the present account over Commonsense Entailment is that it is technically and conceptually simpler. For this reason it provides a better foundation for analyses which depend on Commonsense Entailment, such as Lascarides and Asher's [1993] treatment of rhetorical structure in natural language.

I begin in the next section with a weak first-order modal conditional logic with which to represent defaults.¹ In section 3 I present a syntactic notion of nonmonotonic consequence, *Allowed Consequence*, and analyze several familiar patterns of nonmonotonic reasoning. In section 4 I present *Allowed Entailment* which is an analog, in possible-worlds theory, of McCarthy's [1980] Minimal Entailment. In sections 5 and 6 I generalize these to allow a priority order on defaults and show that (prioritized) Allowed Consequence is sound and complete relative to (prioritized) Allowed Entailment.

2 Default Conditionals

Here is a formal language for expressing defaults. Let L be a countable language of first-order logic with no function symbols. Add a binary conditional operator $>$. The result is $L_>$; where ϕ and ψ are formulas of $L_>$, so is $\phi > \psi$. Familiar monadic modal operators expressing belief (B), knowledge (K) and obligation (O) can be added too; where ϕ is a formula of $L_>$ so are $B\phi$, $K\phi$, and $O\phi$.

Where ϕ and ψ are sentences, $\phi > \psi$ is intended to express that normally ψ , if ϕ . Letting j be an individual constant standing for Jones, for example, letting Dyz express that y is a day in month z , letting Lxy express that x is late for work on day y and letting Fx express that x is fired, $\forall y(D(y, \text{december}) \rightarrow L(j, y)) > OFj$ is intended to express that other things being equal, Jones ought to be fired if he is late every day in december. This formalizes the notion that Jones' being late all december is a *prima facie* reason for him to be fired.

Where ϕ and ψ are predicates or relations, $\forall x(\phi > \psi)$ is intended to express that if something is ϕ then normally it is a ψ . For example, letting Hxy express that y is an event in which x is thoroughly heated, and letting Fxy express that y is an event in which x will catch fire, the monadic predicate $\forall y(Hxy > Fxy)$ expresses that x is *flammable*: it is the sort of thing which in the event that it is thoroughly heated will normally catch fire.

There is something implausible about understanding flammability as a universally quantified conditional: that a piece of wood is flammable wrongly entails that it normally will catch fire in the event that it is thoroughly heated while, say, immersed in water. (Imagine boiling it for hours on end.) There are alternative ways of understanding flammability: that a given piece of wood w is flammable clearly has something to do with events in which w is thoroughly heated. But instead of supposing this claim to quantify over events of this kind it could, say, be taken to express that two properties — the property of being an event in which w is heated, and that of being an

event in which w catches fire — stand in an appropriate relation to one another. The generic *dogs bark* would on this view express that the property of being a dog and being a barking thing stand in this same relation, and so on.

Such a relational account of defaults can be developed on the basis of a suitable theory of properties. And once developed it will be possible to extend to it the account of nonmonotonic reasoning of sections 3 and 4. But that is for another time.

Unlimited nesting of $>$ is allowed. Letting Ax express that x attempts to arrive at work on time (more of the resources of modal logic could be used to analyze A further but that would be a distraction here), the following nested sentence expresses that the fact that Jones' lateness is a reason to fire him is itself a *prima facie* reason for him to try to arrive on time:

$$(\forall y(D(y, \text{december}) \rightarrow L(j, y)) > OFj) > OAj.$$

For another example of nesting, letting Rx express that x ought to be kept well away from radiators we can write $\forall x(\forall y(Hxy > Fxy) > Rx)$ to express that flammable things normally ought to be kept well away from radiators.

What is the meaning of $>$ and what is its logic? As Delgrande, Asher & Morreau and Boutilier all point out, in keeping with the intended interpretation of $\phi > \psi$ there is a standard axiom scheme of conditional logic which is not valid: *modus ponens*, or $\phi > \psi \rightarrow \phi \rightarrow \psi$. It can happen that normally ψ if ϕ , that ϕ holds, and yet that ψ does not hold. These authors helped themselves to interpretations of conditional logic within possible-worlds semantics, as pioneered by Stalnaker and Thomason [1970] and Lewis [1973], making suitable adjustments so as to keep *modus ponens* from being valid. I will do the same.

I begin by interpreting $L_>$ in possible-worlds models (which are presented in Chellas [1980]). The basic notion is that of a *possible-worlds frame*, which is a triple $\langle \mathcal{W}, \Delta, * \rangle$ consisting of a nonempty set $\mathcal{W}$ of possible worlds, a set Δ of individuals over which the quantifiers of $L_>$ range, and a *worlds-selection function* $*$: $\mathcal{W} \times \wp(\mathcal{W}) \rightarrow \wp(\mathcal{W})$. This latter component, which is used to interpret $>$, maps each world and proposition onto a proposition. (*Proposition* is the technical term for a subset of $\mathcal{W}$; a proposition p holds in a possible world w just in case $w \in p$.) Thus $*(w, p)$ is a set of worlds; informally speaking, they are the worlds where things are as they normally are if p holds. Just which set this is depends on w , since how things normally are if p can vary from time to time and place to place. Additional structure can be added to modal frames in order to interpret B, K, O and other modal constructions like counterfactuals; what is needed is quite well known however and I will not discuss it here. See Chellas [1980] for details.

In keeping with its intended purpose $*$ is subject to a very basic constraint of conditional logic:

FACTICITY: A frame $\langle \mathcal{W}, \Delta, * \rangle$ satisfies facticity if

for all $w \in \mathcal{W}$ and $p \subseteq \mathcal{W}$, $*(w, p) \subseteq p$.


Another familiar constraint is however quite inappropriate in this connection. Say that $w \in \mathcal{W}$ is *centered* in a frame $\langle \mathcal{W}, * \rangle$ if for all $p \subseteq \mathcal{W}$, if $w \in p$ then $w \in *(w, p)$. Then we have:

¹ It is neither new nor is it unique to this analysis of nonmonotonic reasoning; Delgrande, Asher & Morreau and Boutilier all rely on some such underlying conditional logic.


CENTERING: A frame $\langle \mathcal{W}, \Delta, * \rangle$ is centered if each of its possible worlds is centered.

We cannot require that all frames be centered since that w is a p -world in no way guarantees that w is a world where everything is as things normally are if p . Along with Delgrande, Asher and Morreau and Boutilier I therefore do not require that modal frames satisfy this constraint. The following examples of frames will be used in later examples:

EXAMPLE 1: The following singleton structure satisfies facticity and is centered:



EXAMPLE 2: The following structure satisfies facticity and is not centered:



Together with any given domain of quantification, each of these structures forms a modal frame.

A modal frame becomes a model $\langle \mathcal{W}, \Delta, *, I \rangle$ for $L_{>}$ on adding an interpretation function I , which in any given possible world assigns elements of Δ to individual constants, and which assigns appropriate intensions to the atomic predicates of $L_{>}$. The conditions of truth of formulas of $L_{>}$ in possible worlds $w \in \mathcal{W}$ of a model $\mathcal{M} = \langle \mathcal{W}, \Delta, *, I \rangle$ are stated relative to assignments a of $L_{>}$ -variables to Δ . The clauses of the truth definition which deal with atomic formulas, sentential connectives apart from $>$ and the quantifiers are all quite familiar; I have not included them here. The clause for $>$ is the standard one in modal conditional logic:

$\mathcal{M}, w \models \phi > \psi [a]$ if and only if $*(w, \models \phi | a) \subseteq \models \psi | a$.

In this expression $\models \phi | a$, the proposition expressed by ϕ in $\mathcal{M}$ relative to a , is just $\{v \in \mathcal{W} : \mathcal{M}, v \models \phi [a]\}$. What is required for a sentence $\phi > \psi$ to be true, then, is that ψ holds in any world where things are as they normally are if ϕ .

EXAMPLE 3: Choosing $\Delta = \{e, w\}$, let $\mathcal{M}$ be obtained by adding, to the singleton structure of EXAMPLE 1, any interpretation function I such that $I_u(e) = e$; $I_u(w) = w$, and $I_u(H) = I_u(F) = \{(w, e)\}$. It is straightforward to verify that $\mathcal{M}, u \models \forall y(Hwy > Fwy)$.

OBSERVATION 4: Let u be centered in $\mathcal{M}$ (as it is in the previous example). Then for any ϕ and ψ , $\mathcal{M}, u \models \phi > \psi \rightarrow \phi \rightarrow \psi$. This follows directly from the truth-conditions for $>$.

These models generate a notion of logical entailment in the standard way. Where Γ is a set of formulas, $\Gamma \models \phi$ means that for any model $\mathcal{M}$, for any possible world w of $\mathcal{M}$ and for any appropriate variable assignment a , if $\mathcal{M}, w \models \Gamma[a]$ then $\mathcal{M}, w \models \phi[a]$. The following example concerns a model showing that *modus ponens* is not valid. That is to say, $\not\models \phi > \psi \rightarrow \phi \rightarrow \psi$:

EXAMPLE 5: Choosing $\Delta = \{e, w\}$, let $\mathcal{N}$ be obtained by adding, to the structure of EXAMPLE 2, any interpretation function I such that $I_u(e) = I_v(e) = e$; $I_u(w) = I_v(w) = w$,

$I_u(F) = \{\}$ and $I_u(H) = I_v(H) = I_v(F) = \{(w, e)\}$. It is straightforward to verify that $\mathcal{N}, u \models \forall y(Hwy > Fwy)$, that $\mathcal{N}, u \models Hwe$, and that $\mathcal{N}, u \not\models Fwe$.

Logical entailment can be given a sound and complete syntactic characterization. I start with a notion of derivability without premises. Let $\vdash \phi$ mean that ϕ can be derived using the following axioms and rules:

- A1. All truth-functional tautologies of $L_{>}$
- A2. $\forall x \phi \rightarrow (t/x)\phi$ where t is any individual term of $L_{>}$
- A3. $\forall x \phi \leftrightarrow \neg \exists x \neg \phi$
- A4. $\forall x(\phi \rightarrow \psi) \rightarrow (\exists x \phi \rightarrow \psi)$
where x is not a free variable of ψ
- A5. $\forall x(\phi > \psi) \rightarrow (\phi > \forall x \psi)$
where x is not a free variable of ϕ
- A6. $\forall x(\phi > \phi)$
- R1. If $\vdash \phi$ and $\vdash \phi \rightarrow \psi$ then $\vdash \psi$.
- R2. If $\vdash (\psi_1 \wedge \psi_2 \dots \wedge \psi_n) \rightarrow \psi$
then $\vdash (\phi > \psi_1 \wedge \phi > \psi_2 \dots \wedge \phi > \psi_n) \rightarrow \phi > \psi$
- R3. If $\vdash \phi \rightarrow (c/x)\psi$ then $\vdash \phi \rightarrow \forall x \psi$
provided the constant c occurs neither in ϕ nor in ψ
- R4. If $\vdash \phi$ then $\vdash (t/x)\phi$
provided the individual term t does not occur in ϕ
- R5. If $\vdash \phi \leftrightarrow \psi$ then $\vdash (\phi/\psi)\chi \leftrightarrow \chi$
where $(\phi/\psi)\chi$ is the result of replacing within χ an occurrence of ψ by an occurrence of ϕ .

$\Gamma \vdash \phi$, or ϕ is derivable from premises Γ , can be defined in terms of the above notion. It means that for some finite $\Gamma^* \subseteq \Gamma$, $\vdash \wedge \Gamma^* \rightarrow \phi$. (This is not the standard definition but it has the advantage that compactness and the deduction theorem involving $\rightarrow$ are obvious.) The following completeness theorem can be proven using a standard, Henkin-style construction of a canonical model:

THEOREM 6: $\Gamma \vdash \phi$ if and only if $\Gamma \models \phi$.

Proof: A variation on a result in Chellas [1980].

3 Allowed Consequence...

In this section the logic of weak conditionals will be extended to a nonmonotonic syntactic consequence notion, written $\vdash_{\mu}$. Informally speaking, the idea is to add *modus ponens* wherever possible, thus allowing the consequents of conditionals to be detached. Technically this is done by adding to the premises of an argument as many instances of the axiom scheme $\phi > \psi \rightarrow \phi \rightarrow \psi$ as are consistent with them. Let MP be the set of all instances of this scheme.

DEFINITION 7: Ω is maximal Γ -consistent within MP iff

- (i) $\Omega \subseteq \text{MP}$, and (ii) $\Gamma \cup \Omega \not\vdash \perp$, while for every Ω^* , if $\Omega \subseteq \Omega^* \subseteq \text{MP}$ and $\Gamma \cup \Omega^* \not\vdash \perp$, then $\Omega = \Omega^*$.

Such sets play a role in the notion of allowed consequence which is analogous to that of extensions in Reiter's Default Logic.

DEFINITION 8: $\Gamma \vdash_{\mu} \phi$ iff for every maximal Γ -consistent set Ω within MP, $\Gamma \cup \Omega \vdash \phi$.

Where $\Gamma \vdash_{\mu} \phi$ I will say that Γ allows the consequence ϕ to be drawn. The following two examples show that interesting instances of *modus ponens* across $>$ are allowed. First, that e is an event in which w is thoroughly heated,

together with the fact that w is flammable, allows the consequence w will catch fire to be drawn:

EXAMPLE 9: $Hw, \forall y(Hwy > Fwy) \vdash_{\mu} Fwe$.

Proof: Let the set of the premises $Hw, \forall y(Hwy > Fwy)$ be called Γ . From EXAMPLE 3 and OBSERVATION 4 it follows that Γ is satisfiable together with all of MP, so (with THEOREM 6) there is just a single maximal Γ -consistent subset of MP, and that is MP itself. Now $\Gamma, MP \vdash Fwe$. ■

The next example is similar but concerns the part of the language in which $>$ is nested. Intuitively speaking, from the facts that w is flammable and that flammable things normally ought to be kept well away from radiators follows that w ought to be kept well away from radiators:

EXAMPLE 10:

$\forall y(Hwy > Fwy), \forall x(\forall y(Hxy > Fxy) > Rx) \vdash_{\mu} Rw$.

Proof: Let the set of the premises $\forall y(Hwy > Fwy), \forall x(\forall y(Hxy > Fxy) > Rx)$ be called Γ . Adapt EXAMPLE 3, requiring in addition of I that $I_{\mu}(R) = \{w\}$. It is a matter of routine to show that Γ is satisfied in the resulting model. From OBSERVATION 4 it again follows that Γ is satisfiable together with all of MP. Now $\Gamma, MP \vdash Rw$. ■

From the above demonstrations, which require only that the premises are satisfiable in the frame of EXAMPLE 1, it will be plain that "irrelevant" additional information can be added to the premises without losing these conclusions. For example, the previous examples can be repeated on adding, say, an extra premise Pw expressing that w is a piece of pine. Other analyses of such reasoning which depart from weak conditional logics, such as the previously cited suggestions of Delgrande and Boutilier, have had to go to great lengths to allow such irrelevant additional premises to be added.

EXAMPLE 11: $Hw, \neg Fwe, \forall y(Hwy > Fwy) \not\vdash_{\mu} Fwe$.

Proof: Let the set of the premises be called Γ . EXAMPLE 5 shows that Γ is satisfiable. It follows with the deduction theorem for $\rightarrow$ and the compactness of $\vdash$ that there is a maximal Γ -consistent subset — call it Ω — of MP. Now $\Gamma, \Omega \not\vdash Fwe$. ■

The following example is similar to the well-worn *Nixon Diamond*, which could have been analyzed in its place. Let Wx express that x is wet, so that $\forall x \forall y (Wx > \neg Fxy)$ expresses that wet things normally don't burn. The following example shows that neither the inference that w will catch fire in e , nor that w will not catch fire, is allowed from premises expressing that w is flammable, that e is an event in which w is thoroughly heated, that w is wet, and that wet things normally do not catch fire.

EXAMPLE 12: Let Γ be $\{Hw, Ww, \forall y(Hwy > Fwy), \forall x \forall y (Wx > \neg Fxy)\}$. Then (i) $\Gamma \not\vdash_{\mu} Fwe$ and (ii) $\Gamma \not\vdash_{\mu} \neg Fwe$.

Proof: (i) With a model construction based on the frame of EXAMPLE 2 it is straightforward to show that Γ is satisfiable together with $(Ww > \neg Fwe) \rightarrow (Ww \rightarrow \neg Fwe)$. It follows that there is a maximal Γ -consistent $\Omega \subseteq MP$ such that $(Ww > \neg Fwe) \rightarrow (Ww \rightarrow \neg Fwe) \in \Omega$. Now $\Gamma, \Omega \not\vdash Fwe$.

(ii) Analogous ■

4 ... and Allowed Entailment

In this section I will generalize the syntactic notion of Allowed Consequence of the previous section. I will then define the nonmonotonic semantic notion of *Allowed Entailment*. Later I will prove that Allowed Consequence is sound and complete relative to Allowed Entailment.

Let $\mathcal{F}$ be any set of sentences called the reasoner's *focus of attention*. (For the meantime $\mathcal{F}$ will always be some or other subset of MP.) Now redefine Allowed Consequence relative to $\mathcal{F}$, by substituting $\mathcal{F}$ for MP in definition 6. The resulting notion $\vdash_{\mathcal{F}}$ of Allowed Consequence relative to a focus of attention clearly is a generalization of $\vdash_{\mu}$. Trivially, choosing $\mathcal{F}$ as MP, $\vdash_{\mathcal{F}}$ and $\vdash_{\mu}$ coincide. To go to the other extreme, choosing $\mathcal{F}$ as the empty set, $\vdash_{\mathcal{F}}$ reduces to logical consequence. Intuitively speaking the focus of attention is a parameter in nonmonotonic reasoning which reflects some of the reasoner's relative concerns: to include a conditional $\phi > \psi \rightarrow \phi \rightarrow \psi$ in $\mathcal{F}$ is in a sense to let it matter that normally ψ if ϕ . It is a crude way of dividing defaults into two priority classes: defaults appearing in $\mathcal{F}$ are alike in having top priority; defaults excluded from $\mathcal{F}$ are alike in having no effect whatsoever on the nonmonotonic consequences of a theory. This part played by the parameter of focus will be greatly extended in the next section, where it will be filled by a vector of sets of defaults. The following notion is central in defining a preference relation on possible worlds which can be used to characterize Allowed Consequence:

DEFINITION 14: $\mathcal{L}, v \leq_{\mathcal{F}} \mathcal{M}, w$ means

$$\mathcal{Th}(\mathcal{L}, v) \cap \mathcal{F} \subseteq \mathcal{Th}(\mathcal{M}, w).$$

Here $\mathcal{Th}(\mathcal{M}, w)$ is the *theory* $\{\phi: \mathcal{M}, w \models \phi\}$ of w in $\mathcal{M}$. One world lies underneath another in the sense of this relation if it "gets more of $\mathcal{F}$ right". With $\mathcal{F}$ chosen within MP, though, more can be said about $\leq_{\mathcal{F}}$. Consider any given sentence $\phi > \psi \rightarrow \phi \rightarrow \psi$ in $\mathcal{F}$. We can say that w is *irregular in regard to $\phi > \psi$ in $\mathcal{M}$* if $\mathcal{M}, w \models \phi$ and $\mathcal{M}, w \not\models \phi > \psi$, but $\mathcal{M}, w \models \psi$. Equivalently, of course, w is irregular in regard to $\phi > \psi$ iff $\mathcal{M}, w \not\models \phi > \psi \rightarrow \phi \rightarrow \psi$. So $\mathcal{L}, v \leq_{\mathcal{F}} \mathcal{M}, w$ just means that v is no more irregular than is w ; not, at least, as far as the defaults we are focusing on are concerned.²

² Notice that I allow possible worlds of *different* models $\mathcal{L}$ and $\mathcal{M}$ to be compared by $\leq_{\mathcal{F}}$, not just possible worlds of one and the same model. This has raised the eyebrows of some who are used to seeing worlds compared — say in regard to their similarity to a third world — but only within models. In standard developments of modal conditional logic, such as Lewis [1973], such similarity relations are of course what make conditionals true or false, as the case may be. I don't know why my comparing the worlds of different models is thought odd; maybe people confuse the *metalinguistic* task of defining a nonmonotonic consequence notion, which is what concerns me here, with the task of giving the truth conditions of an object-language conditional. Indeed one of the differences between defaults and other conditionals on the one hand and consequence notions on the other, namely that whereas the former can be embedded the latter cannot, has been hid from view by the

This relation will be used to define a semantic notion of nonmonotonic consequence analogous to John McCarthy's [1980] notion of Minimal Entailment. First there is the notion of a set of premises Γ being minimally satisfied, relative to $\mathcal{F}$, at a possible world w of a model $\mathcal{M}$:

DEFINITION 15: $\mathcal{M}, w \models_{\mathcal{F}} \Gamma$ just in case:

- (i) $\mathcal{M}, w \models \Gamma$ and
- (ii) if $\mathcal{L}, v \models \Gamma$ and $\mathcal{L}, v \leq_{\mathcal{F}} \mathcal{M}, w$, then $\mathcal{M}, w \leq_{\mathcal{F}} \mathcal{L}, v$.

Second, a sentence ϕ is an *Allowed Entailment* of premises Γ relative to $\mathcal{F}$ (in symbols, $\Gamma \models_{\mathcal{F}} \phi$) if ϕ holds in every world where Γ is minimally satisfied relative to $\mathcal{F}$:

DEFINITION 16: $\Gamma \models_{\mathcal{F}} \phi$ just in case for every $\mathcal{M}$ and w ,
if $\mathcal{M}, w \models_{\mathcal{F}} \Gamma$ then $\mathcal{M}, w \models \phi$.

Now I can state the completeness theorem for Allowed Consequence relative to Allowed Entailment:

COMPLETENESS THEOREM 17: $\Gamma \vdash_{\mathcal{F}} \phi \Leftrightarrow \Gamma \models_{\mathcal{F}} \phi$

This theorem provides Allowed Consequence with a semantic underpinning: the consequences which it is allowed to draw from a given set of premises are just the things which you can be sure are true in all least irregular worlds where the premises are true. It is a special case of the corresponding completeness of Prioritized Allowed Consequence with respect to Prioritized Allowed Entailment, which I will prove in section VI, so the proof will have to wait until then.

One consequence of either theorem which it is interesting to note at this point is that if Γ is consistent, so are its allowed consequences and entailments. This follows from THEOREM 17 together with DEFINITION 7 since, with the deduction theorem for $\rightarrow$ and the compactness of $\vdash$, for any consistent theory Γ there is at least one set Ω which is maximal Γ -consistent within $\mathcal{F}$. As Davis [1980] points out the corresponding thing does not hold in general for Circumscription: there are examples of theories which are classically consistent but from which contradictions follow using Circumscription.

5 Prioritized Allowed Consequence...

Nonmonotonic reasoning needs to be sensitive to the fact that some defaults take precedence over others. Thus, to take the threadbare example, when it comes to drawing conclusions about the ability or otherwise of any given penguin to fly, the default that penguins normally cannot fly takes precedence over the default that birds normally can fly. Laws, interpreted following Raz [1990] as *prima facie* reasons to act, provide a less familiar example. The example is of added interest in that in the case of defaults which express *prima facie* reasons to act which derive from laws, precedence orders do not simply mirror the taxonomic order of the kinds which the defaults are about (as it does in the case of the birds, where the default about the subkind penguin takes precedence over the default about the kind, birds). Instead, precedence orders are in the legal case often simply declared in legal codes. So, for example, in virtue of the law which requires this, the fact that a traffic light

shows red is a *prima facie* reason to stop at an intersection. Additionally, however, the law requires drivers to follow the directions of traffic officers, so the fact that an officer directing traffic is waving you on is a reason to pass through an intersection. And supposing you are waved on through a red light? The traffic code determines explicitly that the directions of traffic officers take precedence over traffic signs, which include lights, so in that case you ought to ignore the light and follow the directions.

In this section I will generalize the notion of Allowed Consequence to take into account priority orders between default conditionals. In the next section the notion of Allowed Entailment will be similarly generalized, and I will prove the completeness theorem according to which these two notions coincide. The first step is to generalize the notion of a focus of attention from a single set of defaults to a linearly ordered (finite) set of such sets, or a *focal vector*.

DEFINITION 18: $\underline{G} = \langle G_1, G_2, \dots, G_k \rangle$ is a *focal vector* just in case each G_i is a focus of attention in the earlier sense.

The idea is that defaults in G_1 will take precedence in conflicts with any defaults from $G_2, G_3, \dots$; defaults in G_2 will take precedence in conflicts with any defaults from $G_3, G_4, \dots$ and so on. Letting $\underline{G} = \langle G_1, G_2, \dots, G_k \rangle$ be a focal vector and $\mathcal{H}$ be a set of sentences, $\underline{G} \models \mathcal{H}$ is a useful notation for the focal vector $\langle G_1, G_2, \dots, G_k, \mathcal{H} \rangle$. Using this notation definitions and proofs can proceed by induction on the length of focal vectors. For a trivial start, $\underline{G}$ can be "flattened" into a set $\cup \underline{G}$ of sentences as follows: $\cup \langle \rangle = \{ \}$; and $\cup \underline{G} \models \mathcal{H} = \cup \underline{G} \cup \mathcal{H}$. The following notion generalizes to the case of focal vectors the earlier notion of maximal Γ -consistency within $\mathcal{F}$:

DEFINITION 19: Ω is *maximal Γ -consistent within $\underline{F}$* is defined by induction on the length of $\underline{F}$:

Base step: $\{ \}$ is maximal Γ -consistent within $\langle \rangle$;

Induction step: Ω is maximal Γ -consistent within $\underline{G} \models \mathcal{H}$ if (i) $\Omega \subseteq \cup \underline{G} \models \mathcal{H}$;

(ii) $\Omega \cap \cup \underline{G}$ is maximal Γ -consistent within $\underline{G}$; and

(iii) $\Gamma \cup \Omega \not\vdash \perp$. And if $\Omega \subseteq \Omega^* \subseteq \cup \underline{G} \models \mathcal{H}$ and $\Gamma \cup \Omega^* \not\vdash \perp$, then $\Omega = \Omega^*$.

OBSERVATION 20: Let $\mathcal{M}, w \models \Gamma \cup \Omega$, where Ω is maximal Γ -consistent within $\underline{F}$. Then $\mathcal{Th}(\mathcal{M}, w) \cap \cup \underline{F} = \Omega$. This follows directly from (iii) in the induction step of DEFINITION 19.

EXAMPLE 21: Let it be agreed that in reasoning about the properties of different kinds of birds, precedence orders on defaults must parallel the taxonomic order of the kinds which they are about. Thus where K_1 is a subkind of K_2 , any defaults about members of K_1 which are in focus must have higher priority than any defaults about members of K_2 which are in focus. And let " $\phi \triangleright \psi \rightarrow \phi \rightarrow \psi$ " be abbreviated " $\phi \triangleright \rightarrow \psi$ ". Finally, let an open formula in a focal vector stand for the set of its ground instantiations. Then for example the focal vector

$\langle \text{PENGUIN}(x) \triangleright \rightarrow \neg \text{FLIES}(x), \text{BIRD}(x) \triangleright \rightarrow \text{FLIES}(x) \rangle$

will be acceptable, but

$\langle \text{BIRD}(x) \triangleright \rightarrow \text{FLIES}(x), \text{PENGUIN}(x) \triangleright \rightarrow \neg \text{FLIES}(x) \rangle$

practice of concentrating on unnested defaults. For more on the differences between consequence notions and conditionals see for example Quine [1950].

will not. Call the first of these vectors $\mathcal{F}$. And let our premises Γ include just

$$\begin{aligned} &\forall x(\text{BIRD}(x) \supset \text{FLIES}(x)) \\ &\forall x(\text{PENGUIN}(x) \supset \neg \text{FLIES}(x)), \\ &\text{BIRD}(\text{tweety}), \text{PENGUIN}(\text{tweety}). \end{aligned}$$

Using the frame of EXAMPLE 2 it is not hard to show that Γ is satisfiable along with $\forall x(\text{PENGUIN}(x) \rightarrow \neg \text{FLIES}(x))$. So Γ is satisfiable along with all of $\mathcal{F}_1$, which is just the set of all sentences of the form $\text{PENGUIN}(c) \supset \neg \text{FLIES}(c)$. So there is a unique maximal Γ -consistent set within $\langle \mathcal{F}_1 \rangle$, and that is $\mathcal{F}_1$ itself. $\mathcal{F}_2$ contains only sentences of the form $\text{BIRD}(c) \supset \text{FLIES}(c)$, none of which are consistent together with $\Gamma \cup \mathcal{F}_1$. It follows from DEFINITION 19 that $\mathcal{F}_1$ is the unique maximal Γ -consistent set within $\mathcal{F}$.

DEFINITION 22: (*Prioritized Allowed Consequence*): $\Gamma \vdash_{\mathcal{F}} \varphi$ iff for every Ω which is maximal Γ -consistent within $\mathcal{F}$, $\Gamma \cup \Omega \vdash \varphi$.

CONTINUATION OF EXAMPLE 21: With Γ and $\mathcal{F}$ as above, $\Gamma \vdash_{\mathcal{F}} \neg \text{FLIES}(\text{tweety})$ but (since $\Gamma \cup \mathcal{F}_1$ is consistent) $\Gamma \not\vdash_{\mathcal{F}} \text{FLIES}(\text{tweety})$.

It is not hard to see that here as in the previous examples a range of "irrelevant" background premises can be added to Γ without changing this outcome. Also, once this example has been understood it will be clear how others involving prioritization of defaults can be analyzed, such as the traffic-law case discussed informally above.

In the following section I go on to provide Prioritized Allowed Consequence with its semantic counterpart, *Prioritized Allowed Entailment*. All that needs to be done is to generalize DEFINITION 14 to an order $\leq_{\mathcal{F}}$ on possible worlds relative to a focal vector $\mathcal{F}$. Substitution of $\mathcal{F}$ for $\mathcal{F}$ in the earlier definitions then results in the right notions. I also prove that Prioritized Allowed Consequence is sound and complete with respect to Prioritized Allowed Entailment.

6 ... and Prioritized Allowed Entailment

Intuitively it is clear what needs to be done in order to define $\leq_{\mathcal{F}}$ for any given focal vector $\mathcal{F} = \langle \mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k \rangle$. It is necessary to combine the different orders $\leq_{\mathcal{F}_1}, \leq_{\mathcal{F}_2}, \dots, \leq_{\mathcal{F}_k}$ appropriately, allowing $\leq_{\mathcal{F}_1}$ to order worlds first, then allowing $\leq_{\mathcal{F}_2}$ to further discriminate between worlds which as far as $\leq_{\mathcal{F}_1}$ is concerned are not discriminable, and so on. Such combinations of partial orders are the subject of the following definition, in which v is a finite vector of partial linear orders, each defined on the same domain, $\mathcal{D}$:

DEFINITION 23: The *alphabetic composition* $\leq_v$ of v is defined by induction on the length of v :

base step: $\leq_{\langle \rangle} = \mathcal{D}^2$;

induction step: $\leq_{d \leq} = \{(d_1, d_2) : d_1 \leq_v d_2 \text{ and if } d_1 \approx_v d_2, \text{ then } d_1 \leq d_2\}$.

A straightforward induction on the length of v shows that $\leq_v$ is a partial linear order on $\mathcal{D}$. The familiar alphabetic

order of words in a dictionary can be construed as such an alphabetic composition.

OBSERVATION 24: Let $v = \langle \leq_{v_1}, \leq_{v_2}, \dots, \leq_{v_k} \rangle$. For all $d_1, d_2 \in \mathcal{D}$, $d_1 \approx_v d_2$ just in case for each $i \leq k$, $d_1 \approx_{v_i} d_2$.

This observation follows with a simple induction on the length of v . It and the following observations will be needed later:

OBSERVATION 25: If $(\mathcal{H}(\mathcal{M}, w) \cap \cup \mathcal{F}) \subseteq \mathcal{H}(L, v)$, then $L, v \leq_{\mathcal{F}} \mathcal{M}, w$. This can be seen with a simple induction on the length of $\mathcal{F}$. The base clause is trivial: $\leq_{\langle \rangle}$ is the universal relation on possible worlds of different models. For the inductive step, if $(\mathcal{H}(\mathcal{M}, w) \cap \cup \mathcal{G} \mid \mathcal{H}) \subseteq \mathcal{H}(L, v)$, then $(\mathcal{H}(\mathcal{M}, w) \cap \cup \mathcal{G}) \subseteq \mathcal{H}(L, v)$ and $(\mathcal{H}(\mathcal{M}, w) \cap \mathcal{H}) \subseteq \mathcal{H}(L, v)$. By induction hypothesis $L, v \leq_{\mathcal{G}} \mathcal{M}, w$. By definition 14 also $L, v \leq_{\mathcal{H}} \mathcal{M}, w$, so $L, v \leq_{\mathcal{G} \mid \mathcal{H}} \mathcal{M}, w$.

OBSERVATION 26: $L, v \approx_{\mathcal{F}} \mathcal{M}, w$ if and only if

$$(\mathcal{H}(L, v) \cap \cup \mathcal{F}) = (\mathcal{H}(\mathcal{M}, w) \cap \cup \mathcal{F}).$$

This is a direct consequence of DEFINITION 14, OBSERVATION 24 and OBSERVATION 25.

Now the following notions of minimal satisfaction and Prioritized Allowed Entailment are completely analogous to the earlier definitions relative to a simple focus:

DEFINITION 27: Let $\mathcal{F}$ be a focal vector. Then

(a) $\mathcal{M}, w \models_{\mathcal{F}} \Gamma$ means:

(i) $\mathcal{M}, w \models \Gamma$; and

(ii) for all L, v , if $L, v \models \Gamma$ and $L, v \leq_{\mathcal{F}} \mathcal{M}, w$, then $L, v \approx_{\mathcal{F}} \mathcal{M}, w$.

(b) $\Gamma \models_{\mathcal{F}} \varphi$ means that for every $\mathcal{M}, w \models_{\mathcal{F}} \Gamma$, $\mathcal{M}, w \models \varphi$.

Now we have:

THEOREM 28: $\Gamma \vdash_{\mathcal{F}} \varphi$ just in case $\Gamma \models_{\mathcal{F}} \varphi$.

Putting the focal vector as $\langle \mathcal{F} \rangle$, THEOREM 17 is a special case of this theorem. THEOREM 28 follows immediately from the following lemma, which characterizes the possible worlds in which any given set of premises is minimally satisfied:

LEMMA 29: $\mathcal{M}, w \models_{\mathcal{F}} \Gamma$ just in case for some Ω which is Γ -consistent within $\mathcal{F}$, $\mathcal{M}, w \models \Gamma, \Omega$.

The proof of this lemma makes use of the following:

FACT 30: $\mathcal{M}, w \models_{\mathcal{G} \mid \mathcal{H}} \Gamma$ just in case

(i) $\mathcal{M}, w \models_{\mathcal{G}} \Gamma$, and

(ii) for any Ω^* , if $(\mathcal{H}(\mathcal{M}, w) \cap \cup \mathcal{G} \mid \mathcal{H}) \subseteq \Omega^* \subseteq \cup \mathcal{G} \mid \mathcal{H}$ and $\Gamma \cup \Omega^* \not\models \perp$, then $(\mathcal{H}(\mathcal{M}, w) \cap \cup \mathcal{G} \mid \mathcal{H}) = \Omega^*$.

Proof of FACT 30 $\Rightarrow$ Suppose $\mathcal{M}, w \models_{\mathcal{G} \mid \mathcal{H}} \Gamma$.

Then (i) By definition of $\models_{\mathcal{G} \mid \mathcal{H}}$, $\mathcal{M}, w \models \Gamma$. Now let $L, v \models \Gamma$ and $L, v \leq_{\mathcal{G}} \mathcal{M}, w$, and suppose for the contradiction that not $L, v \approx_{\mathcal{G}} \mathcal{M}, w$. Then $L, v \leq_{\mathcal{G} \mid \mathcal{H}} \mathcal{M}, w$ so $L, v \approx_{\mathcal{G} \mid \mathcal{H}} \mathcal{M}, w$. From this follows by OBSERVATION 24 that $L, v \approx_{\mathcal{G}} \mathcal{M}, w$.

Also (ii) Suppose $(\mathcal{H}(\mathcal{M}, w) \cap \cup \mathcal{G} \mid \mathcal{H}) \subseteq \Omega^* \subseteq \cup \mathcal{G} \mid \mathcal{H}$ and $L, v \models \Gamma \cup \Omega^*$. Then $\Omega^* \subseteq (\mathcal{H}(L, v) \cap \cup \mathcal{G} \mid \mathcal{H})$, so by OBSERVATION 25, $L, v \leq_{\mathcal{G} \mid \mathcal{H}} \mathcal{M}, w$. Now by assumption $\mathcal{M}, w \models_{\mathcal{G} \mid \mathcal{H}} \Gamma$. Since $L, v \models \Gamma$, it follows that $L, v \approx_{\mathcal{G} \mid \mathcal{H}} \mathcal{M}, w$. By OBSERVATION 26 then $(\mathcal{H}(\mathcal{M}, w) \cap \cup \mathcal{G} \mid \mathcal{H}) = (\mathcal{H}(L, v) \cap \cup \mathcal{G} \mid \mathcal{H})$, so $(\mathcal{H}(\mathcal{M}, w) \cap \cup \mathcal{G} \mid \mathcal{H}) = \Omega^*$.

$\Leftarrow$ Assume (i) and (ii). Clearly $\mathcal{M}, w \models \Gamma$; now suppose $\mathcal{L}, v \models \Gamma$ and $\mathcal{L}, v \leq_{\mathcal{G}|\mathcal{H}} \mathcal{M}, w$. Then (by definition of $\leq_{\mathcal{G}|\mathcal{H}}$) $\mathcal{L}, v \leq_{\mathcal{G}} \mathcal{M}, w$ so (by i.) $\mathcal{L}, v \models_{\mathcal{G}} \mathcal{M}, w$. By OBSERVATION 26, $(Th(\mathcal{M}, w) \cap \cup \mathcal{G}) \subseteq Th(\mathcal{L}, v)$, and (by definition of $\leq_{\mathcal{G}|\mathcal{H}}$) $\mathcal{L}, v \models_{\mathcal{H}} \mathcal{M}, w$, so (by DEFINITION 14) also $(Th(\mathcal{M}, w) \cap \mathcal{H}) \subseteq Th(\mathcal{L}, v)$. Thus $(Th(\mathcal{M}, w) \cap \cup \mathcal{G}|\mathcal{H}) \subseteq (Th(\mathcal{L}, v) \cap \cup \mathcal{G}|\mathcal{H}) \subseteq \cup \mathcal{G}|\mathcal{H}$. By ii. then $(Th(\mathcal{M}, w) \cap \cup \mathcal{G}|\mathcal{H}) = (Th(\mathcal{L}, v) \cap \cup \mathcal{G}|\mathcal{H})$. Now by OBSERVATION 26, $\mathcal{L}, v \models_{\mathcal{G}|\mathcal{H}} \mathcal{M}, w$. This completes the proof of fact 30. ■

Proof of LEMMA 29 The proof is by induction on the length of $\mathcal{F}$. $\mathcal{F} = \langle \rangle$: Trivial, since $\leq_{\langle \rangle}$ is universal and $\{\}$ is the only maximal Γ -consistent set within $\langle \rangle$. So let $\mathcal{F} = \mathcal{G}|\mathcal{H}$. For convenience write Σ instead of $Th(\mathcal{M}, w) \cap \cup \mathcal{G}|\mathcal{H}$. (So, very trivially, $\Sigma \subseteq \cup \mathcal{G}|\mathcal{H}$ and $\Sigma \not\models \perp$) We have the following equivalencies:
 $\mathcal{M}, w \models_{\mathcal{F}} \Gamma \Leftrightarrow \text{FACT 30}$
 $\mathcal{M}, w \models_{\mathcal{G}} \Gamma$, and
 if $\Sigma \subseteq \Omega^* \subseteq \cup \mathcal{G}|\mathcal{H}$ and $\Gamma \cup \Omega^* \not\models \perp$, then $\Sigma = \Omega^*$
 $\Leftrightarrow$ INDUCTION HYPOTHESIS, OBSERVATION 20
 $\Sigma \cap \cup \mathcal{G}$ is maximal Γ -consistent within $\mathcal{G}$, and
 if $\Sigma \subseteq \Omega^* \subseteq \cup \mathcal{G}|\mathcal{H}$ and $\Gamma \cup \Omega^* \not\models \perp$, then $\Sigma = \Omega^*$
 $\Leftrightarrow$ DEFINITION 19
 Σ is maximal Γ -consistent within $\mathcal{G}|\mathcal{H}$
 $\Leftrightarrow$ OBSERVATION 20
 For some Ω which is maximal Γ -consistent within $\mathcal{F}$ $\mathcal{M}, w \models \Gamma, \Omega$. This completes the proof of lemma 29. ■

7 Strengthening Conditional Antecedents

Allowed Consequence and Entailment solve a problem about reasoning with conditionals which I will now very briefly sketch.

If Shoemaker-Levy 9 had hit Earth then there would have been an explosion equivalent to at least 500 megatons of TNT. It seems reasonable to infer that there would have been an explosion of this magnitude if Shoemaker-Levy 9 had hit Earth right after the world-cup final, but the logic of conditionals does not support this inference; Lewis [1973] argues convincingly that the logic of counterfactuals cannot validate the strengthening of conditional antecedents.

There is also the hypothetical syllogism, or *transitivity*. If Shoemaker-Levy 9 had hit Earth, then there would have been a 500 megaton explosion. And if there had been a 500 megaton explosion, then life as we know it would have come to an end. So, you might think, if Shoemaker-Levy 9 had hit Earth, then life as we know it would have come to an end. But again the logic of counterfactual conditionals offers no support. It cannot do since, as can easily be shown, under minimal assumptions validity of transitivity entails that counterfactual antecedents can always be strengthened.

Let TR be the axiom scheme $(\phi > \psi \wedge \psi > \chi) \rightarrow \phi > \chi$. The logic of counterfactuals cannot be allowed to validate TR, but the notion of allowed consequence suggests another way to account for the plausibility of inferences like the examples just mentioned. Just as in this paper the

logic of defeasible conditionals was augmented by adding to any given premises as many instances of MP as one consistently can, the logic of counterfactual or for that matter any other conditionals can be augmented by adding instances of TR. This is achieved by including instances of TR in $\mathcal{F}$. Then the defeasible notion of allowed consequence can account for reasonable examples of strengthening of the antecedent and the hypothetical syllogism, while unreasonable examples, such as inferring *the coffee would have tasted good if there had been sugar and diesel oil in it* from *the coffee would have tasted good if there had been sugar in it*, are "blocked" by background knowledge. Working through this and other examples in detail is something for another time.

References

- Asher, N. and M. Morreau: 1991, 'Commonsense Entailment: a modal theory of nonmonotonic reasoning.' In J. Mylopoulos and R. Reiter, eds., *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, Morgan Kaufman, Los Altos, California, pp. 387-392.
- Boutilier, C.: 1992, *Conditional Logics for Default Reasoning and Belief Revision*. Technical Report KRR-TR-92-1, Computer Science Department, University of Toronto, Ontario, 1992.
- Chellas, B.: 1980, *Modal Logic, an Introduction*. Cambridge University Press, Cambridge.
- Davis, M. 1980, 'The Mathematics of Nonmonotonic Reasoning.' *Artificial Intelligence* 13.
- Delgrande, J.: 1988, 'An Approach to Default Reasoning Based on a First-Order Conditional Logic: revised report.' *Artificial Intelligence* 36 (1988), pp. 63-90.
- Lascarides, A. and N. Asher.: 1993, 'Temporal Interpretation, Discourse Relations and Commonsense Entailment.' *Linguistics and Philosophy* 16, pp. 437-494.
- Lewis, D.: 1973, *Counterfactuals*. Harvard University Press.
- McCarthy, J.: 1980, 'Circumscription — a Form of Nonmonotonic Reasoning.' *Artificial Intelligence* 13, pp. 27-39.
- Quine, W.V.O.: 1950. *Methods of Logic*. New York, Holt.
- Raz, J.: 1990, *Practical Reason and Norms*. 2nd edition, Princeton University Press.
- Reiter, R.: 1980, 'A Logic for Default Reasoning.' *Artificial Intelligence* 13, pp. 81-132.
- Stalnaker R. and R. H. Thomason: 1970, 'A Semantic Analysis of Conditional Logic.' *Theoria*, 36, pp. 23-42.
- Thomason R. H.: 1990, 'Propagating Epistemic Coordination through Mutual Defaults I.' In R. Parikh, ed., *Theoretical Aspects of Reasoning about Knowledge, proceedings of the third conference*. Morgan Kaufmann, San Mateo, 1990, pp. 29-40.