What's in a Default?
Thoughts on the Nature and Role of Defaults in Nonmonotonic Reasoning

James Delgrande
School of Computing Science,
Simon Fraser University,
Burnaby, B.C.,
Canada V5A 1S6.
jim@cs.sfu.ca

Abstract

This paper examines the role and meaning of defaults in nonmonotonic reasoning (NMR). Defaults, that is, statements that express a condition of normally such as "adults are normally employed", are crucial in commonsense reasoning and in artificial intelligence in general. The majority of research concerning defaults has focussed on (default) inference mechanisms, rather than representational issues involving the meaning of a default. I suggest that, despite the very impressive formal work in the area, it would be useful to (re)consider defaults with respect to the phenomena that they are intended to model.

To start, I briefly consider how defaults have been represented in NMR, along with informal interpretations of defaults. Two major distinctions are explored. The first considers the view of a default as an assertion about some domain, as opposed to an inferential procedure for deriving properties of individuals. The second distinction considers the manner in which default application is informally treated, whether as a weak "rule" or essentially as a weak material implication. I suggest that the "weak material conditional" interpretation is not suitable in the case of defaults; this is problematic since most existing approaches take this latter interpretation.

Subsequently, I argue that defaults of normality are best regarded as statements in a naïve scientific theory. A theory of the meaning of such defaults can be given by a logic of weak conditionals, in which a default is treated as a counterfactual normative statement. From this vantage, nonmonotonic reasoning with such conditionals can be re-examined. To this end, the notion of relevant properties emerges as a key factor in drawing default conclusions about an individual. As well, other phenomena, such as reasoning about norms, or deontic assertions, or counterfactuals may be addressed in a similar fashion.

1 Introduction

Classical reasoning is *monotonic*, which is to say it adheres to a principle of monotonicity:

Monotonicity: If $\Gamma \vdash \phi$ then $\Gamma, \Delta \vdash \phi$

Thus, having proven a result in geometry, say, it is absurd to suggest that learning more information about the problem would invalidate the conclusion. On the other hand, our commonsense, everyday knowledge is for the most part *nonmonotonic*, in that it fails to satisfy monotonicity. Hence, for example, on being told that an individual is a bird, one will conclude that it flies, while on being later informed that it is a penguin or is a nestling, one will conclude that it does not fly.

The area of nonmonotonic reasoning (NMR) in artificial intelligence (AI) studies such reasoning. NMR then is a central and crucial area of AI, and is fundamental to commonsense reasoning. The past 30 years have seen much impressive and important work in NMR, beginning with the seminal Artificial Intelligence Journal issue on the topic [AIJ, 1980]. At this stage, 30 years on, we now have a good understanding of principles underlying NMR, and it would appear that the major approaches to NMR have been identified and are well explored.

In NMR, a fundamental notion is that of a *default*, where a default can be thought of as a weak, or defeasible, conditional. The principal use of defaults in NMR is to ascribe default properties to individuals. A default is generally expressed in English in the form "X's are (normally) Y's".[1] Examples include the hackneyed "birds (normally) fly", as well as "adults are employed" or "snow is white".[2] So a primary task in NMR is to come up with a principled means of dealing with such statements.

To this end, it can be observed that much research in NMR has dealt with the development of formal approaches for reasoning with defaults. Paradigmatically,[3] particularly during the early days of the 1980's and 1990's, a research program would involve proposing a formal mechanism, examing its suitability with respect to dealing with defaults, locating glitches in the representation, modifying the approach, and so proceed. The approaches developed, whether default logic, circumscription, nonmonotonic inference relations, conditional logics, or others,[4] are arguably among the most impressive and important formalisms developed in AI.

On the other hand, part of the task of AI researchers is to apply such approaches to real-world problems. That is, these formalisms are intended to be used, and resulting knowledge bases will be used to encode information

[1] In linguistics, such sentences are examples of the broader class of *generics* [Carlson and Pelletier, 1995]. Insofar as possible, linguistic issues and, particularly, linguistic conventions or communication conventions are avoided here. Rather, the focus is on purely representational and reasoning issues.

[2] This last is the well-known example attached to Tarski's theory of truth. But of course it isn't the case that snow is (unreservedly) white, but rather that snow is normally white.

[3] and stereotypically

[4] See the next section for references and brief descriptions

about and to reason about the world. The main thesis of this paper is that in some important cases default formalisms don't capture the phenomenon that they're intended to model. Arguably, a large part of the problem is that representational issues have received insufficient attention.

In this paper, I first review the notion of a (normality) *default*, specifically how defaults have been represented in NMR, along with informal interpretations of defaults. Subsequently, two major distinctions are explored. The first contrasts, on the one hand, the notion of a default as a general assertion about some domain with, on the other hand, "applying" a default to an individual to derive a property of that individual. That is, one can consider a default such as "birds fly" as asserting something about a domain; on the other hand, one can use a default to obtain a default conclusion, such as that a given bird flies. It can be observed that most approaches ignore this first aspect, or conflate these aspects. The second distinction considers the manner in which default application is informally regarded, whether as a weak "rule" or whether as a weak material implication. I suggest that the "weak material conditional" interpretation is not suitable in the case of normality defaults. This, insofar as reasoning with defaults goes, is problematic since most existing approaches take the latter interpretation. The overall conclusion is that, despite the very impressive formal apparatuses developed and despite the remarkable success in applying these approaches in areas such as reasoning about action and planning, diagnosis, database systems, and logic programming, nonetheless there remains a general problem with defaults of normality with "getting the inferences right".

I suggest, toward a direction for a solution, that defaults are best regarded as naïve scientific statements. After developing this argument, I also suggest that this view will lead to a better understanding of reasoning with defaults and that perhaps it will also allow a wider application of default reasoning to other types of weak conditionals, including counterfactuals, deontics, and statements of causality.

2 Defaults

To begin, it seems fair to ask, *What is a default?* This question will be addressed in part throughout this section. Commonly, defaults are expressed in the form "X's are Y's" or "If X then normally Y". As a starting point, a default can be taken as an assertion about the world. Thus "birds fly" says something about the class of birds. We can also ask *How are defaults used?*; and here it can be noted that the standard use of defaults is to draw plausible conclusions, or conclusions in situations where we have only incomplete information.

Various approaches have been proposed for inference involving defaults, notably, Default Logic [Reiter, 1980] (and encompassing, for our purposes, the stable models semantics and answer set programming [Gelfond and Lifschitz, 1988, 1991, Baral, 2003], as well as autoepistemic logic [Moore, 1985, Denecker, Marek, and Truszczyński, 2003]), circumscription [McCarthy, 1980, Lifschitz, 1985, McCarthy, 1986], nonmonotonic inference relations (and associated closure operations) [Kraus, Lehmann, and Magidor, 1990, Lehmann and Magidor, 1992], and conditional logics [Delgrande, 1988, Lamarre, 1991, Boutilier, 1994].[5] As described, a general problem (then and now) is getting the inferences right: obtaining plausible, commonsense conclusions given a set of defaults and general assertions about a domain.

2.1 Default Inference: Encodings

We give the briefest of introductions to approaches to nonmonotonic reasoning here; for details the reader should consult the aforecited references, or general accounts such as [Brewka, 1991b, Antoniou, 1997, Brewka, Niemela, and Truszczynski, 2007].

A default can be encoded according to several quite different schemes:

"Rule of Inference": This is the approach of Default Logic. "Birds fly" can be encoded either in propositional or first order logic as follows:

$$\frac{Bird\ :\ Fly}{Fly} \qquad \text{or} \qquad \frac{Bird(x)\ :\ Fly(x)}{Fly(x)}.$$

In the first case, if $Bird$ is true and Fly is consistent with what is believed, then Fly is concluded. The notion of "is consistent with" is subtle, and leads to an intricate, elegant fixed point definition. In the second case, the same intuitive account can be given, except that x is instantiated to some ground term. That is, a default rule with free variables can be regarded as a schema, standing for the set of its ground instantiations.

Several points can be noted:

1. Despite its name, Default Logic is not a logic of defaults per se, since it doesn't give an account of a notion of truth of a default. Instead what is provided is a means of drawing default conclusions in the absence of information.
2. Since Default Logic isn't a logic of defaults, but rather is a general and powerful mechanism, one must "program" desirable properties for defaults. For example, given the additional default that penguins don't fly, along with the information that birds are penguins, one has to stipulate explicitly that the more specific penguins-don't-fly default takes priority over the less specific birds-fly default.
3. Quite frequently in the literature a default is expressed propositionally. However it is not clear what a default such as $\frac{Bird\,:\,Fly}{Fly}$ means; specifically, it is not clear what the propositions $Bird$ and Fly refer to. Probably the most intelligible gloss is that, in default reasoning one most often is reasoning about an individual, say x, and the defaults are phrased with reference to this individual. Hence the default may be more mnemonically encoded propositionally as $\frac{x\text{-}is\text{-}a\text{-}Bird\,:\,x\text{-}Flys}{x\text{-}Flys}$.

[5] Comments on these approaches will be seen to apply to other approaches, including conditional entailment [Geffner and Pearl, 1992] and abductive approaches such as Theorist [Poole, 1988], as well as inheritance networks [Horty, 1994].

Via Classical Logic: This is the approach taken by circumscription. The "birds fly" example can be encoded propositionally or in first-order logic as:

$$(Bird \wedge \neg Ab_F) \supset Fly \qquad \text{or} \qquad \forall x.(Bird(x) \wedge \neg Ab_F(x)) \supset Fly(x). \quad (1)$$

Thus, in the propositional case, if $Bird$ is true, and Ab_F is not, then one can derive Fly. The intended meaning of Ab_F is that the individual in question is not abnormal with respect to flight. In *circumscribing* the atom Ab_F, essentially if Ab_F *can* be taken to be false then it *is* taken to be false. For the predicate $Ab_F(\cdot)$, one analogously minimises the extension of the predicate. Semantically this is carried out by defining an ordering over models of a knowledge base, preferring those models where the Ab atoms are false (or the extension of the Ab predicates is smallest, respectively), and then just considering the minimal models of a knowledge base.

It can be noted that (1) leads to the very strong conclusion that unless a bird can be shown to not fly, one concludes that it flies. Hence for example, if one knew nothing about the birds of Madagascar, one would conclude that they all fly.

The points made concerning Default Logic also apply to circumscription: Circumscription is not a logic of defaults, but rather provides a means by which defaults can be encoded. Similarly, to make default inferences have the "right" properties one needs to enhance the approach. Thus to deal with specificity information such as that implicit between penguins and birds, priorities are introduced into circumscription. Last, of course, the same comments apply to the meaning of atoms like $Bird$ and Fly in the propositional formula in (1).

Inference Relation: The area of nonmonotonic inference relations can be regarded foremost as providing a general framework whereby general principles of nonmonotonicity may be studied, and via which other formalisms may be compared. However, one might also examine a specific nonmonotonic inference relation with regard to its suitability as an approach to dealing with defaults. The base approach which has been used for representing defaults is called *preferential reasoning*. Our canonical example would be expressed via a nonmonotonic inference relation as $Bird \mathrel{|\!\sim} Fly$. In this case, one can specify relations between nonmonotonic inferences, for example:

$$\text{From} \quad Bird \mathrel{|\!\sim} Fly, \quad Bird \mathrel{|\!\sim} Nest \quad \text{infer} \quad Bird \wedge Nest \mathrel{|\!\sim} Fly. \quad (2)$$

Hence from "birds fly" and "birds build nests", one can infer that "birds that build nests fly". The resulting systems are inferentially weak, at least with regards to obtaining desirable nonmonotonic consequences; for example one cannot infer $Bird \wedge Green \mathrel{|\!\sim} Fly$ from $Bird \mathrel{|\!\sim} Fly$. This is addressed by extending the set of inferences via a (nonmonotonic) *closure* operator.

Thus, it appears that a nonmonotonic inference relation is *about* defaults, in the sense that one might read (2) as saying (despite the phrasing as an inference relation) that from defaults "birds fly" and "birds build nests" one can infer the default "birds that build nests fly". The closure operator then, so it might seem, extends reasoning to that of deduction involving individuals. We discuss this point further in the next section.

Modal Operator: Last, one might consider a default to be a "real" assertion, carrying a truth value. In this case, our example could be encoded using a new connective, as $Bird \Rightarrow Fly$. In this instance $\Rightarrow$ is a binary modal operator, with intuitive meaning "in the most normal of worlds in which $Bird$ is true, Fly is also true." The example (2) can then be expressed as a formula:

$$(Bird \Rightarrow Fly \wedge Bird \Rightarrow Nest) \supset (Bird \wedge Nest) \Rightarrow Fly. \qquad (3)$$

It proves to be the case that there are very close connections between conditional logics of defaults and nonmonotonic inference relations. In fact, the central approach in each case has been shown to be translatable to the other, fully preserving inferences. As well, the central closure operator in each case has also been shown to be symmetrically translatable. It might seem that these two approaches are merely syntactic variants of each other. However, we later suggest that, despite these formal intertranslations, there are significant differences between the approaches with respect to their suitability for representing and reasoning with defaults.

None of the above schemes appears to be immediately suitable for fully dealing with defaults. In the case of Default Logic and circumscription, for example, one has powerful inference mechanisms, and the challenge is to modify the inference mechanism, or how it is applied, in order to get the "right" properties. This led to the development of variants of the basic approach, see for example [Łukaszewicz, 1988, Brewka, 1991a, Mikitiuk and Truszczyński, 1993, Delgrande, Schaub, and Jackson, 1995, Delgrande and Schaub, 1997] with respect to Default Logic. This also led to a general methodology for determining suitable default inference during the 1980's and 1990's which can be called "test-and-refine": Typically a nonmonotonic inference mechanism would be proposed or modified; it would be shown to work on a set of troublesome examples; later other troublesome examples would arise; the approach would be modified, and the process continued. Thus in dealing with specificity, in Default Logic semi-normal defaults were employed, while in circumscription a notion of prioritisation was introduced. The so-called Yale Shooting Problem [Hanks and McDermott, 1986]) is a good example of a problem for which the obvious encodings didn't work, but that also spurring significant resarch and results in the area. On the other side, nonmonotonic inference relations and modal approaches provide a semantically-justified account of the notion of a default, but in this case the difficulty lies in getting a nonmonotonic counterpart that has the "right" properties.

It can also be observed that the above paradigmatic schemes for expressing defaults are very different with respect to their form. Moreover, one would expect to obtain different conclusions depending on which approach is used to express a default. This raises several key questions: *Why should one prefer one approach over another?* And: *which approach is most suitable for dealing with defaults?* And moreover: *if two approaches lead to different conclusions, how does one judge which is "correct"?* From a formal point of view,[6] the answer to these questions is clear: One needs a theory of defaults in order to be able to determine what the properties of defaults should be. Consequently, in the next subsection, we examine the question of what a default informally means.

2.2 Interpretations of Defaults

Let's reconsider what it is we're trying to deal with and, to be specific, consider various possible informal interpretations of "birds fly". Among other alternatives, the following are possible readings of "birds fly":

1. Most birds fly
2. A bird that can be consistently assumed to fly does fly.
3. Birds normally fly
4. The prototypical bird flies
5. Birds generally/usually fly

A fair question to ask at this point is: *Do any of these interpretations align with the encodings that we've seen in the previous subsection, and if so, which and in what fashion?* While it isn't immediately clear which, if any, of the previous approaches fit with these interpretations, we *can* consider these interpretations with respect to how they fit with an informal notion of default.

Consider the first interpretation of "birds fly" as "most birds fly". This is clearly a statistical assertion: one has some population of birds in mind, and over half of them fly. One can develop an approach where, given that some large proportion of birds fly and individual x is a bird, one accepts the belief that x flies – indeed the late Henry Kyburg has addressed this interpretation. (See for example [Kyburg, 1994].) Without going into detail, we will note that such probabilistic approaches may be seen as being orthogonal to nonmonotonic reasoning. Kyburg expressed the difference as follows:

Schema for probabilistic inference: $\dfrac{\mathcal{BK}, E}{C,\ \text{hedged}}$

Schema for nonmonotonic inference: $\dfrac{\mathcal{BK}, E}{C}$ hedged inference

That is, in the probabilistic case, one makes a (monotonic) inference that is nonetheless "hedged". In the nonmonotonic case, a consequence is accepted while the inference itself is defeasible.

[6] Which is to say that there are also informal considerations. In particular, any theory of defaults will need to produce plausible or commonsense conclusions. Arguably it is the job of formalisation to precisely capture such informal notions.

In any case, the representation schemes that we have reviewed resist a probabilistic reading. In Default Logic, for example, a rule $\frac{Bird(x):Fly(x)}{Fly(x)}$ applies to individuals and (roughly) rests on a notion of consistency, not probability.[7] Similar considerations apply in circumscription. As well, the nonmonotonic inference relations or conditional logics that have been proposed to represent defaults cannot be given a probabilistic reading. That is, in preferential reasoning, $Bird \mid\!\sim Fly$ cannot be coherently interpreted as "most birds fly".[8] We return to this issue in Section 4.

The second reading, "a bird that can be consistently assumed to fly does fly" is clearly epistemic in nature. Autoepistemic logic addresses this interpretation; given the results of [Denecker et al., 2003] linking autoepistemic logic and Default Logic, Default Logic can also be interpreted in this light. However, such an interpretation clearly doesn't express the meaning of "birds fly"; instead it presents a way that a reasoner may conclude a default property of an individual. Note however that in this case it is not clear how one may draw *appropriate* conclusions; for example one would have to "program" a notion of specificity between defaults.

The next interpretation ("birds normally fly" or perhaps "the normal bird flies") is arguably closer to what is *meant* by "birds fly", since it seems to be simply true that birds normally fly. Conditional logics of defaults take this interpretation, and preferential or rational nonmonotonic inference relations can also be seen in this light. As indicated earlier, the issue here, assuming that one is happy with a given logic of defaults, is how to reason about the properties of specific individuals.

It can be noted however that this notion of normality has the following problem when it comes to reasoning about default properties: Consider yet again birds and their (default) properties. Presumably one would agree that a penguin should be concluded to have feathers. However, a penguin is clearly not a normal bird (it doesn't fly, for one thing) and so one could not use the statement "the normal bird has feathers" to reason about penguins. On the other hand, there seems to be no problem in asserting that penguins should be concluded to have feathers, since birds have feathers. This problem is pointed out in [Carlson and Pelletier, 1995],[9] where it is noted that presumably and hopefully every human being is exceptional in some fashion. But then no human is "normal" per se, and so one cannot directly appeal to a global notion of normality in concluding default properties. We expand on this also in Section 4.

The fourth interpretation refers to a prototype. In saying that "the prototypical bird flies", roughly one has an idea of the notion of a prototypical bird, or best or typical instance of the class of birds [Rosch, 1978]. Notions of prototypicality then are descriptive or contingent; the prototypical bird is

[7] However [Reiter, 1980] suggests that the default rule $\frac{Bird(x):Fly(x)}{Fly(x)}$ is intended to represent "most birds fly". What seems to be more appropriate is to say that if one accepts that "most birds fly" or "the large majority of birds fly" is true, then the default rule will let one jump to the conclusion that a specific bird flies.

[8] Some approaches propose the reading of "birds fly" as meaning "all but an infinitesimal number of birds fly". However, such a reading isn't just inaccurate; it seems to be simply wrong: Clearly there are significant numbers of birds that do not fly, while "birds fly" is true.

[9] See also [Poole, 1991]

essentially the "best" representative of the set of birds. Arguably, in reasoning about default properties, we want to go beyond notions of similarity to a prototype, which is to say, default reasoning is more than similarity to a given prototype.

The final interpretation, "birds generally or usually fly" is perhaps ambiguous. On the one hand, it can be read as a statement in qualitative probability, analogous to "most birds fly", in which case the earlier comments apply. On the other hand, it might be read as "the usual bird flies" which would seems to be roughly synonymous with "the normal bird flies".

This discussion of possible interpretations of defaults is not intended to be exhaustive; and it is quite possible that there are interpretations that have been missed. The discussion does emphasise the (obvious) point that "birds fly" is ambiguous. Moreover, given that we are interested in the meaning of a default such as "birds fly", we can rule out some of these informal interpretations. So, while we might agree that "most birds fly" is true, it doesn't capture the meaning of "birds fly" – for example if we were to arrange that all existing birds be held down, "birds fly" would still be true. Similarly, autoepistemic interpretations are inadequate to represent the meaning of "birds fly" (since certainly "birds fly" would be true even if there were no believers to hold beliefs about birds). Consequently we focus on the "normality" interpretation, and suggest that "birds normally fly" is what is *meant* by "birds fly".

3 Defaults: Two Issues

We next consider two issues regarding defaults. The first issue is representational, and concerns the dual aspects of defaults, as bearers of truth values and as things that are used for drawing inferences. As a specific consequence of this distinction, we also discuss the way in which formalisms for dealing with defaults have handled individuals. The second issue concerns reasoning, specifically whether a default is best informally regarded as a "weak rule of inference" or a "weak material conditional".

3.1 Defaults: Representation vs. Reasoning

Let's reconsider the notion of a default. We can note several facts about a default such as "'birds fly". First, it asserts something about the external world: "birds fly" is clearly true while equally "cows fly" is false. As well, if one accepts that "birds fly" is true, then one rationally would accept other defaults, such as "birds fly or swim". Second, the truth or falsity of a default is independent of there being any believers. If human beings (and their knowledge bases) were to disappear, birds would still fly, and "birds fly" would still be true. Third, a default expresses a property of individuals belonging to a particular class, while not mentioning any specific individuals. Finally, a default says nothing about how one may obtain a default conclusion about an individual.

This suggests that a default asserts a general property about members of a class, and it is the task of nonmonotonic inference to draw conclusions about specific individuals based on a collection of defaults. So we can distinguish a default *assertion* or *proposition*, from an *inference* involving a default. To spell things out:

- A default is either *true* or *false*. One may derive defaults from a set of defaults; however such derivations say nothing about particular individuals nor properties of specific individuals.
- A default inference is either *sound* or *unsound* (or better perhaps, since we're dealing with nonmonotonicity, *rational* or *not rational*). A default inference ascribes a property to an individual.

Consequently we distinguish two types of reasoning:

- With defaults (as assertions).

  This is the realm of conditional logics of normality. Nonmonotonic inference relations can also be seen in this light (given the correspondence results between conditional logics and nonmonotonic inference relations).

- Applying defaults, to give conclusions about individuals.

  This is the realm of Default Logic, and circumscription, along with the rational closure for rational nonmonotonic inference relations.

Clearly, "traditional" approaches to nonmonotonic inference (as exemplified by Default Logic and circumscription) have nothing to say about the meaning of a default, and arguably *this* is what has led to issues with obtaining the "right" inference. Conditional logics, and by extension nonmonotonic inference relations, deal with the meaning of a default as an objective entity talking about classes; it is not surprising then, as we later discuss, that such approaches have difficulties when it comes to expressing default properties about individuals.

An interesting distinction that can be made concerning default assertions compared to default inference, is that the former is essentially *semantic* while the latter is *syntactic*, in the following sense: Defaults are things that are either true or false, with respect to some larger theory. Two formulas that are true under precisely the same conditions express the same thing, and the fact that they may be written differently can be seen as an irrelevant syntactic commitment. Hence, in a logic of defaults, the formulas $(Bird \Rightarrow Fly) \wedge (Bird \Rightarrow Nest)$ and $Bird \Rightarrow (Fly \wedge Nest)$ are true in exactly the same models, and so express the same *proposition*. Arguably this is as things should be.

This is not the case for default inference. In Default Logic and circumscription, one would expect quite different outcomes given the set of defaults

$$\left\{\frac{Bird(x)\ :\ Fly(x)}{Fly(x)},\ \frac{Bird(x)\ :\ Nest(x)}{Nest(x)}\right\}$$

and the set

$$\left\{\frac{Bird(x)\ :\ Fly(x) \wedge Nest(x)}{Fly(x) \wedge Nest(x)}\right\},$$

or the two circumscriptive theories:

$$\{\forall x.(Bird(x) \wedge \neg Ab_F(x)) \supset Fly(x), \quad \forall x.(Bird(x) \wedge \neg Ab_N(x)) \supset Nest(x)\}$$

and

$$\{\forall x.(Bird(x) \wedge \neg Ab_{FN}(x)) \supset (Fly(x) \wedge Nest(x))\}.$$

In particular, if one knows of a bird that it doesn't fly, the former theories would allow one to conclude that it nonetheless builds nests. The overall observation then is that a default, as an instrument for inference, is a syntactic notion; "logically equivalent" sets of defaults may give different default conclusions. Again, this is as things should be.

We next examine this distinction with respect to how approaches to default reasoning address first-order issues.

3.1.1 Defaults and First-Order Concerns

We can observe that virtually all approaches to default reasoning have problems (or at least a certain awkwardness) in the first order case. Indeed, first-order issues are often ignored (with the possible exception of circumscription) in that defaults are usually expressed in a propositional language, as

$$\frac{Bird : Fly}{Fly} \quad \text{or} \quad (Bird \wedge \neg Ab_F) \supset Fly \quad \text{or}$$

$$Bird \mathrel{|\!\sim} Fly \quad \text{or} \quad Bird \Rightarrow Fly$$

It is unclear what is meant in these cases, unless a default is understood as applying to a specific individual. That is, the rule $\frac{Bird : Fly}{Fly}$ only makes sense if $Bird$ is regarded as standing for $x\text{-}is\text{-}a\text{-}Bird$ for understood individual x (and similarly for Fly). Indeed, in the literature, this is just how such propositional glosses are taken, with the understanding that first-order issues are orthogonal to whatever a particular paper at hand is about.

In Default Logic and circumscription, there is no problem expressing a default in first-order terms. Thus in Default Logic, one can write $\frac{Bird(x) : Fly(x)}{Fly(x)}$. This rule applies to instances only, and so can be regarded as a rule schema, standing for the set of its ground instances. This has been a point of criticism of Default Logic in the past. However, given the distinction between default assertions and default inference, such a criticism, at least with regards to reasoning about default properties, seems misplaced: Default Logic has nothing to say about a default as an assertion (i.e. default rules are not things that can be true or false) but rather solely concerns inference; as we suggest below, default inference is most appropriately regarded as involving individuals. Similarly, in circumscription we can write $\forall x.(Bird(x) \wedge \neg Ab_F(x)) \supset Fly(x)$. Hence every bird, except for known exceptions, flies. As an assertion about the world, this is clearly false,[10] unless one is talking about a constrained domain such as the birds at some zoo. Nonmonotonic inference relations on the other hand have representational problems in the first-order case, since if the symbol $\mathrel{|\!\sim}$ stands for an inference relation, it is not clear what an expression $Bird(x) \mathrel{|\!\sim} Fly(x)$ would mean. In particular, there is no formal relation between the occurrences of free variable x on either side of the $\mathrel{|\!\sim}$ symbol, although informally there is.

[10] In this regard then, the circumscription of such a formula has the same epistemic flavour as autoepistemic logic, in that it appears to talk about individuals not *known* to be exceptional.

On the other hand, in a conditional logic, since a default is part of the object language, there is no problem in adding quantification. However, it is not immediately clear how semantically this should be carried out. Something like $\forall x.Bird(x) \Rightarrow Fly(x)$ is problematic [Delgrande, 1998]; as well, intuitively this formula doesn't seem to capture the idea that birds fly since, among other things, for any bird x it isn't the case that x normally flies (a penguin doesn't for example). To address issues concerning quantification and modalities, [Delgrande, 1998] suggests that the conditional connective $\Rightarrow$ be a variable-binding operator, and so our canonical example would be expressed $Bird(x) \Rightarrow_x Fly(x)$.

Yet another alternative is to embrace *concepts* as objects in a domain of discourse, and declare that $B \Rightarrow F$ is a formula in some logic of concepts. Since the area of description logics can be seen as addressing (monotonic) logics of concepts, a possible course of action is to define a description logic for defaults. A goal then would be to define a suitable notion of "default subsumption", writing something like $Bird \sqsubseteq_d Fly$. However, in these cases we are back to regarding defaults as assertions, and so inference regarding individuals would be a separate issue.

So to conclude this subsection, we suggested previously that defaults are expressed at a level independent of individuals, and that inferences *about* defaults are similarly independent of individuals. Moreover, defaults concern open domains, that is they encompass all past, present, future, and possible individuals. On the other hand, default inference concerns *specific* individuals, and, putting it more strongly, default inference involves reasoning about a specific individual or individuals. Thus, the argument: "*Adults are normally employed; therefore adults are normally employed or happy*" is independent of any particular individual. A default conclusion about (adult) Chris is a different matter and is on a different level. Otherwise, if these levels are conflated, this can lead to undesirable conclusions such as the example from circumscription "*every bird except for the known exceptions flies*".

It can be observed that this is exactly the same situation that one has in databases. Thus for example a relation schema is defined prior to there being any database instances, and integrity constraints provide general constraints, and are also expressed independently of a database instance. Querying a database involves reasoning about individuals, in that for a simple query, a set of instances satisfying the query is returned.

3.2 Defaults: Rules vs. Conditionals

If we consider default inference, there are two distinct interpretations of a default with respect to its applicability. On the one hand, applying a default can be regarded as employing something like a defeasible rule of inference. Default logic falls into this category and, indeed, defaults in Default Logic have been referred to as "domain specific rules of inference". Hence if one knows that an individual is a bird then, lacking information to the contrary, one concludes that it flies. If one knows of an individual that it does not fly, then nothing can be concluded about birdhood.

On the other hand, applying a default can be regarded as reasoning with a weak or defeasible material conditional. Circumscription is in this category; for a formula such as $(Bird \wedge \neg Ab_F) \supset Fly$, if circumscribing yields that Ab_F is false, then one ends up effectively with a material conditional $Bird \supset Fly$ which, if $Bird$ is true, allows one to deduce Fly. And if one knows that $\neg Fly$ is true, one can conclude $\neg Bird$. Without going into details, the standard way of closing a (rational) nonmonotonic inference relation, given by the rational closure, also exhibits material-conditional-like behaviour in the absence of exceptional conditions as does, for example, conditional entailment.

Both interpretations have received criticisms or can be shown to lead to unfortunate properties. For example, in the case of Default Logic, one cannot reason by cases: given that birds normally fly, as do bats, and given that an individual is either a bird or bat, one cannot conclude that it flies. However, approaches that behave like weak material conditionals also have difficulties. Consider the defaults that if someone gets a salary increase then they're normally happy, and if they break their leg then they're not happy; also assume that $Chris$ gets a salary increase. In a circumscriptive abnormality theory this can be expressed as follows:

$$\begin{aligned} &\forall x.(Raise(x) \wedge \neg Ab_R(x)) \supset Happy(x),\\ &\forall x.(BreakLeg(x) \wedge \neg Ab_B(x)) \supset \neg Happy(x),\\ &Raise(chris) \end{aligned}$$

Given nothing else, we conclude $\neg BreakLeg(chris)$. This is clearly an undesirable consequence. We get similar problems with the rational closure, conditional entailment, and other such approaches.[11] This last example also appears to be fatal. In circumscription, for example, it is not at all clear how this behaviour can be blocked, or even if it *can* be blocked. Consequently, we take this example as being decisive and so accept that:

> *With respect to defaults, inference is rule-like and not (material) conditional-like.*

4 Defaults as Naïve Scientific Theories

The conclusion of the previous section is problematic: Most approaches to default reasoning fall into the weak-material-conditional category and those that don't, namely Default Logic and related systems, provide only a basic mechanism for inference that does not reflect how one would wish to reason with defaults. This then suggests that it would be instructive to first study the *phenomena* modelled by approaches to defaults – that is, consider what it is in the world that's being represented, and then use such a study to drive

[11] A possible rejoinder is that such approaches are not intended to be used for reasoning about normality properties of individuals. Such a rejoinder is well taken. However, different approaches may nevertheless be examined with respect to their overall applicability in different situations.

a study of what constitutes *desirable* default inferences. To this end, in this section I argue that defaults are best regarded as statements in a naïve scientific theory. (For an excellent discussion, see [Putnam, 1975].) From this I argue that *relevance* is the key notion needed to formalise default inference.

4.1 What *is* a Default?

Assume that we live in a Newtonian universe, and consider the following assertion.

Example 1 *Planets move in ellipses.*

In our Newtonian universe, this statement would be accepted as true. However, on the other hand, no planet would *ever* be observed to move in an ellipse. Rather, if one plotted the path of a planet, it would be observed to more or less follow an ellipse. If asked about this discrepancy, an astronomer would excuse the error by saying that the measuring instruments weren't exact, or that there was atmospheric interference, or that there were other bodies whose gravitational influence needed to be taken into account, or some such conditions interfered with the observations. If pressed, the astronomer might assert that if the universe consisted only of a star and its orbiting planet, *only then* would the planet would move in a perfect ellipse. Nevertheless (back in the real world), the statement "planets move in ellipses" is nonetheless useful: it can be accepted as true, in that any deviation from an ellipse can be explained in principle by other real or hypothesised bodies, instrument errors, etc. Moreover the statement has predictive value: the orbit of the moon can still be calculated very accurately, and in fact deviations in Uranus' orbit led to the discovery of Neptune.

"Planets move in ellipses" is clearly a scientific assertion, and can be considered as part of a naïve scientific theory, in that it is a qualitative outcome of a more precise expression (using the inverse square law of gravitation) of an underlying theory. As well, it clearly has the flavour of a default.

Consider the next example:

Example 2

1. *Brass doorknobs disinfect themselves of bacteria within eight hours.*
2. *Copper conducts electricity.*
3. *Copper has atomic number 29.*

The first statement certainly sounds like a default, as does the second. Both in fact are true[12] though both allow exceptions. Copper wire immersed in water does not conduct electricity, for example. However, the third statement is quite different in character. In particular, it is *definitional*, and specifies an *essential* property of copper. (Thus a mass of atoms each with atomic number 28 isn't an exceptional chunk of copper; rather it is nickel.) In fact, the properties of

[12]Brass is a copper compound and copper has germicidal properties.

copper can be determined or justified via its atomic structure. That is, atomic structure allows a precise definition of an element, and this can be used to give an *account* of the previous default statements.

So arguably, a default such as “birds fly” is an assertion in a naïve scientific theory, the same way that “copper conducts electricity” is. Similarly, “birds fly” asserts that in some sense *flight* is part of the meaning of *bird.* Or, phrased differently, if we had a complete theory of birds, we could exactly account for flight as a property of birds. Note that, by this account, it is possible for “birds fly” to be true, while no existing bird in fact flies. So while this provides a means of determining the meaning of “birds fly”, it has nothing to say about default inference.

4.2 Representing Defaults

The previous section raises the question: how does one reason about sentences in a “naïve scientific theory”? Arguably there is already a logic for defaults, given by the so-called “conservative core” [Pearl, 1989] and (re)discovered by various researchers or appearing under different guises, see [Adams, 1975, Pearl, 1988, Kraus, Lehmann, and Magidor, 1990, Lamarre, 1991, Boutilier, 1992, Dubois, Lang, and Prade, 1994], among others. This also is the system of preferential reasoning [Kraus et al., 1990] referred to earlier with respect to nonmonotonic inference relations. In the version described here, a (weak) conditional operator $\Rightarrow$ is introduced into propositional logic, as we’ve already seen. The operator $\Rightarrow$ is a binary modal operator, and its semantics is given in terms of a standard Kripke structure, where the accessibility relation is given by a preorder over possible worlds. This preorder reflects a notion of relative normality between possible worlds. Then, informally, $\alpha \Rightarrow \beta$ is true at a world, just if β is true at all least α worlds. The intuition then is that $\alpha \Rightarrow \beta$ is true at a world, just if, looking at the “most normal” α worlds, β is true at all these worlds. Hence “birds fly” is true just if in the least exceptional worlds (and so ignoring things like being a penguin, having a broken wing, etc.) in which there are birds, birds fly. Formally there is little to add:[13]

Sentences are interpreted in terms of a *model* $M = \langle W, \leq, P\rangle$ where:

1. W is a set (of possible worlds),
2. the accessibility relation $\leq \subseteq W \times W$ is transitive and reflexive, and
3. $P : \mathbf{P} \mapsto 2^W$.

Truth conditions for the standard connectives are as in propositional logic, while for the weak conditional we have:

$$\models^M_w \alpha \Rightarrow \beta \text{ iff: for every } w_1 \in min(\alpha, \leq) \text{ we have } w_1 \models \beta$$

[13] The point in providing a sketch of a formal development isn’t necessarily to establish a definitive logic for defaults; while it (or a slightly stronger logic) is the accepted account, it is possible that someone will come along with a superior account of defaults. If this were the case then the discussion here would remain unchanged.

where $min(\alpha, \leq)$ is the set of least worlds according to $\leq$ in which α is true.

So based on this and the previous discussion, we can regard a default as a *counterfactual normative* statement. That is "birds fly" can be interpreted as, "for any individual bird x, in the most normal of possible affairs, x would fly" or "if x were a normal bird, then x would fly". One can then make a nonmonotonic inference by assuming that, given a set of defaults, states of affairs are ranked as normal as consistently possible with those defaults, and that given contingent information, the actual world is among those ranked least in which the contingent information is true.[14]

The difficulty is that this doesn't quite work for inference. Thus, given that one agrees that a normal bird flies, has feathers, builds a nest, etc. then if one knows only that an individual is a bird, then indeed it will be concluded that the individual flies, has feathers, builds a nest, etc. A problem arises however if one knows that an individual bird does not fly. Then this individual can't be a normal bird, and so one can't use the default that normal birds build nests. This suggests that nonmonotonic inference based on a notion of strict minimality of worlds, based in turn on aggregated normality information, is not entirely appropriate for inference involving defaults. The next subsection proposes an alternative.

4.3 Reasoning with Defaults of Normality

Consider again naïve scientific theories, and consider a length x of copper wire. We would conclude that x conducts electricity if

1. we had no further information;
2. we knew only that it was mined at Copper Mountain;[15] or
3. we knew only that it had bends in it.

We would not conclude that x conducts electricity if

1. we tested it and it didn't conduct electricity;
2. we knew it had significant impurities; or
3. it was immersed in water.

Thus, for limiting cases, if all we knew was that the antecedent of a default were true then we would apply the default; if we knew that the consequence were false, we would not apply it. Otherwise, one might note that the location where the wire was mined and the fact it has bends are *irrelevant* with respect to conducting electricity, while the presence of impurities and water are clearly *relevant*. So essentially we want to say of x that it conducts electricity if there is nothing known that is *relevant to it not conducting electricity*. This indicates

[14] In fact this is a sketch of the intuitions underlying [Pearl, 1990]. The rational closure [Lehmann and Magidor, 1992] is founded on differing intuitions, though in a strong sense the same inferences are obtained. [Lehmann and Magidor, 1992, p. 28] also suggests that "any reasonable system should endorse any assertion contained in the rational closure".

[15] in southern British Columbia, Canada where indeed copper is mined.

that a theory of *relevance* (or perhaps *reasons* [Horty, 2010]) is the appropriate notion needed for default inference.

Relevance An incorporation of relevance represents a shift in how defaults would be handled. Previously, for default $\alpha \Rightarrow \beta$ (or nonmonotonic inference relation $\alpha \mid\!\sim \beta$), default inference was effected by assuming that the present state of affairs was as normal as consistently possible. Thus if one knew of an individual only that it was a bird, then one would conclude by default that it flies. If one knew also that it didn't build nests, then the normality assumption would no longer hold (since a non-nest-building bird isn't normal) and so either one would lose the inference about flying, or else it would have to be restored by other means, such as the lexicographic closure [Benferhat, Cayrol, Dubois, Lang, and Prade, 1993, Lehmann, 1995].

Relevance on the other hand appears most naturally, at least in this context [Delgrande and Pelletier, 1998], to be a ternary relation: one might say for example that having a broken wing is relevant to a bird flying. Intuitively, a default $\alpha \Rightarrow \beta$ provides a means of concluding β from α. Roughly, one would want to say that β can be concluded on the basis of α if there is nothing blocking the inference, or if there is no *reason* that the inference should be blocked, or if there is nothing *relevant* that would lead one to not draw the conclusion given in the consequent.

It might seem that this line of intuitions ultimately leads back to circumscriptive abnormality theories perhaps, or consistency conditions as found in Default Logic. Thus, so the argument might run, in an abnormality theory one says that a bird flies unless it is in some fashion abnormal with respect to flight and surely (so the argument might run) this is just another way of saying that there is no known reason for it to not fly, or there is nothing relevant known concerning flight. While it is true that things could be phrased in terms of abnormality (or, for that matter, consistency), there is a big difference: we are now working within a logic of conditionals, and not an augmentation of classical logic.

Consider how we might formalise a notion of relevance toward default inference. In outline, one wants to say something like:

Informal Definition:

Given: a set of defaults $\mathcal{T}$ and facts $\mathcal{F}$.

Conclude β by default if:

1. There is $\alpha \Rightarrow \beta$ where $\mathcal{T} \models \alpha \Rightarrow \beta$ and $\mathcal{F} \models \alpha$.
2. If $\mathcal{T} \models \gamma \Rightarrow \neg\beta$ where $\mathcal{F} \models \gamma$ then $\mathcal{T} \cup \mathcal{F} \models \alpha \supset \gamma$.

The informal definition says that β can be concluded if

1. there is a reason to do so, and
2. there is nothing *relevant* blocking the inference.

Consider where we have the defaults that birds fly, animals do not fly, and birds with broken wings don't fly, along with the fact that birds are animals:

$$Bird \Rightarrow Fly,\ Animal \Rightarrow \neg Fly,$$
$$Bird \wedge BrokenWing \Rightarrow \neg Fly,\ \Box(Bird \supset Animal)$$

If we are given that an individual is a bird (α in the informal definition), then we will conclude that it flies (β). Although the individual is also an animal (γ) and animals don't fly, the notion of being a bird is more specific than that of being an animal (i.e. $\mathcal{F} \models \alpha \supset \gamma$). On the other hand, if the bird has an injured wing, then clearly there is relevant information (viz. $\gamma = Bird \wedge BrokenWing$) as to why it does not fly by default.

Relevance: Other Conditionals In the previous section, we described the "standard" logic of defaults, given as a specific conditional logic. This logic is but one of a large family of conditional logics, where conditional logics have been used also to formalise notions including counterfactuals, deontics, hypotheticals, causality, etc. [Lewis, 1973, Chellas, 1980, Nute, 1984]. Although such notions haven't received the attention of normality defaults, it seems clear that one can reason by default in any such logic. Thus, for example, if one should not speed when driving a car, but that it is permissible to speed if it allows one to avoid an accident, then if in fact one is driving then it is a reasonable conclusion, all other things being equal, that one should not speed.

As well, a notion of *relevance* seems equally pertinent in reasoning with these other conditionals. Thus although in general one should not speed, but one may speed if speeding allows one to avoid an accident, then avoiding an accident is a relevant factor in determining how fast one may drive. On the other hand, the colour of one's car is not relevant, and the fact that one is late for an appointment should not be relevant. This suggest that *relevance* is *the* appropriate mechanism for weak conditionals and default inference in general. As well, it seems that a general account of relevance may lead to a satisfactory account for default reasoning, as well as defeasible reasoning with counterfactuals, deontics, causality, etc. Last, it can be noted that the informal definition given above is expressed independently of any specific logic, and so an overarching account of relevance (as a ternary relation with respect to weak conditionals) may provide a unifying framework for an extended notion of defeasible reasoning.

5 Conclusion

This paper has examined the notion of (normality) defaults in nonmonotonic reasoning. We noted that early work, as exemplified by Default Logic and circumscription, focussed on developing inference mechanisms and then on using such mechanisms to try to suitably encode reasoning with defaults. Similar remarks apply to subsequent work, represented by applications of nonmonotonic inference relations and conditional logics to defaults, even though such work began with a semantic account of defaults. Since no extant approach satisfactorily captures default reasoning, we suggested that a suitable strategy is to step back and consider first the phenomenon that is being modelled, that is, determining what a default such as "birds fly" means. This requires distinguishing the *representation* (or assertion) of a default from an inference involving the application of the default. The former is either true of false, while the latter is (in the case of defaults) rational or not rational.

Along the way, I suggested or noted that:

1. A default is essentially a semantic notion, in that a default or set of defaults can be replaced by logically equivalent defaults without altering the meaning of the theory. Default inference on the other hand is syntactic, in that replacing a set of defaults with an equivalent set of defaults may result in different default conclusions.

2. A general first-order default, as a proposition, is independent of specific individuals and applies to open domains. Default reasoning on the other hand concerns inference of properties of individuals. Thus, "birds fly" is expressed independently of any individual, but default inference concerns specific individuals. Reasoning about a population as a whole is best approached at the level of reasoning *about* defaults. Mixing these two levels yields conclusions such as "every bird except known exceptions flies".

   This split is analogous to that in database systems where the database schema and integrity constraints corresponds to defaults (along with other general information concerning a domain), while querying a database instance corresponds to default reasoning about individuals over contingent information.

3. Default reasoning is analogous to reasoning with a rule, not a version of a weak (material) conditional. This suggests that we have a ways to go with respect to getting the inferences right, since most approaches to default inference are closer to the weak-(material-)conditional interpretation. This latter group includes approaches using circumscription, as well as conditional entailment, and the rational closure and related approaches. Default Logic obviously involves rules, but in this case one has a general inference mechanism only, but where there is little connection between the inference mechanism and how one might want to reason with defaults.

I argued that an appropriate theory of defaults involves adopting (or specifying) an appropriate *conditional logic* for representing normality defaults. This logic might well be provided by the so-called "conservative core", or a slightly stronger variant given by a conditional logic based on a notion of normality reflected by a total preorder over possible worlds.[16]

Given such a logic, one can then ask *What are the principles that justify a default inference?* For normality defaults, default inference hinges on a formalisation of *relevance* or *reasons*. As well, this notion of founding default

[16]And so corresponding to the *rational* systems of [Lehmann and Magidor, 1992].

inference on relevant properties also appears applicable to the full range of conditional logics, and so applicable to approaches for reasoning with counterfactuals, norms, deontics, causality, etc.

Acknowledgements

I would like to thank Marc Denecker, whose discussions on this topic got me thinking about the subject of this paper. As well, I thank Jeff Horty and Jeff Pelletier along with the other Nonmon@30 participants (and in particular Didier Dubois, Alex Bochman and Daniel Lehmann) for helpful and often enlightening discussions. I thank a reviewer of this paper for helpful comments. The fact that "snow is white" is a generic statement is noted in [Carlson and Pelletier, 1995], from where I got it. Finally, thanks to Gerd, Victor, and Mirek for the terrific conference that led to this book.

References

Special issue on nonmonotonic reasoning. *Artificial Intelligence*, 13(1-2), 1980.

E. Adams. *The Logic of Conditionals.* D. Reidel Publishing Co., Dordrecht, Holland, 1975.

G. Antoniou. *Nonmonotonic Reasoning.* MIT Press, 1997.

C. Baral. *Knowledge Representation, Reasoning and Declarative Problem Solving.* Cambridge University Press, 2003.

S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 640–645, Chambéry, Fr., 1993.

C. Boutilier. *Conditional Logics for Default Reasoning and Belief Revision.* PhD thesis, Department of Computer Science, University of Toronto, 1992.

C. Boutilier. Conditional logics of normality: A modal approach. *Artificial Intelligence*, 68(1):87–154, 1994.

G. Brewka. Cumulative default logic: In defense of nonmonotonic inference rules. *Artificial Intelligence*, 50(2):183–205, 1991a.

G. Brewka. *Nonmonotonic Reasoning: Logical Foundations of Commonsense.* Cambridge University Press, Cambridge, 1991b.

G. Brewka, I. Niemela, and M. Truszczynski. Nonmonotonic reasoning. In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, pages 239–284. Elsevier Science, San Diego, USA, 2007.

Gregory N. Carlson and Francis Jeffry Pelletier, editors. *The Generic Book.* University of Chicago Press, Chicago, 1995.

B.F. Chellas. *Modal Logic*. Cambridge University Press, 1980.

J.P. Delgrande. An approach to default reasoning based on a first-order conditional logic: Revised report. *Artificial Intelligence*, 36(1):63–90, 1988.

J.P. Delgrande. On first-order conditional logics. *Artificial Intelligence*, 105 (1-2):105–137, 1998.

J.P. Delgrande and J. Pelletier. A formal analysis of relevance. *Erkenntnis*, 49 (2):137–173, 1998.

J.P. Delgrande and T. Schaub. Compiling specificity into approaches to nonmonotonic reasoning. *Artificial Intelligence*, 90(1-2):301–348, 1997.

J.P. Delgrande, T. Schaub, and K. Jackson. Alternative approaches to default logic. *Artificial Intelligence*, 70(1-2):167–237, October 1995.

Marc Denecker, Victor W. Marek, and Miroslaw Truszczyński. Uniform semantic treatment of default and autoepistemic logics. *Artificial Intelligence*, 143(1):79–122, 2003.

D. Dubois, J. Lang, and H. Prade. Possibilistic logic. In D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Nonmonotonic Reasoning and Uncertain Reasoning*, volume 3 of *Handbook of Logic in Artifical Intelligence and Logic Programming*, pages 439–513. Oxford, 1994.

H. Geffner and J. Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artificial Intelligence*, 53(2-3):209–244, 1992.

M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In R. Kowalski and K. Bowen, editors, *Proceedings of the Fifth International Conference and Symposium of Logic Programming (ICLP'88)*, pages 1070–1080. The MIT Press, 1988.

M. Gelfond and V. Lifschitz. Classical negation in logic programs and deductive databases. *New Generation Computing*, 9:365–385, 1991.

S. Hanks and D. McDermott. Default reasoning, nonmonotonic logics, and the frame problem. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 328–333, 1986.

J.F. Horty. Some direct theories of nonmonotonic inheritance. In D.M. Gabbay, C.J. Hogger, and J.A. Robinson, editors, *Nonmonotonic Reasoning and Uncertain Reasoning*, volume 3 of *Handbook of Logic in Artifical Intelligence and Logic Programming*, pages 111–187. Oxford: Clarendon Press, 1994.

John F. Horty. Reasons as defaults. Draft, 2010.

S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1-2):167–207, 1990.

H.E. Kyburg, Jr. Believing on the basis of evidence. *Computational Intelligence*, 10(1):3–20, 1994.

P. Lamarre. S4 as the conditional logic of nonmonotonicity. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning*, pages 357–367, Cambridge, MA, April 1991.

D. Lehmann. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15(1):61–82, 1995.

D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60, 1992.

D. Lewis. *Counterfactuals.* Harvard University Press, 1973.

V. Lifschitz. Computing circumscription. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 121–127, Los Angeles, 1985.

W. Łukaszewicz. Considerations on default logic: An alternative approach. *Computational Intelligence*, 4(1):1–16, Jan. 1988.

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

J. McCarthy. Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence*, 28:89–116, 1986.

A. Mikitiuk and M. Truszczyński. Rational default logic and disjunctive logic programming. In A. Nerode and L. Pereira, editors, *Proceedings of the Second International Workshop on Logic Programming and Non-monotonic Reasoning.*, pages 283–299. The MIT Press, 1993.

R.C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25:75–94, 1985.

D. Nute. Conditional logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 2, pages 387–439. D. Reidel Pub. Co., 1984.

J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufman, San Mateo, CA, 1988.

J. Pearl. Probabilistic semantics for nonmonotonic reasoning: A survey. In *Proceedings of the First International Conference on the Principles of Knowledge Representation and Reasoning*, pages 505–516, Toronto, May 1989. Morgan Kaufman.

J. Pearl. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In R. Parikh, editor, *Proc. of the Third Conference on Theoretical Aspects of Reasoning About Knowledge*, pages 121–135, Pacific Grove, Ca., 1990. Morgan Kaufmann Publishers.

D.L. Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36(1):27–48, 1988.

D.L. Poole. The effect of knowledge on belief: Conditioning, specificity, and the lottery paradox in default reasoning. *Artificial Intelligence*, 49(1-3):281–307, 1991.

H. Putnam. The meaning of 'meaning'. In *Mind, Language and Reality: Philosophical Papers Volume II*, pages 215–271. Cambridge University Press, 1975.

R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1-2):81–132, 1980.

E. Rosch. Principles of categorisation. In E. Rosch and B.B. Lloyds, editors, *Cognition and Categorisation.* Lawrence Erlbaum Associates, 1978.