Reprinted from

HANDBOOK OF LOGIC AND LANGUAGE

edited by

Johan van BENTHEM
University of Amsterdam & Stanford University

Alice ter MEULEN
Indiana University, Bloomington, Indiana

1997

ELSEVIER
AMSTERDAM · LAUSANNE · NEW YORK
OXFORD · SHANNON · TOKYO

1997

THE MIT PRESS
CAMBRIDGE, MASSACHUSETTS

CHAPTER 20

Generics and Defaults*

Francis Jeffry Pelletier

Department of Philosophy and Department of Computing Science, University of Alberta, Edmonton, Alta., Canada T6G 2E5
E-mail: jeffp@cs.ualberta.ca

Nicholas Asher

316 Waggener Hall, The University of Texas, Austin, TX 78712, USA
E-mail: nasher@bertie.la.utexas.edu

Contents

*This paper is strongly influenced by – indeed, identically overlaps at some points – Krifka, Pelletier, Carlson, Ter Meulen, Chierchia, and Link (1995). Not all those authors agree with everything we say here, but it was through interaction with them that most of the linguistic ideas behind this paper came into existence, and we gratefully acknowledge their help. The paper also owes some of its content to work done jointly with Len Schubert: Schubert and Pelletier (1987, 1989). Other parts of this paper overlap work done jointly with Michael Morreau: Asher and Morreau (1991, 1995). We also wish to express our appreciation to the editors of this volume, Alice ter Meulen and Johan van Benthem, for their encouragement and for their comments on earlier versions of this paper.

HANDBOOK OF LOGIC AND LANGUAGE
Edited by J. van Benthem and A. ter Meulen

1. Linguistic and epistemological background

1.1. Generic reference vs. generic predication

Historically speaking, there have been two quite distinct phenomena that have been called genericity. The first is reference to a kind, a *genus*, as exemplified in (1). The italicized NPs in (1) do not denote or designate some particular potato or group of potatoes, but rather the kind Potato (*Solanum tuberosum*) itself. In this usage a generic NP is one which does not refer to an ordinary individual or object, but instead refers to a kind.

(1) a. *The potato* was first cultivated in South America.

b. *Potatoes* were introduced into Ireland by the end of the 17th century.

c. The Irish economy became dependent upon *the potato.*

NPs such as *potatoes* or *the potato* in these sentences are called kind-referring NPs (sometimes generic NPs), as opposed to object-referring NPs, and the predications in sentences involving such NPs are called kind predications as opposed to object predications.[1]

The second phenomenon commonly associated with genericity are propositions which do not express specific episodes or isolated facts, but instead report a kind of *general property*, that is, report a regularity which summarizes groups of particular episodes or facts. One common sort of generic proposition of this type occurs when the regularity holds across the individual instances of a kind; in such a case the regularity is stated as one holding for the kind itself. Examples can be found in the natural readings of the sentences in (2). Here (2a,b,c) do not state something about a specific potato, but about potatoes in general. This notion of genericity is clearly a feature of the whole sentence (or clause), rather than of any NP in it; it is the whole generic sentence that expresses regularities which transcend particular facts.

(2) a. A potato contains vitamin C and amino acid.

b. Potatoes contain protein.

c. The potato contains thiamin.

Closely related to such sentences as these are habituals, in which a regularity of action is predicated of an ordinary individual. In such cases we (mostly) intend that the regularity is predicated of the object at different times. An example is in (3), which does not report a particular episode, but a habit.

(3) Mary smokes a cigar after dinner.

---

[1] Sentences employing such kind-referring NPs are possibly what Barth (1974) has in mind by "logophoric use of the articles". She says that "*The/An M is P* is a logophoric proposition iff the article is not being used anaphorically, the NP is neither used referentially nor attributively, and this sentence does *not* mean any of:

all Ms are P; every M is P; an arbitrary M is P;

the class/set of Ms is included in the set of Ps; some M is P.

And furthermore The/An M is M has to be regarded as logically true."

It is not clear whether or not there are any logophoric sentences in this sense, at least not if some of the proposed reductions of generic NPs to quantified statements are correct.

The difference between (2) and (3) is a matter of whether one is generalizing over objects or over what an object does. In (2) we generalize over objects and say that a typical instance of such an object has a property, whereas in (3) we generalize over what an object does and say that it typically acts in a certain way. Notice then an ambiguity in such sentences as

(4) Violinists smoke after concerts

which can be seen saying either that the typical violinist has the property of smoking after each concert or as saying of all violinists that they have the habit of typically smoking after concerts. Of course it can even more naturally be seen as generalizing over *both* objects and actions, saying that the typical violinist has the property of typically smoking after concerts. Sentences like (2), (3) and (4) are called characterizing sentences (or sometimes generic sentences) since they express a characterizing property. They are opposed to particular sentences, which express particular events, properties of particular objects, and the like.[2] We therefore can distinguish characterizing predications and particular predications. Other common terms for characterizing sentences found in the literature are (g)nomic, dispositional, general, or habitual. Much of our knowledge of the world, and most of our beliefs about the world, are couched in terms of characterizing sentences. According to this attitude such sentences would be either true or false, *not* indeterminate or figurative or metaphorical or sloppy talk.

Obviously, sentences containing reference to kinds have something in common with characterizing sentences: when we refer to kinds we abstract away from particular objects, whereas when we use characterizing sentences we abstract away from particular events and facts. Furthermore, it seems natural that one way to express a general law or regularity about the specimens of a kind is to state it for the kind itself. Nonetheless, as it is argued in detail in (Krifka, Pelletier, Carlson, Ter Meulen, Chierchia and Link, 1995), it is necessary to distinguish them, since not only are there intuitive differences concerning their logical form but also there are linguistic differences between them. In this article we concentrate exclusively on characterizing sentences, since it is the phenomenon manifested by such sentences that has attracted the most interest from the logical representation wing of the artificial intelligence world. In this article we will not be concerned with the problem of delimiting characterizing sentences from kind-referring predications, nor with distinguishing all sorts of generic phenomena from non-generic phenomena. Instead we will merely assume that we know what sentences are the type of generic sentences in which we are interested; and we will proceed to investigate their semantic and logical properties. Our interest is in generic sentences of the characterizing sort, and it is these sentences that we will call generic statements in what follows.

1.2. Why are there any generic sentences at all?

A crucial feature of characterizing sentences indeed, perhaps their most interesting semantic feature and that feature which brings them to the attention of this volume is that

[2] They are also opposed to general sentences in which there is explicit quantification over particular objects manifesting particular events or properties such as *Each potato in this room was grown in Alberta.*

they admit exceptions.[3] That is, sentences in (2) and (3) are true even though there are potatoes which do not contain those nutrients and even though Mary sometimes does not smoke a cigar after dinner. In our opinion, it is precisely this feature that explains why all natural languages have generic sentences. It is an innate disposition of humans ... perhaps this is a consequence of having intelligence or perhaps it is one of humanity's survival characteristics ... that they desire to understand and characterize the world immediately surrounding them. People notice regularities in nature and form (what we might call) folk-laws to codify these regularities and to predict what the future will bring. So, although not all potatoes contain the nutrients mentioned in (2), and although Mary does not always smoke cigars after dinner, still such folk-laws are intellectually satisfying and practically useful because Mary (and potatoes) *typically* or *usually* or *normally* or *nomically* ... smokes after dinner (or contain those nutrients). Perhaps it is a feature of having finite, fallible minds that makes us often notice regularities that have exceptions, or perhaps it is more a matter of needing to be able to choose regularities quickly in order to get on with other aspects of our survival. Whatever the underlying reason, the fact is that people notice those regularities that can be used to predict actions of others and changes (or constancies) in one's environment. And such regularities commonly have exceptions; either ones that are noticed later or ones that we think we can safely ignore (for whatever reason).

Once one sees generic sentences and the role they play in guiding ordinary actions and beliefs, one discovers that such a construct is used in many different areas of academic research. For example, one can see deontic laws that have *ceteris paribus* clauses as being generic in this way; one can see presuppositions as being the exceptions to statements that don't count in the literal meaning of the statement; one can see arguments over realism/antirealism in philosophy of science as an argument over whether generic statements (with their exceptions) can ever enjoy the privileged status of a true scientific law; and one can understand the status of (legal) laws as being generic in nature as they admit exceptions and special pleas.

1.3. Generics and exceptions, two bad attitudes

From the semantic point of view it is crucial to the understanding of generics to be able to specify how the generic statement is related to the particular statements that are specific instances of it. For example, we might ask: How many exceptions can a generic statement tolerate?, or more generally, What is the relationship between a generic statement and explicitly quantified statements? In addition to the various reasonable attitudes that different theorist have taken concerning this relationship, there are two bad attitudes that are sometimes found in the literature from which we wish to dissent.

[3] This is not their only interesting semantic feature. In (Carlson, 1977) and (Lawler, 1973) it was pointed out that such generics seem to admit only predications where the property described is in some sense essential to the subject. Thus, *Madrigals are polyphonic* is a true characterizing statement whereas **Madrigals are popular* is not, perhaps because polyphonicity (but not popularity) is essential to madrigals. It is clear that this is related in some way to the nomic nature of generics that we discuss later, but it is far from clear how.

One attitude, an attitude taken by various logic texts and by some earlier AI practitioners, is that the generic statements are strictly speaking false, but are acceptable despite the exceptions because they are close enough to being universally true. So, *Potatoes contain vitamin C* is acceptable despite being literally false, because almost all potatoes, or all the noticeable or important ones, contain vitamin C.[4] We claim that this is not just an unhelpful way to look at the topic, but is totally incorrect. Most of our knowledge about the world is encoded in characterizing sentences, and it seems pointless to claim that such encodings are merely acceptable and not really true. Furthermore, if this were a correct attitude then one would expect that generic statements would be more acceptable or truer the fewer exceptions there be. Yet as we shall shortly see, this is not how generics act.

A second attitude that might be taken toward such exception-tolerating sentences is that they are *neither* true nor false; that these concepts simply do not apply to them. Instead, characterizing sentences are directives or rules as to how to proceed, how to draw inferences, and so on. In such a view, a characterizing statement like *Potatoes contain vitamin C* gives us information about what to expect in our experience, how to act in the presence of certain vegetables, and what to infer about the items in the grocery store namely, we should expect, and act in accordance with, and infer that: in the absence of information to the contrary, any potato we encounter will contain vitamin C. This attitude, though common in certain areas of AI, seems to have two flaws in it. First, in treating characterizing statements in effect as *rules*, as directives about what to infer, it denies truth or falsity to such statements. Instead, the relevant concept is one of truth-preservingness or of validity. Such a view denies that the classic *Snow is white* is either true or false! And it does so not because of the existence of some yellow snow or the like, but because it holds that this sentence is intended only to guide our expectations as to what color of snow we will encounter normally. It was remarked above that much (most?) of our information about the world is captured by these characterizing statements; but it would be a consequence of the present view that this information is not knowledge, since it is not true (indeed, is neither true nor false, is neither accurate nor inaccurate, etc). Instead, it would merely be a matter of directing our actions and inferences. Perhaps there are ways to ease ones conscience about this matter, but it seems very harsh.

The second conceptual difficulty with this approach is that if generic statements are neither true nor false, and instead are rules of inference, then they cannot be embedded within one another. And although some researchers have asserted that this is indeed the case and have used this claim to bolster their position, they fail to notice sentences like

(5) a. Usually, if a person smokes after dinner, he also drinks brandy before bed.

b. It is common that countries which do not honor women's rights also do not honor general human rights.

[4] A variant on this first bad attitude would be to claim that the acceptability of such "strictly speaking false" statements is dependent upon our ability to specify (if asked) what the exceptions are. Although the statements are strictly false, we accept them because we know where to find the exceptions. Our examples below will show that this optimistic outlook on exceptions is totally without merit.

Note that both the antecedent and the consequent of these conditionals are themselves characterizing statements (they each admit of exceptions, they each express regularities over events, etc.), and that each statement asserts there to be a generic connection between the two (the connection admits of exceptions, expresses a general regularity, etc.). Such examples seem to support strongly the view that generics are not directives or rules of inference, but instead are either literally true or literally false. We will consider this attitude and these objections again below, in Section 3.

1.4. Exceptions and generics, some other attitudes

If we adopt the view that generics *can* be either true or false, there is then a question about the logical form of generics. And this brings up once again the question of the relationship between a characterizing sentence and the specific episodes or cases of which it is a generalization. Put bluntly, we wish to know: How many exceptions can a generic statement tolerate and remain true? Different kinds of characterizing generics call for different types of cases to be relevant, and therefore determine different sorts of legitimate exceptions. This can be seen in the differences between episodic generics (where there is a regularity asserted about events),

(6) a. Tabby usually lands on her feet.
b. Marvin normally beats Sandy at ping-pong.

non-episodic generics (where there is a regularity asserted to hold over a class of objects),

(7) a. Bears with blue eyes are normally intelligent.
b. A grade school student is usually a child.

and sentences which can be seen as ambiguous between the two

(8) a. People who have a job are usually happy.
b. People who live far from their work usually drive.

(These last can be seen as saying either (a) that most people who have a job are happy [or: most people who live far from work are drivers], or (b) that a person is happy for most of the time that s/he holds a job [or: a person who lives far from work drives most of the time].)[5] We will not, however, delve into these distinctions. Instead we will concentrate mostly on the naive question of how many exceptions can be tolerated by a characterizing sentence and have it still be true. We will argue that all quantificational analyses fail: that there is no univocal quantifier which can work for all generics. We then briefly examine whether there can be some probabilistic analysis that will work;

[5] The sentences in (7) can also be seen as ambiguous: not only might it be a generalization over the objects in the class (most bears with blue eyes / most grade school children) but also as generalizing over states (such bears are intelligent most of the time but maybe sometimes change into being stupid / a grade school student is a child for most of its life). These sentences should be considered in concert with the earlier examples (2) and (4).

and again we argue in the negative: there is no sense in which a characterizing sentence asserts a unique, or even a rough, probability. Having argued against any relationship – whether quantificational or probabilistic – between generics and their instances, we offer a brief explanation of why this should be so. The failures of the quantificational and probabilistic analyses have consequences for certain AI accounts of the logical form of characterizing sentences, and we shall therefore turn to the topic of these AI techniques.

Consider these true characterizing sentences:

(9) a. Snakes are reptiles.
b. Telephone books are thick books.
c. Guppies give live birth.
d. Italians are good skiers.
e. Crocodiles live to an old age.
f. Frenchmen eat horsemeat.
g. Unicorns have one horn.

Obviously we understand the truth of these sentences as calling for different relative numbers or proportions of instances of the subject terms satisfying the predicate term. In (a) it is all; in (b) most; in (c) some subset of the females; in (d) some small percentage, but a greater percentage than in other countries (or maybe: the very best of the Italian skiers are better than the very best from other countries); in (e) it is strikingly few, since of the hundreds born to one female at a time, most are eaten within a few weeks of birth; in (f) there need be only a very small percentage – from the vantage point of North America, the mere fact of its happening at all is striking; and in (g) no unicorns have one horn. Examples such as these show that there is no univocal quantifier which will serve in all characterizing sentences.

Even attempts to employ vague, probabilistically-oriented quantifiers such as *most* or *generally* or *in a significant number of cases* are misguided. Consider such false characterizing sentences as

(10) a. Leukemia patients are children.
b. Seeds do not germinate.
c. Books are paperbacks.
d. Prime numbers are odd.
e. Crocodiles die before they attain an age of two weeks.

Each of these false characterizing sentences becomes true when prefixed with *In a significant number of cases*. Indeed, the actual extension of the subject within the predicate has almost nothing to do with whether the generic is true or not. Even if all children ever born in Rainbow Lake, Alberta, were left-handed, this would have nothing to do with whether

(11) Children born in Rainbow Lake, Alberta, are left-handed.

To determine the truth of (11) we would instead look to the water supply, or to the mineral deposits near the town, or to the electromagnetic fields surrounding the hospital, or to the ancestry of the parents of such children, etc., to determine whether some special causal factor was at work and made (11) a true characterizing sentence. The actual numbers or percentages of such children that are left-handed is simply not relevant, except as evidence that some other force might be at work. In this, generic statements are akin to scientific laws: they must be carefully distinguished from accidental generalizations. If it were accidentally to happen that all the remaining pandas in the world lost a leg, then even though every panda had three legs, still the characterizing sentence *Pandas have three legs* would be false. A sentence that is an accidental generalization is not a true characterizing sentence.

1.5. Generics and intensionality

Furthermore, characterizing sentences are inherently intensional, and this too is an obstacle to any extensional analysis, whether quantificational or probabilistic. Characterizing sentences like

(12) a. This machine crushes oranges,

b. Kim handles the mail from Antarctica,

c. Members of this club help one another in emergencies,

can be true despite there never being an opportunity for the generic episode to take place. (12a) can be true even if the machine is immediately destroyed upon coming off the assembly line; (12b) can be true even if there happened never to be any mail from Antarctica; and (12c) can be true even if no emergency ever occurred. Such sentences show that what is required for the truth of generics is intensional: in these cases it amounts to the design of the machine, or the job-description of the mail-sorters, or the preparedness to act in certain ways in certain situations. Surely this shows the complete implausibility of trying to capture genericity with a quantifier, no matter how inherently vague or probabilistically-determined one tries to make the quantifier. No such extensional analysis *can* be correct. One of the earliest philosophers to make this point was Bosanquet (1888), who remarked that if a generic judgment is formulated by means of the quantifier *all*, then it is helpless in the face of the most trivial exception and for this reason he used as examples only statements not containing a quantifier. (For example, Man is an animal capable of social life, A society organized on a purely commercial basis treats the working classes as little better than slaves, The bacillus is a septic organism.) His idea seems to be rather congenial to our position: we notice that *nearly all* of the Xs are Y; we notice some exceptions; but we nonetheless doubt that the nearly all could be a mere coincidence. Therefore there must be some sort of connection, however circuitous and the exceptions can be accounted for by special conditions. And to express this sort of nomic dependency we use the generic forms of bare plural (Xs are Y), indefinite (An X is Y), and definite (The X is Y).[6]

[6] For a discussion of Bosanquet, see (Barth, 1974, pp. 364–367).

Another reason that there must be some operator other than an implicit quantifier is that there are ambiguities in certain characterizing sentences ambiguities between different generic readings. For one thing, there is a clear ambiguity between (what might be called)[7] a habitual reading and a dispositional reading. For example,

(13) John drinks beer,

can be understood as meaning that beer is John's favorite alcoholic beverage (habitual) or that John does not object to drinking beer (dispositional). Other examples of ambiguities are illustrated by sentences like[8]

(14) a. Typhoons arise in this part of the Pacific.
1. Typhoons in general have a common origin in this part of the Pacific.
2. There arise typhoons in this part of the Pacific.

b. A computer computes the daily weather forecast.
1. Computers in general have the task of computing the daily weather forecast.
2. The daily weather forecast is computed by a computer.

c. A cat runs across my lawn every day.
1. Cats in general run across my lawn every day.
2. Every day, a cat runs across my lawn.

The 1- and 2-readings of these sentences are both generic readings, and so we need some method by which to represent this type of ambiguity. There is also a difference between readings of generics that is imposed by emphasis or focus. Consider the difference amongst (italics indicates primary stress, which can be either focal or contrastive)

(15) a. Leopards usually attack monkeys *in trees*.

b. Leopards usually attack monkeys *in* trees.

c. Leopards usually attack *monkeys* in trees.

d. Leopards usually *attack* monkeys in trees.

(15a) says that when leopards attack monkeys somewhere, that somewhere is usually in trees; (15b) says that when leopards attack monkeys in, around, or near trees, it is usually *in* the tree; (15c) says that when leopards attack something in trees, it is usually monkeys; and (15d) says that when leopards are near monkeys in trees, they usually attack them. Another sort of ambiguity occurs in sentences such as

(16) Bullfighters are often injured.

[7] These different readings were called universal and existential by Lawler (1973), on the grounds that the first case is a quasi-universal reading and the latter case asserts an instance of John drinking beer. But we wish to move away from the extensional implications of these sentences in either of their senses and concentrate on their intensional force.

[8] The first of these sentences is from Milsark (1974), the second from Carlson (1979) where it is attributed to Barbara Partee, and the third is from Schubert and Pelletier (1987).

Clearly, often invokes some sort of implicit comparison class, and the ambiguity in (16) has also an effect on what this comparison class is. One can see (16) as saying that bullfighting is a dangerous sport and that, compared to participants in other sporting activities, bullfighters get injured in their sport more often. But one can also see (16) as saying that the sort of person who becomes a bullfighter is the sort of a person who tends generally to get injured compared to other people in their daily lives, bullfighters injure themselves more often in their daily lives. The difference here and in the other ambiguities also is what Schubert and Pelletier (1987) called the restricting cases against which the characterizing statement is made. In (16) the restricting cases can be seen as either situations in which athletes are practicing their trade, or else situations in everyone's ordinary life. In the various sentences of (15), the relevant restricting cases are those mentioned in the explanations following those examples.

1.6. Goals of an analysis of generic sentences

From a semantic point of view, there are three central goals of an analysis: (a) The analysis should give some account of the truth conditions of generic sentences, preferably one that accords with intuitive views of the truth conditions of sample generic sentences, showing how exceptions arise and in which way they can affect the truth of such statements; (b) The analysis should explain the facts about intensionality, especially the law-like nature of generic sentences and how they differ from quantified, extensional sentences; (c) The analysis must give an account of how generic sentences are used in reasoning, especially reasoning that involves the fact that such sentences admit exceptions.

Of course there are other goals one might reasonably ask of a philosophically and semantically adequate theory of generic sentences such as an account of how the analysis gives an explanation of propositional attitudes involving generic sentences, or an account of how generic sentences come to form part of the folk explanation of the universe, or an account of how generic sentences interact with other sentences in any of the fields mentioned above in Section 1.2 (ethics, legal theory, philosophy of science, etc). One might also ask for a clear statement of how the proposed semantic analysis fits in with a syntactic analysis of the same material: how can one syntactically distinguish the different partitions of the generic sentences so as to yield the various ambiguities we find in them? How can we deal with generic anaphora, since sometimes pronouns appear to refer to kinds and other times to instances of the kinds? And without doubt, there are many other reasonable expectations one might have of an adequate theory of generics. Our goals in this essay are modest, however; we concentrate only on the three goals identified in the previous paragraph, and even here we only give outlines. But we do think that the outlined answers will provide the final reasons desired, once the details are worked out.

1.7. A little notation

In general, such characterizing sentences have three parts, joined by an operator that we will simply call **GEN**: a matrix (a main clause) which makes the main assertion of the characterizing sentence, a restrictor clause which states the restricting cases relevant to the particular matrix, and a list of variables that are governed by **GEN**.[9] With such a notation we have a convenient way to represent the various readings of these characterizing sentences.

(14′) a. Typhoons arise in this part of the Pacific.

1. $\mathbf{GEN}[x](x \textbf{ are typhoons}; \exists y[y \textbf{ is this part of the Pacific} \mathbin{\&} x \textbf{ arise in } y])$.

2. $\mathbf{GEN}[x](x \textbf{ is this part of the Pacific}; \exists y[y \textbf{ are typhoons} \mathbin{\&} y \textbf{ arise in } x])$.

(13′) a. John drinks *beer.*

$\mathbf{GEN}[x, y, s](x = \textbf{John} \mathbin{\&} x \textbf{ drinks } y \textbf{ in } s;\ y \textbf{ is beer})$.

b. John *drinks* beer.

$\mathbf{GEN}[x, y, s](x = \textbf{John} \mathbin{\&} y \textbf{ is beer} \mathbin{\&} y \textbf{ in } s \mathbin{\&} x \textbf{ in } s;\ x \textbf{ drinks } y \textbf{ in } s)$.

c. John *drinks beer.*

$\mathbf{GEN}[x, s](x = \textbf{John} \mathbin{\&} x \textbf{ in } s;\ \exists y[\textbf{beer}(y) \mathbin{\&} x \textbf{ drinks } y \textbf{ in } s])$.

(Again, the italicizing indicates stressed pronunciation, or focus.) The universal interpretation in (13′a) can be rendered as: In appropriate situations (for which we use the variable s, and talk about an object being in a situation) in which John drinks something, this is normally beer. The existential interpretation in (13′b) says that in appropriate situations where there is some beer available, John normally drinks it. The habitual interpretation in (13′c) says that, in appropriate situations which contain John, he will drink beer.

The list of those variables which are governed by the **GEN** operator must be allowed to vary from reading to reading, for it is here that some of the ambiguities reside. Not only is it sometimes true that the objects which are governed by **GEN** are different in the different readings, as in (14′a), but also there can be different numbers of arguments, as in (13′c) compared to (13′a) or (13′b). And one can see the ambiguity reported in (8) as a matter of having different arguments to **GEN** objects vs. situations. It is this sort of flexibility which will allow us to capture the different readings of complex generics, where a NP may be within the scopes of several **GEN** operators or explicit quantifiers, any of which may control it, so that a disambiguating syntax is required. Consider

(17) If a person occasionally contributes to a charity when he receives an appeal, he will usually receive further appeals.

Here *a charity* may be controlled by *occasionally* or by *usually*, depending on whether *further appeals* is understood as *further appeals from charities* or as *further appeals from that charity*. Disambiguations similar to those made in (13) and (14a) can be made for the sentences in (15), (16) and (17), as well as the other sentences in (14), given clever enough identification of the situations and objects, and their interrelations.

[9] This tripartite structure comes from Lewis (1975).

1.8. Generics vs. explicit statements of regularities

It has already been remarked that generic statements appear to express law-like regularities, and not mere accidental correlations. This raises some interesting questions about the possible relationship between generics and explicit statements of regularities. What is the relationship between the characterizing statement (18a) and the rest of the sentences in (18)?

(18) a. Birds fly.[10]
b. Most birds fly.
c. Usually, birds fly.
d. Birds typically fly.
e. Birds generally fly.
f. Normally, birds fly.
g. In general, birds fly.

(18a) states a generalization that follows somehow from (assumed) natural laws; it conveys nomic force. The other statements can be used to express this meaning also, but in addition they can be used to assert, on a purely extensional level, that most, or many, etc., birds can fly. This latter is *not* a generic statement – it does not have nomic force although in many cases generic statements imply or implicate these kinds of extensional statements.[11] Nonetheless, many of the analyses of characterizing statements, especially those in the AI literature, are in fact analyses of these extensional sentences. Any attempt to analyze the generic statement in terms of "most" or "the usual" falls into this category; and attempts to use "the typical" or "the normal" fall into this category also, unless these phrases are themselves given a non-extensional meaning ... an option we will consider in the next section. They simply are not analyses of generics, but rather are analyses of these extensional implications and then only in those cases where the specific generic in fact *has* this sort of extensional implication. (As opposed to some of the other types of generic statements mentioned above.) It seems to us that many of the AI theorists have merely cited those specific characterizing statements that in fact happen to have these sorts of extensional consequences, then proceeded to provide an analysis of sorts for these consequences and concluded that they thereby have an analysis of the original generic statement. It is this last inference that we challenge, and it is for this reason that we earlier gave the long list of generics that do not have this sort of extensional consequence. Of course, these AI analyses can be interesting in their own right, because these sort of extensional consequences are important and challenging. But as mentioned below, it seems that at least some of the AI theorists arrive at their analyses due to a confusion between generic nomic necessity and extensional generality.

[10] There is also an ambiguity in 'fly' in these sentences, between the sense that birds have the ability to fly vs. that they *are* flying. This ambiguity is heightened in those sentences with the explicit adverbial quantifiers, (b), (g). We are concerned here with the "ability to fly" sense, but the other sense of these sentences is also a generic meaning and it too requires a suitable treatment. The analysis hinted at above, with situation-variables as part of the possible arguments to the list of controlled variables seems to afford a way to accommodate this ambiguity.

[11] But not always, as various of the aforementioned examples prove.

2. Semantic analyses of generics

The **GEN** operator, and its variable-control list, is just notation. The real question is: what is the semantic interpretation of **GEN**? We have already seen that it cannot be universal quantification, nor indeed any univocal quantifier, not even a vague probabilistic quantifier. The literature contains (at least) seven other suggestions. First, the **GEN** operator might be spelled out as relevant quantification. Second, the apparent quantification embodied in generic sentences might be understood as singular predication where the subject is an abstract object. Third, the notion of prototypical entities might be employed as an improvement on the abstract object analysis. Fourth, **GEN** can be seen as a statement of a stereotype. Fifth, **GEN** might be analyzed as a modal conditional, to be interpreted in a possible world semantics. Sixth, **GEN** might be analyzed as combining with a sentence to express a constraint in the theory of Situation Semantics. And seventh, **GEN** might be given a default reasoning analysis in one form or another.

Perhaps with suitable work each of these approaches can accommodate the view that characterizing statements have nomic force and are more than mere statements of what happens more often than not.[12] We content ourselves with making brief comments on these approaches. After these comments we turn in Section 3 to a somewhat more detailed analysis of default theories and the extent to which they can be used to account for the semantic phenomena of generics.

2.1. Relevant quantification

The generic operator might be spelled out as a universal quantification over relevant entities. When a statement is made about a natural group of objects, the hearer uses world knowledge to restrict the statement to just those members of the natural group to which it can be applied in a suitable way (Declerck, 1991). For example

(19) Whales give birth to live young.
$\forall x[[\mathbf{whale}(x)\ \&\ \mathbf{R}(x)] \rightarrow x\ \mathbf{gives\ birth\ to\ live\ young}]$

will be a predication over female, non-sterile whales, as only they could possibly give birth to live young in the first place. In the formulation above, this would be expressed by specifying the restriction variable **R** in a suitable way. One obvious problem with this approach is that this principle, as it stands, can easily justify all kinds of generic sentences ... since it is easy to find restrictions which would make *any* quantification come out true. For example, the analysis could make

(20) Whales are sick

be a true characterizing sentence, since we can take **R** to be the predicate *sick*, hence to restrict the quantification to sick whales. So this approach calls for a theory of suitable restrictions, and it is unclear how or whether this can be developed. Or, if it can be developed, perhaps it would just be circular in that one can find values of **R** only if one already knows what the true characterizing sentences are.

[12] Of course giving an account of what happens more often than not is itself no trivial matter. But interesting as this may be, it is not the same task as giving an account of the semantics of generic statements.

2.2. Abstract objects

It might be supposed that generic NPs refer to an indeterminate or arbitrary object, and that characterizing statements are in fact singular statements about this object. This view has the apparent advantage of enforcing a close link between our two types of genericity: Generic NPs and characterizing sentences. But it also seems not to cover the whole range of characterizing sentences such as habituals like *Mary smokes when she is nervous* that do not appear to mention any arbitrary objects.[13] Furthermore, it is not obvious how to enforce the distinction between an accidental universal generalization and a true generic with its nomic force in any theory that appeals to arbitrary objects. In assessing the plausibility of any such analysis, much depends on what such a theory counts as an arbitrary object.

Lewis (1972, p. 203) attributes to the dark ages of logic the story:

> The phrase ‘some pig’ names a strange thing we may call the *existentially generic pig* which has just those properties that some pig has. Since some pig is male, some pig (a different one) is female ... the existentially generic pig is simultaneously male, female ... The phrase ‘every pig’ names a different strange thing called the *universally generic pig* which has just those properties that every pig has. Since not every pig is pink, grey, or any other color, the universally generic pig is not of any color. (Yet neither is he colorless since not every – indeed not any pig – is colorless) ... He is, however a pig and an animal, and he grunts; for every pig is a pig and an animal and grunts. There are also the *negative universally generic pig* which has just those properties that no pig has (he is not a pig, but he is both a stone and a number), the *majority generic pig* which has just those properties that more than half of all pigs have, and many more ...

The history of philosophy has not been kind to arbitrary objects. For example, Lewis follows the above quotation with: “The story is preposterous since nothing, however recondite, can possibly have more or less than one of a set of incompatible and jointly exhaustive properties”.[14] Berkeley (1710) complained of (what he took to be) Locke’s (1690) version of this theory according to which an arbitrary person (or general idea of person) would be of no particular height, of no particular sex, of no particular age, and so forth. How, Berkeley asks, can such an item represent anything in experience such as a person? Further, Berkeley continues, the whole notion seems contradictory. If the arbitrary object or general idea has just those properties enjoyed by all people, then it will be *either male or female*. But as we have seen, it is not male, nor is it female; it is either male or female, but neither male nor female. Frege (1904), commenting on the use of arbitrary numbers in mathematics, asks whether there can be two *distinct* arbitrary numbers. If not, the theory of arbitrary numbers would be useless for mathematics (which often asks us to consider two arbitrary numbers); but if so, the theory is committed to distinct objects that are indiscernible, since the alleged two arbitrary numbers share all properties. Russell

[13] Perhaps one could invoke arbitrary Mary-when-nervous objects?

[14] At least, cant have them as its properties. But of course a *set* of properties could contain contradictory properties as members, and it is this use that Lewis recommends for natural language semantics. In the theory developed in his (1972), such sets are *characters*, and are what is in the extension of noun phrases; and a character is individual in his sense just in case it is a maximal compatible set of properties so that something could possess all and only the properties in the set. Otherwise he calls the character *generic*. But this sense of the term is not what is under discussion in the present article.

(1905) points to similar difficulties with Meinongian (1904) incomplete objects, which are arbitrary objects in the sense used here.

A modern account of (universally) arbitrary objects is provided by Fine (1985), who also brings out the similarities (and differences) of his approach to that of the ε-calculus (Hilbert and Bernays, 1934; see also Routley, 1967; and Smirnov, 1976). It is possible that a similar theory of *majority* arbitrary objects or maybe even *typically* arbitrary objects could be specified along the same lines, but we do not think it would be of use in explicating natural language generics, for it would share a number of features of the actual Fine theory which we find unsuitable. In his explicit consideration of all the preceding criticisms, Fine arrives at a theory in which the principle of bivalence fails and in which disjunctions are not semantically evaluated directly by a disjunction rule (and similarly, we presume, with other connectives). Further, there needs be a distinction between (what Fine calls) generic predicates and classical predicates. It is only the generic predicates which can give rise to arbitrary objects (and thus to generic NPs). All these restrictions and caveats seem necessary to Fine's theory, and they certainly do not take away from his goal of giving an account of how to employ arbitrary objects in formal logical proofs that invoke Existential Instantiation and Universal Generalization as rules of inference. But it seems to us that it still requires a lot of further research to employ them in an account of natural language generic statements. For example, although mathematics and logic can perhaps do without bivalence in their claims concerning arbitrary numbers or arbitrary instances or arbitrary triangles, the same does not seem true of natural language generic statements. Given that one is willing to assign truth values at all to generic statements, it seems clear that any particular generic sentence is either True or is False. 'Ravens are black', 'Sally smokes after dinner', 'Birds fly', and 'Hamsters eat dogs' are each of them true ... or if one should happen not to be true, then it is false. Certainly none of them are intermediate or neither true nor false.

2.3. *Prototypes*

It can be difficult to distinguish an arbitrary object approach from a prototypical object approach, despite Barth's (1974, p. 200) claim that "the two concepts should not be regarded as logical synonyms and ... not even as closely related notions". The real difference, as we see it, is that an arbitrary object theory proposes it to be an object that is of a different nature than "ordinary objects" (especially in the sorts of properties it has), whereas prototype theories presume that prototypes are of the same nature as ordinary objects. Recall that a universally arbitrary X-object has only those properties that all Xs share, that an existentially arbitrary X-object has properties that any X has, majority arbitrary X-object would have only those properties that are enjoyed by most Xs, and that a typically arbitrary X-object would have only those properties that a typical X has. But this means that universally arbitrary X-objects, majority arbitrary X-objects, and typically arbitrary X-objects could all be *missing* properties that are necessary for (ordinary) existence. For example, although the typical human has hair, there may be no typical hair color for humans ... and so the typically arbitrary person would have hair but no hair color. Yet it is necessary that any person with hair (who exists in the ordinary sense) must have some hair color or other. A prototype theory would make different claims about such cases. It presumes that prototypes are "ordinary objects", and so the prototypical human is an ordinary human. A prototypical human will have hair, as we said; thus, a prototypical human will have some specific hair color. But since there is no prototypical hair color, it follows that there will be distinct prototypical humans (each with hair) that exemplify the different prototypically possible hair colors. It may be that one of these people will have black hair, another brown hair, another blond hair, and a fourth red hair. From this it follows that the prototypical hair color is: either black or brown or blond or red.

In the quote cited just above, Barth was complaining about Jespersen (1924), who uses the term 'generic person'[15] when talking about NPs with the indefinite article (and indefinite pronouns). Jespersen says that such an article "is a weaker version of *any* ... where a dog is taken to represent a whole class". Barth takes this to mean that an arbitrary dog is to be regarded as paradigmatic, which in turn is the "misconceived" moment in the area.

It is not so obvious to us that Jespersen is conflating arbitrary and paradigmatic objects. To us, the difference between the two is that an A-arbitrary object of type X has all and *only* the properties that A of the Xs have, whereas a paradigmatic object of type X has all the prototypical properties that Xs have, but it will also have other properties. Our idea is that for such other properties there is a prototypical object enjoying each of the various alternative properties. Thus, a prototypical person has two arms, two legs, and hair on the head; and it has these because they are the prototypical properties of people. But being male, or being female, is not such a property; thus there will be one prototype that is female and one which is male. Having blonde, brunette, red or black hair are not prototypical properties, and this will give rise to four prototypes for each of the sexes, and so on.

The prototype approach we are interested in assumes that, from among the entities which are categorized as being an instance of a certain concept, we can choose those entities which are best representations of that concept. The fundamental idea behind these approaches is that all the entities which are the most typical representatives of a concept are called prototypes ... a concept popularized in cognitive psychology especially by Rosch (1978) and Osherson and Smith (1981). A characterizing sentence is seen as a universal quantification over the prototypical elements of a concept. For example, a sentence such as *A cat has a tail* can be paraphrased as *Every prototypical cat has a tail.* If we adopt an operator **TYP** which restricts the extension of a predicate to the entities which are prototypical for that predicate, then the prototype-analysis of generics becomes:

(21) A cat has a tail.
$\forall x[\mathbf{TYP}(\mathbf{cat})(x) \rightarrow \exists y[y \textbf{ is a tail } \& \ x \textbf{ has } y]]$.

When developing a unified treatment for all characterizing sentences, however, we must assume a very general prototypicality operator, because the contents of a Restrictor can

[15] Barth says, "Jespersen's use of 'person' here is quite general and covers dogs and other logical individuals as well as human beings".

vary widely; and we must also allow such a **TYP** to be applied to predicates of different adicities. One feature of **TYP**, which has been widely noted (Smith and Osherson, 1984), is that it is not a compositional operator. For example, **TYP(pet fish)** is not a function of **TYP(pet)** and **TYP(fish)**. Note also that this operator cannot be defined in terms of sets or other extensional entities, but must be specified as an operator whose arguments are intensional expressions. For, if **GEN** were defined in terms of extensional entities, then, in a world in which all birds except penguins became extinct, the notions of typical birds and typical penguins would coincide.

But even this intensional approach does not give a fine-grained enough representation, as shown by the following true characterizing sentences:

(22) a. A duck has colorful feathers.

b. A duck lays whitish eggs.

The problem here is that only male ducks have colorful feathers, and only female ones lay whitish eggs. As the set of male and female ducks are disjoint, the predicate **TYP(duck)** does not apply to any object at all. And this would have the untoward logical consequence that *any* characterizing sentence of the form *A duck Fs* would be true (because the antecedent of the universal quantifier would always be false). Clearly, the notion of prototypicality must be relativized to the property being expressed in order to save this approach. But it is far from clear how to accomplish such a relativization.

2.4. *Stereotypes*

A related approach is to see characterizing sentences as expressing **stereotypes** (Putnam, 1975; Scribner, 1977). Look at the following contrast:

(23) a. A lion has a mane.

b. A lion is male.

Why are we ready to accept (23a) but not (23b)? Note that arbitrary lions are more likely to be male than to have a mane, since only, but not all, male lions have a mane (e.g., not male cubs). Nonetheless, (23a) is definitely true and (23b) is definitely false. Why? One possible answer is to say that (23a) but not (23b) expresses a stereotype about lions in our culture: it is part of our stereotypical knowledge about the kind *Leo leo* that it has a mane. The general idea is to break down the meaning of a lexical expression into several components, including its extension and some stereotypical properties. These properties are considered to be the core facts (about the extension of the entities) which everyone speaking the language in question must know.[16] If **GEN** expresses stereotypical knowledge, then in order to understand **GEN** we need to investigate the formation of stereotypes. But there is little hope that we will find principles of *general* logical interest in such an investigation. For example, one reason why having a mane is part

[16] Note that here, as in the case of prototypes, the stereotype of a compound expression like *pet fish* is not a function of the stereotypes of its parts, *pet* and *fish*.

of the stereotype of *lion* is that the lion is the only cat that has a mane, making this a distinguishing property for lions. Our task would be to search out other stereotype formations in which distinguishing properties play a role. But it seems unlikely that there is anything of general *logical* interest (as opposed to anthropological interest) in stereotype formation.

Another potential difficulty for a stereotype analysis arises if one assumes cultural norms are the source of stereotyping properties (Scribner, 1977). For instance, suppose it is the norm in some culture to assume that snakes are slimy. Even there, the sentence *Snakes are slimy* is a false sentence, no matter how much it is believed to be true by most members of the culture, since snakes, themselves, are in fact not slimy. That is, generics are construed as making claims about the world, rather than what is considered a cultural norm. Furthermore, a stereotype analysis cannot be the entire story, for stereotypes are tied to single words or well-known and fixed concepts, but the Restrictor of **GEN** can be made up by novel concepts as well. A sentence such as *Mary smokes when she comes home* requires a generalization over situations in which Mary comes home. This sentence can be understood and believed to be true even if the hearer does not have a stereotype about situations in which Mary comes home. If the stereotype theory were correct, **GEN** would not have a uniform interpretation after all, or, equivalently, there would have to be numerous different generic operators. In either case there is no general theory of the semantics of **GEN**.

2.5. Modal conditional approaches

The modal conditional approach uses a possible-worlds semantics in the analysis of characterizing sentences, especially treating such sentences as modal conditionals (Delgrande, 1987; Asher and Morreau, 1991; Boutilier, 1994). This seems a more promising approach than any of the foregoing approaches. For one thing, it has often been remarked that characterizing generic sentences resemble conditional sentences. (For example, a characterizing sentence such as *A lion has a bushy tail* can be rephrased as: *If something is a lion, it has a bushy tail.* And restrictive *when*-clauses, which are often used in characterizing sentences – *When a cat is dropped, it usually lands on its feet* – are similar to conditional clauses in some respects.) So the extensive literature on the modal treatment of conditionals and counterfactuals is relevant here as well. Furthermore, the philosophical literature on dispositional predicates is related to the modal approach. Dispositional predicates, like *be soluble in water*, are generally reduced to law-like sentences (*If x is put into water, it will dissolve*), and law-like sentences are in turn analyzed as modalized sentences. Thus modal approaches perhaps promise to accommodate many of our desiderata for generics.

What sort of modality should be used? Since generics resemble counterfactuals, they seem to be best handled as a sort of variably strict conditional. Counterfactuals have the property noted by Lewis (1973) that two conditionals may have contradictory consequents even though one antecedent entails the other. For example,

(24) a. If you were to drop that plate, it would break.

b. If you were to drop that plate and it were made of rubber, it would not break.

c. If you were to drop that plate and it were made of rubber and its temperature were $-40°$ C, it would break.

Analyzed in terms of nonvariable modality and a material conditional, such statements are jointly unsatisfiable, whereas intuitively we take them to be all true. Generics statements exhibit the same sort of behavior:

(25) a. Birds fly.

b. Birds that are penguins don't fly.

c. Birds that are penguins and have rockets strapped to their backs fly.

So a natural suggestion (introduced first by Delgrande, 1987) is to analyze generics in terms of some sort of quantified conditional. Delgrande (1987) first proposed that generic statements might have the logical form of a quantified conditional, and his ideas have certainly been influential in the AI community.[17] Indeed the semantics from that we propose in Section 4 is a direct descendant of his.[18]

Yet there are some problems that need to be addressed. It is most tempting to analyze the relevant type of generic sentence in terms of "the most normal possible worlds". Sentences like (25a) above would be analyzed as: *In any of the most normal possible worlds, every bird flies.* Yet it is far from clear that there is any sense at all to the notion of being "a most normal possible world" in an unrestricted sense. (Isn't ours the most normal, by definition?) And even if we restrict the notion of normality to just being normal-with-respect-to-bird-flying-ness, it still seems implausible to say that it is more normal to have all birds fly than for some of them not to fly. Is it more normal for penguins to fly? Is it really more normal if there are no penguins? Or kiwis or emu or ostriches? Or fledglings? Or clipped-winged or broken-winged or dead birds? And as sentence (22a) reminds us, if there were a world in which all ducks had brightly colored feathers then there would be no female ducks and hence no brightly-colored males. (Or would it be more normal for the female ducks to be brightly-colored??) And does (22b) presume that it would be more normal for male ducks to lay eggs?? It seems pretty clear that one overall notion of "normalcy" just is not what is called for here, for these imagined scenarios simply are *not* more normal. And as the case of the interacting defaults in (22a,b) shows, even on the extensional level it is pretty difficult to imagine how such "conflicting defaults" can be accommodated when the so-called normal worlds must choose one or the other of the groups to be eliminated as abnormal, and yet if one is eliminated then the other becomes highly abnormal because its existence depended on the existence of the other group.

[17] Other people who have advocated at least implicitly that generics be treated as some sort of conditional are Pearl (1988), who analyzes conditionals in the manner of Adams (1975) as ε-probability statements.

[18] Like most others in the AI community, Delgrande's primary interest is the formalization of commonsense reasoning and not an analysis of generics.

2.6. *Situation semantics*

According to the situation-semantic approach (Barwise and Perry, 1983), it is convenient to use domains "smaller" than complete possible worlds as conversational background. A useful domain might be one in which we look only at single turtles or only at fertilized female ducks. The fifth approach does exactly this by considering *situations* instead of possible worlds, modeling characterizing sentences as *constraints* on situations. Constraints have already been used in Situation Semantics to model conditional sentences (Barwise, 1986), and since characterizing sentences are very similar to conditionals, this constraint technique would seem to be applicable to them as well.

Constraints are relations between types of situations. A constraint such as $\Sigma \Rightarrow \Sigma'$, where Σ and Σ' are situation types, says that Σ *involves* Σ', or that whenever Σ is realized, then Σ' is realized as well, i.e. whenever there is a situation σ of type Σ there is also a situation σ' of type Σ'. The situation types may contain parameters which can be *anchored* to specific entities, locations or types. (In such an approach, parameters are similar to variables, and anchors similar to variable assignments.) An important rule is that whenever f is an anchor for Σ (i.e. $\Sigma(f)$ differs from Σ only insofar as some parameters of Σ are anchored) and we have a constraint $\Sigma \Rightarrow \Sigma'$, then we also have a constraint $\Sigma(f) \Rightarrow \Sigma'(f)$. This rule captures the dependencies between parameters, and can be used to express the dependencies in sentences with adverbial quantifiers if we analyze the variables as parameters. It is important to note that it may be the case that a constraint holds only with respect to some background **B**. Such a conditional constraint is given as $(\Sigma \Rightarrow \Sigma') \mid \mathbf{B}$, read as: Σ involves Σ', given that **B**, where **B** is a situation type as well. The background **B** can be more specific than possible worlds could possibly be. For example, in sentence (22b) **B** can be restricted to situations containing female animals because the sentence tells us something about the mode of giving birth, and therefore only female animals should count. But our considering here only situations that contain no male ducks in no way commits us to denying that there are also mating situations in which there are both male and female ducks. As with the modal approach, the pertinent background is often left unspecified; and again we must assume some rule of accommodation. We discuss such a rule below, in presenting our theory in Section 4.

2.7. *Default reasoning approaches*

The sixth approach for handling the semantics of characterizing sentences is actually a set of related approaches being developed under the rubric non-monotonic reasoning both in logic and in AI. Although this literature does not contain an explicit discussion of generic sentences, the crucial examples that guide its development are always characterizing sentences, and so, if these theories are successful, we should expect from them an adequate semantics of characterizing sentences.

There are four main accounts of non-monotonic reasoning. Consider the following example:

(26) Birds are feathered.

a. If x *is a bird* is true, and if the fact x *is feathered* can be consistently assumed, then conclude that x *is feathered* is true.

b. If x is a bird and it is not known that x is not feathered, then x is feathered.

c. If x is a bird, and x is not abnormal for a bird with respect to being feathered, then x is feathered.

d. if x is a bird, and the probability of x's flying conditional upon it's being a bird is α, then the probability that x flies is α.[19]

The reasoning in (a) is an example of a "default rule" (Reiter, 1980). This is an inference rule that allows us to reach a conclusion C (here, x *is feathered*) from a sentence A (here, x *is a bird*), given that another sentence B (the so-called justification; here, x *is feathered*) is consistent with the facts assumed so far, i.e. we cannot infer its negation from these or other facts. The reasoning in (b) is an example of "autoepistemic reasoning" (McDermott and Doyle, 1980; Moore, 1984). The various approaches which fall into this category differ with respect to the epistemic modality they invoke, as well as in certain technical details. In general, this type of approach can be characterized as reasoning with the absence of positive knowledge, and is similar to default logic in this respect. Unlike default logic, however, modalized nonmonotonic reasoning allows the default rule to be made explicit in the object language by the modal operator it *is not known that.* The reasoning in (c) represents "circumscription" or "minimal entailment" (McCarthy, 1980, 1986). The central idea is to cover all exceptions by one predicate (one stating that the exceptional cases are abnormal) and to restrict the domain of this abnormality predicate to those entities which *must* be abnormal given the positive knowledge we have. This minimization of predicate domains is called circumscription. Clearly, such an abnormality predicate must be interpreted relative to some other property. For example, the emu, Bruce, is abnormal for a bird because he cannot fly, but not abnormal in the sense of having no feathers. The reasoning involved in (d) illustrates one of the ways in which probabilistic logic may be brought to bear on the issue of reasoning with uncertainty. Authors such as Bacchus, Grove, Halpern and Koller (1993) explicitly make the link between probabilistic logic and generics in the way glossed in (d); others such as Pearl (1988) have used the probabilistic logic only indirectly – by adapting the probabilistic semantics for conditionals given by Adams (1975). In the next section we examine these four approaches in more detail.

3. Default approaches and generics

In Section 2.7 above, we briefly introduced four different ways of using default or non-monotonic reasoning formalisms as possible ways to understand the semantics of generic, characterizing sentences – or at least, as ways to understand a *part* of the semantics of such sentences. In this section we explore these methods in more detail, with particular attention to how the inner workings of these formalisms interact with the phenomena

[19] α is a number that is typically taken to be high and in the case of Pearl (1988) $1 - \varepsilon$.

of genericity. We will concentrate on two aspects of this interaction – first, what sort of truth conditional semantics, if any, these formalisms imply for generics; and second, how these formalisms account for various intuitively acceptable forms of nonmonotonic reasoning in which characterizing sentences figure as premises or as conclusions.

3.1. A general characterization of nonmonotonicity and patterns of nonmonotonic reasoning

The term 'non-monotonic' indicates that these frameworks give a formal mechanism for retracting a previous conclusion when given new evidence. Intuitively speaking, an inference is non-monotonic if the set of premises Γ generates conclusion ϕ, but the premises $(\Gamma \cup \psi)$ do not generate conclusion ϕ. For example, learning that the supermarket is open and knowing that it carries wine and that I have enough money to buy wine, I can conclude that I can buy wine. However, I might (later) learn in addition that city laws prohibit the selling of alcoholic beverages between midnight and 6:00 in the morning, and as it is after midnight now, I must retract that conclusion. And this is done without retracting any of the original premises, and without saying that the original conclusion was invalidly inferred. All the approaches we have characterized as default approaches are centrally concerned with drawing inferences, and as such differ from, e.g., the modal conditional approach which is *not* a theory of inferences but rather a theory of the truth conditions of an object language connective.

The problem of how to treat exceptions is the basic question of the enterprise. A typical problem is that we may infer from the fact that Bruce is a bird the fact that he can fly, even though there are many types of birds which are exceptions to this rule, such as kiwis, penguins, ostriches, emu, dead birds, and birds with clipped wings. We might try to include all these exceptions in our rule[20] by saying that: if Bruce is a bird that is not a kiwi, not a penguin, not an ostrich, not an emu, is not dead, doesn't have clipped wings, etc., then it can fly. But it is not possible in general to give a complete list of exceptional properties; and, in a particular case, we may not know whether the entity has an exceptional property. Still, we often reason that, if Bruce is a bird, then he can fly, and retract this conclusion at a later time if we learn that Bruce is, in fact, an emu.

Research in nonmonotonic reasoning is relevant to the semantic analysis of generic statements or characterizing sentences. Over the past twenty years, those working on nonmonotonic reasoning have isolated a number of patterns deemed to be defeasibly acceptable, though deductively invalid, reasoning; and these patterns appear to crucially use generic statements as premises. These patterns constitute, we believe, an important aspect of the use of generics, hence one which a theory of meaning for generics should capture. On the other hand, it seems incontrovertible that the acceptability of these patterns of reasoning depends upon the meaning of the generics which figure in the premises. So an adequate semantics for generics should use the truth-conditional meaning of generics to define a notion of consequence that captures and justifies these patterns of reasoning, in some sense of justification. This notion of consequence, however, is

[20] Thereby adopting the variant (mentioned in footnote 5 above) of our first bad attitude.

not to be confused with the *ordinary* notion of validity that we have referred to at least implicitly in discussing the truth conditions of generics.

Here are some of the best known patterns, listed sometimes together with a similar but invalid argument form that shows how the original pattern can be defeated. By using our **Gen** notation (suppressing the mention of some occurrences of variables), we can represent the argument form *Defeasible modus ponens* (DMP) as follows:

DMP	**The "Defeat" of DM**
$\mathbf{Gen}[x](\mathrm{A};\mathrm{B})$	$\mathbf{Gen}[x](\mathrm{A};\mathrm{B})$
$\dfrac{\mathrm{A}[d/x]}{\mathrm{B}[d/x]}$	$\dfrac{\mathrm{A}[d/x],\ \neg\mathrm{B}[d/x]}{?}$

(This is the most famous example of the default inference patterns, and is enshrined in the Tweety example.)[21] We cannot expect commonsense entailment to behave like the two displayed patterns for completely arbitrary predicates A and B. For example, if A and B are chosen to be one and the same formula, then when adding the premise $\neg\mathrm{B}(t)$ to "defeat" the modus ponens we end up with the following premises: $\forall x(\mathrm{A} > \mathrm{A}),\ \mathrm{A}(t),\ \neg\mathrm{A}(t)$. These premises *contain* $\mathrm{A}(t)$, and for this uninteresting reason $\mathrm{A}(t)$ will continue to follow (possibly along with everything else, since the premises have even become inconsistent). Similar comments may be made in connection with the other patterns of defeasible reasoning.

Another pattern is the so-called *Nixon Diamond*, the name of which stems from another famous example from the nonmonotonic reasoning literature. Unlike the preceding example, this one exhibits a *failure* to draw conclusions from conflicting defaults.

Nixon is a Quaker.
Nixon is a Republican.
Quakers are pacifists.
Republicans are not pacifists.
?

And there is the *extended Nixon diamond* in which the Nixon diamond gets added premises and a different conclusion, which now *is* validly drawn:

Nixon is a Quaker.
Nixon is a Republican.
Quakers are doves.

[21] This most celebrated and widely accepted pattern of nonmonotonic reasoning does not work equally well with all generics, however. Consider the argument,

Frenchmen eat horsemeat.
Pierre is a Frenchman.
So, Pierre eats horsemeat.

This argument does not seem so acceptable; but arguably the difference stems from an interpretation of this generic as being of a piece with the generics that Lawler (1973), we feel inaccurately, called existential. Such sentences or such readings of generic sentences are not relevant to reasoning with exceptions, and they are not captured by this semantics. It remains an open question as to what determines such readings or the preference for such readings. Perhaps a finer lexical analysis of the predicates involved will yield an answer, but this is only a guess. Perhaps the analysis suggested in Section 1.7 holds the key.

Republicans are hawks.
Doves are not hawks.
Doves are politically active.
Hawks are politically active.
Nixon is politically active.

Two other important patterns are given below that researchers in AI have felt an essential part of defeasible reasoning with generics; they reflect the idea that information about subkinds should take precedence over information about the kinds which subsume them.

PENGUIN PRINCIPLE

Birds fly.	$\mathbf{Gen}[x](\mathrm{B};\mathrm{F})$
Penguins do not fly.	$\mathbf{Gen}[x](\mathrm{P};\neg\mathrm{F})$
Penguins are birds.	$\forall x(\mathrm{P}\rightarrow\mathrm{B})$
Tweety is a penguin.	$\mathrm{P}(\tau)$
Tweety does not fly.	$\neg\mathrm{F}(\tau)$

That penguins do not fly is a defeasible fact about penguins, a rule which admits exceptions. That penguins are birds, on the other hand, is a matter of taxonomic fact for which there are no exceptions.[22] It is interesting that swapping this taxonomic fact for the softer defeasible fact does not change our intuitions about the PENGUIN PRINCIPLE:

WEAK PENGUIN PRINCIPLE

Adults are employed.	$\mathbf{Gen}[x](\mathrm{A};\mathrm{E})$
College students are not employed.	$\mathbf{Gen}[x](\mathrm{S};\neg\mathrm{E})$
College students are adults.	$\mathbf{Gen}[x](\mathrm{S};\mathrm{A})$
Sam is a college student.	$\mathrm{S}(\sigma)$
Sam is not employed.	$\neg\mathrm{E}(\sigma)$

Our comments in Sections 1.2 and 1.3 suggest – and we will argue in more detail for this in Section 3.2 – that generics must be allowed as conclusions to arguments. Once we countenance generics as conclusions to defeasible arguments, then the following patterns naturally suggest themselves.[23]

DEFEASIBLE TRANSITIVITY

Defeasibly Valid	*Defeasibly Invalid*
Birds fly.	Slow eaters enjoy their food.
Sparrows are birds.	Those disgusted by their food are slow eaters.
Sparrows fly.	Those disgusted by their food do not enjoy it.
	Those disgusted by their food enjoy their food.

[22] Such taxonomic facts are understood as being meaning postulates, or modally necessary conditionals (of the S_5 variety), or we could simply stipulate that taxonomic facts are taken to be true in all worlds.

[23] Note the difference in the following patterns from what might be called "pointwise" versions of the same arguments. DEFEASIBLE TRANSITIVITY below yields a generic statement as a conclusion whereas the pointwise defeasible transitivity would add 'x is a sparrow' as a premise and yield 'x flies' as conclusion, and would do this for each x.

DEFEASIBLE STRENGTHENING OF THE ANTECEDENT

Defeasibly Valid	*Defeasibly Invalid*
Birds fly.	Birds fly.
White birds fly.	Dead birds do not fly.
	Dead birds fly.

Finally, we mention a type of inference that is invalid in nonmonotonic logics based on conditionals but which is valid in Reiter's default logic and in McCarthy's circumscription. It is widely believed that instances of this type should be valid in any nonmonotonic reasoning system.[24] Consider first the following pattern, in which we infer that Fido, even though we know he is abnormal in one respect (he only has three legs), is hairy, and so presumed normal in another respect.

(27)

$\mathbf{Gen}[x](\mathrm{D};\mathrm{H})$	dogs are normally hairy.
$\mathbf{Gen}[x](\mathrm{D};\mathrm{F})$	dogs normally have four legs.
$\mathrm{D}[f/x],\ \neg\mathrm{F}[f/x]$	Fido is a three legged dog.
$\mathrm{H}[f/x]$	Fido is hairy.

As mentioned, this inference is widely accepted in the AI literature; nevertheless it is suspect on empirical grounds. As Elio and Pelletier (1993, 1996) mention, it seems that ordinary people are rather unhappy about drawing these default inferences about objects that are known to violate other default rules. Here, Fido violates the default rule concerning number of dog-legs; and people in general are thereby less willing to allow Fido to obey the other default rule, about dog-hair.[25] And there are other inferences that seem undeniably acceptable but which nevertheless fail for several nonmonotonic logics that are based on conditionals. For example, nonmonotonic logics that are equivalent to the rational closure defined by Lehmann (1989) fail to predict the following inference, which seems intuitively acceptable (where B, D, F, K are logically independent in a sense to be made specific later):

(28)

$\mathbf{Gen}[x](\mathrm{B};\mathrm{F})$	Birds fly.
$\mathbf{Gen}[x](\mathrm{D};\mathrm{K})$	Dogs bark.
$\mathrm{B}[t/x],\ \mathrm{D}[c/x],\ \neg\mathrm{F}[t/x]$	Tweety is a non-flying bird; Charlie is a dog.
$\mathrm{K}[c/x]$	Charlie barks.

Note the difference between the inferences (27) and (28). In (27) we have some particular object f that violates some rule, and the inference nonetheless says to infer that the object f obeys some other default rule. As mentioned, one might say that an object failing to obey one default rule is itself a reason for being less confident in that object obeying

[24] See Lifschitz (1989, Problem 3).

[25] In the present example, people probably *would* agree to this specific conclusion, but that's because they know that number of legs doesn't affect hairiness. But as an abstract *pattern*, people are unhappy about the inference.

other rules. But in (28) we have two distinct objects, one of which violates a default rule. It is difficult to see why this should have any relevance to the other object's obeying or disobeying a different default rule. And the empirical results mentioned above do not cast doubt upon these sort of cases.

As explained in (Benferhat, Dubois and Prade, 1992; Benferhat, Cayrol, Dubois, Lang and Prade, 1993), most of the known rational systems of nonmonotonic logic fail to predict this inference, because they use a ranking of all the default rules according to which rules are grouped into clusters, and these clusters are then ranked. The clusters are defined by the logical relations between the antecedents. All those rules with antecedents that either strictly or by default imply a formula that is an antecedent in another rule X are in a cluster with a higher rank than the cluster containing X has. Some of the rules in a cluster may have exceptions; in that case all the rules in the cluster are blocked from any application. The preferred model takes all the clusters down the ranking to the first cluster that has an exception and treats them as universally quantified material conditionals. In effect, the blocked rules are simply treated as atoms. In the case of (28), both rules are put in the same cluster; since one of the rules has an exception, both are rendered inert.

The failure to validate the pattern in (28) seems disastrous for any theory of the role of generics in nonmonotonic reasoning. Nonmonotonic formalisms that do not verify this pattern appear to be completely unusable, since in most realistic applications of nonmonotonic reasoning we cannot presume that there are no exceptions to any generic premises that might be added to our argument.

3.2. *Evaluation of default approaches I: Default logic and update semantics*

The default logic approach (Reiter, 1980) in (26a) differs from the other nonmonotonic approaches we consider in that it uses a formula of the meta-language rather than of the representation language to state how characterizing sentences are to be understood. For example *Birds are feathered* would be for Reiter a "normal" default rule of the form:

$$\frac{\mathrm{bird}(x)\colon\ \mathrm{feathered}(x)}{\mathrm{feathered}(x)}.$$

Informally, this rule says that if x is a bird and it is consistent to assume that it is feathered, then you may conclude that it is feathered. A set of such rules help make up a *default theory*, which is a pair containing a set of facts and a set of default rules, $\langle \mathrm{W}, \mathrm{D}\rangle$. Reiter then defines the notion of an extension of a default theory to formalize commonsense reasoning. An extension of a default theory represents a particular sort of fixed point closure of the set of facts under the default rules (for details see the chapter in this handbook by Thomason). An apparently appropriate notion of nonmonotonic consequence for reasoning about generics in this theory is to say that a sentence ϕ follows from a default theory just in case ϕ is a classical logical consequence of every extension of the default theory. This definition of nonmonotonic consequence, like the alternatives proposed below, suffices to derive fly(tweety) from the default theory

$$\left\langle \{\text{bird}(\text{tweety})\},\ \left\{\frac{\text{bird}(x)\colon\ \text{fly}(x)}{\text{fly}(x)}\right\}\right\rangle$$

provided that we understand the default rule above as in effect the infinite set of all its instantiations when applying the notion of an extension. Thus, default logic is able to account for the pattern DMP.

Default Logic also predicts that the Nixon diamond is defeasibly invalid and the extended Nixon diamond is defeasibly valid. The Nixon diamond premises yield the following default theory

$$\Big\langle \big\{\text{republican}(\text{nixon}),\ \text{quaker}(\text{nixon}),\ \forall x(\text{dove}(x) \to \neg\text{hawk}(x))\big\},$$
$$\left\{\frac{\text{republican}(x)\colon\ \text{hawk}(x)}{\text{hawk}(x)},\ \frac{\text{quaker}(x)\colon\ \text{dove}(x)}{\text{dove}(x)}\right\}\Big\rangle.$$

This default theory yields two extensions, one in which Nixon is a hawk (and so not a dove) and one in which Nixon is a dove (and so not a hawk). So neither the conclusion that Nixon is a hawk nor the conclusion that Nixon is a dove follows from this default theory, given our definition of defeasible consequence for default theories.

It also predicts patterns of the form depicted in (28) and even (27) to be acceptable. It fails, however, to verify the Penguin Principle or the Weak Penguin Principle. In effect it treats the premise sets of those patterns on a par with the premise set in the Nixon diamond; it predicts that there are two extensions, one in which Tweety flies and one in which Tweety does not.[26] Influenced by Reiter's default logic, Veltman (1995) proposes a semantics in which all of the nonmonotonic inference patterns mentioned are validated defeasibly.

Default logic and its offspring do a good job at capturing those patterns of defeasible reasoning that we mentioned in the previous section. However, default logic does not provide us with an acceptable formalization of generic statements. Default rules are *rules*, and therefore are sound or unsound – rather than sentences, which are either true or false. If we analyze characterizing sentences using default rules, these sentences would not have truth values, and their meanings could not be specified by an ordinary semantic interpretation function. One consequence of being neither true nor false not being *in* the language is that characterizing sentences would therefore not talk about the world, instead they would talk about which inferences to draw. And this seems to us to be a strike against such an account.

Another possible strike against such an account is that since generics are seen as rules, not as statements, they cannot be the conclusions of arguments. But we think that generic conclusions may be among the things we should come to believe on the basis of other

[26] We note that extensions to the original default logic, such as that provided by Brewka (1991) can handle the Penguin and Weak Penguin Principles.

generic statements. In fact, one might even think it inconsistent to be willing to infer the particular conclusion *Sam flies* from the premises *Sparrows are birds, Birds fly* and *Sam is a sparrow* ... and to do this for any individual Sam, while refusing to infer the generic *Sparrows fly*. If it is true that the generic statement *Sparrows fly* summarizes facts about individual sparrows inherent flying abilities, and says something like *Individual sparrows normally fly, as a matter of nomic necessity*, how can we refuse to draw the generic as a conclusion?

Veltman (1995) in fact does claim that characterizing sentences are merely rules of inference, whose significance is not truth-conditional, but lies instead in what is called their "dynamic" meaning: the contribution which they make to how individuals update their beliefs on obtaining more information. Veltman puts his claim this way:

> The heart of the theory ... [of generics within update semantics] does not consist in a specification of truth conditions but of update conditions. According to this theory the slogan "You know the meaning of a sentence if you know the conditions under which it is true" should be replaced by this one: "You know the meaning of a sentence if you know the change it brings about in the information state of anyone who accepts the news conveyed by it".

The question of whether characterizing sentences have truth values, or whether, as Veltman supposes, they merely express rules of default inference, we take to be a very fundamental one. It is unlikely to be settled by amassing intuitive judgments of the type we brought up earlier in this Section and in Section 1, so it is of importance to look for other evidence which might bear on the issue. Here is one consideration.

If we were to accept the view, we would need to find a way to handle cases with nested **GEN** operators such as (5), or the following:

(29) A cat is healthy if it chases an object when it is moved in front of its eyes.

$$\mathbf{GEN}[x]\big(x \textbf{ is a cat } \& \ \mathbf{GEN}[y,s](y \textbf{ is an object } \& \ y \textbf{ is moved in front of } x\text{'s } \textbf{eyes in } s;\ x \textbf{ chases } y \textbf{ in } s);\ x \textbf{ is healthy}\big).$$

Default rules can't work here because in such an approach the embedded **GEN**-formula would have to be spelled out as a default rule which is a statement in the meta-language, and cannot be conjoined with a sentence such as x *is a cat* in the object language as required for the treatment of the outer **GEN**-phrase. This feature of generic statements that they can be embedded or nested within other characterizing statements is not restricted to just a few formulas. Already (5) and (29) illustrate a wide range, but also consider

(30) a. People who work late nights do not wake up early.

b. People who do not like to eat out, do not like to eat out.

Sentences like these are nested in that they attribute properties which involve genericity (as expressed here by the "habitual" predicate *wakes up early*) to kinds which are defined by means of characterizing properties (*people who work late nights*). (30b) is another example of a nested characterizing sentence, this time one which arguably is logically valid.

Theories of generic meaning must be sufficiently general and provide interpretations not only for simple characterizing sentences like *birds fly*, but also for composite sentences in which genericity mixes with more genericity (as in the cases of nested characterizing sentences just mentioned), and those cases where it mixes with counterfactuality, belief, knowledge, and other propositional attitudes ... as for instance in

(31) a. John believes that cassowaries fly.

b. John knows that Mary loves kissing him, and he would be unhappy if she were to like it less.

A theory of what generics mean ought at least to extend to a theory of what they mean in such contexts. And it is extremely difficult to see how any theory denying truth and falsity to characterizing sentences can account for these phenomena.

This is where theories that assign generics truth conditions (like McCarthy's below or the one that we mention later from Asher and Morreau) differ from those which treat characterizing sentences as mere rules of default inference. Where characterizing sentences are assigned truth values it is immediately clear how to interpret nested sentences. And where the truth values are assigned in a possible worlds framework, as they are below, it is clear how to embed characterizing sentences into propositional attitudes, modal contexts, counterfactual and other conditionals, and any other constructions which have been treated within this framework.

But if characterizing sentences are interpreted as mere rules of default inference, or functions from information states to information states, it is not even clear how they could be nested even up a few levels without making the formal account phenomenally complicated. We take this to be strong *theoretical* evidence of the importance of truth conditions in a theory of what characterizing sentences mean, and a strong reminder of the danger of restricting attention to too narrow a range of the phenomena which a theory of generic meaning can be expected to explain.

If one thinks that generic statements should have truth conditions as we do, then one must look at formalisms other than default logic or those accounts of generics which, like Veltman's, are inspired by its proposals.

3.3. *Evaluation of default approaches II: Autoepistemic logic*

Autoepistemic logic introduced by Moore (1984) attempts to formulate "nonmonotonic reasoning" by means of a particular epistemic logic. Autoepistemic logic gives a modal translation (though not a modal conditional translation) to generic statements. A statement like *birds fly* is translated as:

$$\forall x((\text{bird}(x)\ \&\ \text{M}(\text{fly}(x))) \to \text{fly}(x)),$$

where M is an epistemic modality that means *consistent with all my beliefs*. Moore's idea, following Stalnaker (1968), is that nonmonotonic reasoning, or at least one form of nonmonotonic reasoning, can be captured by formalizing the result of an ideal agent's reflecting on his beliefs. Moore calls this a *stable autoepistemic theory*. More particularly, if T is a set of formulas in a language with one modal operator B (for 'it is believed that') that represents the beliefs of an ideally rational agent, then T is a stable autoepistemic theory if

T is closed under classical logical consequence
whenever $\phi \in$ T, then B$\phi \in$ T
whenever $\phi \notin$ T, then $\neg$B$\phi \in$ T

To find out whether a particular formula ϕ nonmonotonically or "autoepistemically" follows an arbitrary set of formulas X, it suffices to see whether ϕ follows classically from every stable autoepistemic expansion of X, which is simply X closed under the clauses of what it is to be a stable autoepistemic theory given above. While Moore thought that his autoepistemic logic was distinct from default logic, Konolige (1988) showed that there was a sound and complete translation from default logic into autoepistemic logic; recently, Gottlob (1993) has shown that the expressive power of standard autoepistemic logic is strictly greater than that of default logic. Autoepistemic logic does as well as (though no better than) Reiter's default logic in capturing the defeasibly valid patterns of nonmonotonic reasoning, though the techniques used by Brewka (1991) to extend default logic to capture the Penguin Principles may also be applied to autoepistemic logic.

The question central to our concerns here is, does autoepistemic logic offer an adequate interpretation of generic statements? Clearly, we think it better than that of default logic, since the autoepistemic translation of *birds fly* is a sentence that can be true or false. But the autoepistemic interpretation of generic statements has some drawbacks which we think make it unacceptable. It fails to capture the appropriate nomic connection between antecedent and consequent of a generic statement for one thing. The autoepistemic interpretation of generic statements makes such sentences as *squares are round*, *tigers are vegetables* or *numbers are animals* be true. This is so because it is not epistemically possible (for any rational agent) that a square be round or that a tiger be a vegetable, and thus those statements are true in every one of these worlds. But such generic statements are clearly false. Further, this proposed analysis makes the truth conditions of generics depend on the epistemic states of particular agents – a hypothesis with which we are uncomfortable. We believe that *birds fly* would be true even if rational agents had never existed. Perhaps there are ways to eliminate the subjectivity from the autoepistemic proposal, but we will not attempt this here. We feel that there are enough problems with this view of generic statements not to proceed any further.

3.4. *Evaluation of default theories III: Circumscription*

McCarthy (1980) proposed the notion of circumscription not to analyze the meaning of generics but rather to formalize commonsense reasoning. Yet because the types of inference that he took to be examples of commonsense reasoning (e.g., "Defeasible Modus Ponens") typically contain generics as premises, we can take his ideas implicitly to define a semantics of generics. To a first order language, McCarthy adds a set of one-place predicates $\{Ab_i\}$. Each predicate Ab_i represents the property of being abnormal in some respect i. The translation into McCarthy's language of circumscription of a generic like *birds fly* then is: *any bird that is not abnormal-with-respect-to-birdiness flies* or symbolically

$$\forall x((\mathrm{Bird}(x)\ \&\ \neg \mathrm{Ab}_i(x)) \to \mathrm{Fly}(x)).$$

Even though this is an extensional treatment of generics, it offers significant improvements over the classical extensional translation, $\forall x(\mathrm{Bird}(x) \to \mathrm{Fly}(x))$. For instance, unlike the classical treatment, this treatment allows that *birds fly* may remain true even though not all birds fly, since there may be birds which are abnormal-with-respect-to-birdiness. McCarthy combines this treatment of generics with a particular approach to commonsense reasoning. To see what commonsensically follows from a set of premises containing generics, McCarthy proposes that we see what holds in those models of the premises in which the extensions of the Ab_i predicates are as small as possible. So for instance, if our set of premises are {birds fly, Tweety is a bird}, then we should look at all those models of the premises in which the extension of Ab_i above is as small as possible. Since nothing in these premises forces Ab_i to have a nonempty extension, we look at those models in which the extension is null; and in those models Tweety flies. Hence, says McCarthy, the inference from {birds fly, Tweety is a bird} to *Tweety flies* is an acceptable form of commonsense reasoning.[27] Circumscription captures all the inferences that Reiter's default logic does, but it too fails to capture the Penguin Principles. Lifschitz's development of prioritized circumscription (Lifschitz, 1985) remedies this difficulty, however.

Again although we think that research along these lines has helped elucidate commonsense reasoning and that commonsense reasoning is an important feature of the use of generics, we still feel that this proposal has some pretty clear defects as an account of the semantics of generics, and we will cite some of them in the last part of this section.

3.5. *Evaluation of default approaches IV: Probabilistic analyses*

Probabilistic analyses of generic statements and their use in nonmonotonic reasoning divide into two camps. The first approach, most recently advocated by Bacchus, Grove, Halpern and Koller (1993), uses probability statements directly to represent generics. Not only do they argue that nonmonotonic reasoning resembles probabilistic reasoning, but in particular they represent generic statements as conditional probability statements where the conditional probability approaches 1 (this has its roots in Adam's (1975) representation of a conditional sentence $\mathrm{A} > \mathrm{B}$ as the conditional probability statement $\mathrm{Pr}(\mathrm{B}/\mathrm{A}) > 1 - \varepsilon$). They then go on to define a nonmonotonic consequence relation that has many attractive properties. For us, however, the decision to represent natural language generics as conditional probability statements seems mistaken for the reasons we gave earlier in Section 1.3, and so we shall not consider these views further.

The other approach, exemplified by Pearl (1988), uses the probabilistic calculus to provide a semantics for a conditional, again in the manner of Adams (1975) as $\mathrm{Pr}(\mathrm{B}/\mathrm{A}) >$

[27] For more details, see Thomason's chapter in this Handbook.

$1 - \varepsilon$. And it is these conditionals that represent generic statements. This is an attractive proposal and close to the one that we shall propose in the next section. The nonmonotonic consequence relation defined by Pearl has many desirable properties. However, a difficulty with Adams's (and Pearl's) semantics for conditionals is that it is impossible to interpret conditionals with antecedents that contain conditionals within it. This casts doubt on the proposal as an adequate account of generic statements, since we have seen that nested generics (both in the antecedent and the consequent) are commonplace and perfectly acceptable. Further, Pearl's nonmonotonic consequence relation Z fails to verify either the patterns in (28) or in (27), and as we have said before, failing to validate at least the (28) pattern makes us seriously doubt that this approach can serve to capture the properties of generics in nonmonotonic reasoning.[28]

3.6. A general evaluation of default reasoning approaches

The main reason that non-monotonic logics appear to be useful for representing the meaning of characterizing statements is that they explicitly allow for exceptions to general rules, which reflects the fact that characterizing sentences typically allow for exceptions. Furthermore, there is a correspondence between the nomic quality of these characterizing sentences and the way generalizations are captured in non-monotonic logic: Part of the nomic force of characterizing sentences is that they make claims about an open (or open-ended) class of entities. For example, the sentence *A lion has a mane* does not make a claim about the closed class of all existing lions, but about every (realistically) possible lion. This excludes the possibility of simply listing the properties of the entities in question, or of formulating universal sentences and enumerating their exceptions. If it could be shown that default reasoning also embodied this feature, that would give a strong reason to favor it as an analysis.

Some nonmonotonic reasoning approaches also seems to fare well in certain cases which are problematic for the pure modal-conditional approach.[29] For example, the sentence *Ducks lay whitish eggs* does not require us to construct biologically inconsistent possible worlds in which there are no male ducks; it simply says that, when x is a duck and we have no information to the contrary (e.g., that it is a male duck), we can assume that x can lay whitish eggs.

The preceding few paragraphs state some of the attractive features of (some of) the nonmonotonic reasoning approaches in their treatments of natural language generic statements. But maybe even this much credit to the nonmonotonic logic approach is too much. For instance, isn't it *always* the case that we *do* have some information to the contrary, and yet we continue to assert the characterizing statement? Consider these characterizing statements:

(32) a. Marvin usually beats Sandy at golf.

b. Italians are good skiers.

---

[28] In (Goldszmidt and Pearl, 1990) an approach is developed that overcomes this difficulty.

[29] Of the approaches considered in the previous section, Pearl's (1990) nonmonotonic logic, which is based on understanding generics as conditionals, would have trouble with this sort of inference.

c. Bees lay millions of eggs at a time.

Clearly it would be incorrect to analyze them as

(32′) a. Given no information to the contrary, conclude that Marvin is beating Sandy at golf.

b. If x is an Italian and there is no information that would say s/he wasn't a good skier, then assume s/he is a good skier.

c. If x is a bee and there is no reason to assume it is not laying millions of eggs, then conclude that it is laying millions of eggs.

These are incorrect because there is *always* information or good reasons to the contrary, for such characterizing statements. We *know* that Marvin and Sandy are *not* usually playing golf, and so by (32′a) we would never conclude that Marvin is beating Sandy. We have an *immense* amount of information that *very, very few* Italians are world-class skiers, so we would never draw the conclusion indicated in (32′b). And we know that hardly any bees ever lay eggs, and further that when they do it is not over a very long period of time, and so we would never conclude the consequent of (32′c). But therefore it follows that the analyses in (32′) are not correct as accounts of the generics of (32).

Also, much of the default reasoning literature and many of the proposed mechanisms are stated at the extensional level. The methods yield minimal models: models where there are the fewest objects that are abnormal-with-respect-to-the-specified-predicate. Yet this is not really what is desired as an analysis of generics, for it does not honor the nomic force of generics. And even in the extensional case it yields incorrect results, for example the minimal model approach would claim that 'Birds fly, Sam is a bird ∴ Sam flies' is valid yet is defeated by the addition of 'There is a non-flying bird' as a premise. Also, little if any work has been done on the interacting defaults of the sort expressed by the sentences in (22). For these reasons, it is difficult to take seriously the work thus far completed in default reasoning as an analysis of the phenomenon of genericity. Much of the problem stems from the differing desires of AI researchers and formal semanticists. The former researchers are trying to show how to reason when we have statements that admit of exceptions. Yet this is not all there is to the semantics of generics. Indeed, it seems to us that this is less significant than issues of nomicity and of trying to specify what the relevant background against which the characterizing statement is being made, and using this information to state the truth-conditions for generic statements. And it is also less significant than the fact that differing generic statements call for radically different numbers or proportions of exceptions. Until such issues can be adequately addressed, it seems improper to investigate merely the notion of reasoning which admits exceptions and expect this to yield an answer to the semantics of generics.

For these reasons we put forward an approach to non-monotonic reasoning and the semantics of generic statements that combines insights from various of the suggestions given above, in the hopes that what was appropriate from each of those methodologies will survive. The approach, laid out in (Asher and Morreau, 1991, 1995) and (Morreau, 1992), uses a conditional logic (with a possible worlds semantics) to characterize certain aspects of generics. In particular, it allows one to specify semantic truth conditions of characterizing sentences and to do so in such a way as to allow them to enjoy their nomic force at least to the extent that nomic force and intensionality can be captured in modal logics. It combines this semantics with a non-monotonic inference relation based on situation semantics to capture default inferences involving characterizing sentences. This use of situation semantics to capture the default inferences avoids the use of minimal models (and the like) that is characteristic of circumscription and related extensional methods. In short, the theory outlined in the next section gives truth conditions to generic statements and allows these statements to have exceptions, but it separates this truth-conditional role of the semantics from the inferential role of drawing conclusions non-monotonically.

4. A formal theory of generics

Section 2 of this chapter suggests several sources for a formalization of the meaning of generics, such as: the theory of prototypes, situation theory, modal and conditional logics, and the theory of nonmonotonic reasoning. In this section, we will rehearse the requirements concerning a formal semantics for generics and then mention one theory that meets these requirements.

4.1. Requirements for the monotonic semantics for generics

We have argued that generics have truth conditions and that a semantic theory of generics should capture them. We also argued that the semantics for generics should be intensional, since extensional and probabilistic accounts succumb to difficulties that the modal approaches avoid. Our strategy is to interpret **GEN** as a conditional operator, although we need to be somewhat on the intuitive side as to what the variable-list is that **GEN** controls. Thus, from a sentence like (33a) we arrive at the logical form (33b).

(33) a. Dogs bark,

b. $\mathbf{GEN}[x](\mathbf{Dog}(x);\ \mathbf{Bark}(x))$,

or if we wish to analyze further the dispositional predicate *bark*: as itself involving a generic predication (which accords with intuitions)

c. $\mathbf{GEN}[x,s]\big(\mathbf{Dog}(x)\ \&\ \mathbf{IN}(x,s);\ \exists e(e < s\ \&\ \mathbf{Barking}(x,e))\big)$.

This logical form has the following intuitive meaning: For any object x if x is a dog and x is in an appropriate circumstance then this circumstance will normally contain an event of barking by x. So understood, **GEN** is definable by means of a universal quantifier binding some variables and a propositional conditional operator, call it $>$, linking the restrictor of **GEN** with its nuclear scope. So (33b) would, for instance, emerge as

d. $\forall x(\mathbf{Dog}(x) > \mathbf{Bark}(x))$.

Now of course, it is notoriously difficult to specify the restrictive argument for event quantification in the analysis of dispositional predicates; still, we would argue that this is an appropriate representation of the original English ... remaining vague just where the English is vague.

We now turn to some remarks concerning desiderata for the semantics of the conditional operator, >. It seems reasonable to exploit the semantics for conditional sentences developed by Stalnaker and Thomason (1967), Stalnaker (1968) and Lewis (1973), which assigns truth values to conditional sentences relative to possible worlds. Roughly then, (33c) is true at a world w just in case for every object d the set of worlds that are normal with respect both to w and to the proposition of d's being a dog is a subset of the set of worlds in which d typically is barking in the appropriate circumstances. More precisely, one can make use of a selection function * that takes worlds and propositions (sets of worlds) as arguments and returns a proposition.

$$M, w, a \models \mathrm{A} > \mathrm{B} \quad \text{iff} \quad {}^*(w, [\mathrm{A}]_{M,a}) \subseteq [\mathrm{B}]_{M,a},$$
$$\text{where } [\mathrm{A}]_{M,\alpha} = \{w \in W_M\colon\ M, w, \alpha \models \mathrm{A}\}.$$

We can think of ${}^*(w, p)$ as all the normal p worlds according to w, that is, all the worlds in which p, along with all the typical consequences of p in w, hold.

What sort of additional logical inferences do we want > to verify? One intuitive principle that a semantics for generics should validate is some sort of "idempotence" principle such as that exemplified by the following sentence:

(34) a. Dogs are dogs.

This generic, though uninteresting, strikes us as valid. And on our tentative understanding of **Gen**, we cannot imagine how such a sentence could come out as false. For anything, if it is a dog, then normally it is a dog. We account for this by FACTICITY, below.

Another pattern of logical entailment, called WEAKENING OF THE CONSEQUENT, is common to all sorts of conditionals with a modal semantics and seems obviously right. We think that it is part of the truth conditions that generics sentences like (34b) always come out true.

b. People who don't like to eat out don't like to eat out.

Our method of accounting for such sentences is by means of the principle: If B logically entails C, then the characterizing sentence *A's are B*, entails the characterizing sentence *A's are C*.

$$\text{WEAKENING:}^{30} \quad \text{if } \models \forall x(\mathrm{B} \to \mathrm{C}) \text{ then } \models \forall x(\mathrm{A} > \mathrm{B}) \to \forall x(\mathrm{A} > \mathrm{C})$$

[30] The $\to$ is the monotonic material conditional of classical logic, while > is the conditional being defined. The notation is that in use in (Asher and Morreau, 1991, 1995) and (Morreau, 1992).

which can be used jointly with facticity to ensure that sentences like (34b) come out true.[31]

Finally, we take another often-adopted principle of conditional logic to be intuitively acceptable. If A's are normally B's and if C's are normally B's, then for any object, if it's either an A or a C it's normally a B. Using the bare plurals formulation, we accept sentences like the following as valid:

(35) If lions are fierce and if tigers are fierce, then (things that are) lions or tigers are fierce.

We might call this principle **OR** to reflect the fact that conjunctions of generic sentences can lead to generics whose representation in terms of **Gen** or $>$ makes the restrictor a disjunctive formula, just as with the antecedent of a classical conditional.

On the other hand, the semantics for the conditional should not support modus ponens, modus tollens, nor principles like $(\alpha \,\&\, \beta) \to (\alpha > \beta)$, which are prevalent in many conditional logics. To do otherwise would violate one of the starting points of the study of generics, the existence of exceptions. None of modus ponens, modus tollens, and the aforementioned principle honor exceptions to true conditional statements.

If we adopt the analysis of **Gen** using the conditional $>$, then we can validate or invalidate the sort of patterns of inference we have been discussing, by placing or not placing certain constraints on the selection function $*$. For instance, the constraint on $*$ known as centering below supports modus ponens for $>$, and so should *not* be adopted in a semantics for generics.

CENTERING: If $w \in p$, then $w \in {}^*(w,p)$.

Other constraints on $*$, like FACTICITY, verify the desirable idempotence inference above in (34a).

FACTICITY: ${}^*(w,p) \subseteq p$.

FACTICITY says that in those worlds where the propositions normally associated with p are true, p is also true. In other words, p is itself one of the things which normally holds when p holds. The following constraint would validate the OR principle

OR: ${}^*(w, p \cup q) \subseteq {}^*(w,p) \cup {}^*(w,q)$.

Delgrande (1987) and others like Boutilier (1992, 1994) who have used modal conditionals both to define a nonmonotonic consequence relation and (at least implicitly) to represent generic statements as well, have often opted for more constraints on the modal frames and the selection function. Were we to follow them, $>$ would reflect properties attributed by Gabbay (1985) and Kraus, Lehmann and Magidor (1990) to a nonmonotonic consequence relation. We would thus strengthen the basic logic of $>$. For instance,

[31] Starting with the instance of FACTICITY, $\forall x((\mathrm{P}x \,\&\, \neg \mathrm{E}x) > (\mathrm{P}x \,\&\, \neg \mathrm{E}x))$, we apply WEAKENING to get $\forall x((\mathrm{P}x \,\&\, \neg \mathrm{E}x) > \neg \mathrm{E}x)$, which is the relevant translation of (34b).

we can force $>$ to obey all the constraints of a "rational conditional" (Nute, 1980) by adding certain constraints on $*$. This rational conditional obeys the following principles:

CUT:	$((A > B) \mathbin{\&} ((A \mathbin{\&} B) > C)) \to (A > C)$.
CAUTIOUS MONOTONICITY:	$((A > C) \mathbin{\&} (A > B)) \to ((A \mathbin{\&} B) > C)$.
RATIONAL MONOTONICITY:	$((A > C) \mathbin{\&} \neg(A > \neg B)) \to ((A \mathbin{\&} B) > C)$.

If $>$ were a rational conditional, then generics would obey the following argument patterns:

(36) [CUT] Dogs bark. Dogs that bark annoy Fred. So, dogs annoy Fred.
[CM] Dogs annoy Fred and dogs bark. So, dogs that annoy Fred bark.
[RM] Dogs bark. It is not true that dogs are not hairless. So, hairless dogs bark.

We are not sure that these forms are intuitively acceptable, although Rational Monotonicity seems very plausible.[32] And so we will not impose these constraints, nor their associated inference patterns on $>$, as a part of our semantics of generics.

A final requirement on the truth conditional semantics for generics is that nestedness should also be handled satisfactorily. Earlier, we argued that sentences in whose translation the **Gen** quantifier had scope over formulas that themselves contained the **Gen** operator were not only grammatical but had determinate truth conditions. While we are unsure what principles involving nested generics the logic of **Gen** or $>$ should capture, we reject any semantics for generics in which the set of "normal" worlds is fixed in advance independently of the antecedent. In some nonmonotonic logics that define conditionals in terms of modal operators, this "independent fixing" occurs. For instance, in the logics of Boutilier (1992) based on variants of S4 or S4.3, the following principle is valid:

$$(A > B) \to ((A > \neg B) > \bot).$$

This is unacceptable in a theory of generics for it would entail the validity of:

(37) If dogs bark, then typically dogs that don't bark are frogs.

The ability to handle nested generics is also important when it comes to problems of quantification. Lehmann, for instance, notes that the following three sentences are entirely consistent:

(38) a. Elephants like zookeepers.
b. Elephants don't like Clyde.
c. Clyde is a zookeeper.

However, formalizing the first generic in a "flat" fashion using a conditional with a conjunction in the antecedent, such as

(38′) a. $\forall x \forall y((\text{Elephant}(x) \mathbin{\&} \text{Zookeeper}(y)) > x \text{ likes } y))$,

[32] For all "generically consistent" A, i.e. where $\neg(A > \bot)$, RM as formulated above entails CM.

would entail in many conditional logics (e.g., Boutilier's) or on an extension of Veltman's (1995) approach (given an ordinary formalization of the second generic and the third premise) that in the normal worlds where e is an elephant, e does not like Clyde only if Clyde is not a zookeeper. We find this counterintuitive, but when we add some means of reasoning nonmonotonically with these premises whereby the conditionals with the logically more specific premises cancel out those with more general antecedents in case of conflict (see Section 4.2 below), we can conclude from the information given in the premises, that indeed elephant e should like Clyde. But this is intuitively what the second premise says is not true!!

Notice that if we follow rules of compositional translation from natural language familiar to linguists, we get a quite different translation for the first generic, namely:

(38″) a. $\forall x(\text{Elephant}(x) > \forall y(\text{Zookeeper}(y) > x \text{ likes } y))$.

If we adopt the constraint of weakening as we believe to be reasonable, then we get (using universal instantiation)

(38‴) a. $\forall x(\text{Elephant}(x) > (\text{Zookeeper}(c) > x \text{ likes } c))$.

Let e be an elephant and let w′ be any normal e-elephant world with respect to the world in which (38‴a) is true. Then this generic implies that, for any normal Clyde-zookeeper world with respect to w′, Clyde is liked by e in that world. Consider the translation of the second generic:

(38′) b. $\forall y(\text{Elephant}(y) > \neg y \text{ likes } c)$.

There is no incompatibility with these translations. It is true that even Clyde, in those worlds where he is a normal zookeeper, is liked by the elephants. But that is perfectly consistent with the disposition of the elephants normally not to like Clyde, as he is (in the actual world). Clyde, being a nasty zookeeper in the actual world, is normally not liked by the elephants. There is in effect a de re/de dicto distinction familiar from other intensional constructions; Clyde may, of course, have different properties in the worlds where he is a normal zookeeper from those that he has in the actual world or from those that he has in a normal e-elephant world for any elephant e. If this is right, then we see a reason for insisting that the inference, exportation, is not valid:

$$\text{EXPORTATION:}\quad \forall x(\mathrm{A} > \forall y(\mathrm{B} > \mathrm{C})) \rightarrow \forall x\forall y((\mathrm{A}\ \&\ \mathrm{B}) > \mathrm{C})).$$

This yields yet another reason to abjure extensional translations of generics, for in such translations (like that of circumscription) exportation is perforce valid. It also gives us a reason to not take the step often made in modal analyses of normalcy where one insists that the normal e-elephant worlds with respect to w are those worlds that *resemble* the actual world as much as is consistent with e's being a normal elephant. We note that such a constraint, in a definition of the generic conditional as in (Boutilier, 1994), would yield the unintuitive conclusion that the closest worlds compatible with e's being a normal elephant are those in which Clyde is not a zookeeper.

4.2. Interim evaluation

The system CE (commonsense entailment) of Asher and Morreau (1991, 1995) meets most of the desiderata above. It is a first order language augmented with a conditional connective $>$, which is interpreted in modal frames using a selection function that obeys only FACTICITY and OR. (Hence the axiomatization of $>$ contains only $(A > A)$ and $((A > C) \& (B > C)) \to ((A \vee B) > C))$.) It has two consequence relations, $\models$ and $\approx$, corresponding to classical and default consequence. The axiomatization is shown to be (argument) sound and complete in (Morreau, 1992).

We now turn to a discussion of some of the merits of this truth conditional proposal. First of all, it is nonextensional and so avoids the objections we leveled at extensional approaches. It makes generics modal but contingent, which is what we would like. That is, the truth value of a generic statement depends on worlds other than the actual one (generics are modal statements), but it may also vary from world to world (they are contingent); our generic truths may not be the generic truths of some other world. This analysis of truth conditions meets the desiderata mentioned in Section 4.1. It verifies FACTICITY, OR and WEAKENING, but it fails to verify modus ponens for $>$, and the unintuitive claims about nested conditionals that follow from Boutilier's definition of the generic conditional. This is because in our analysis of the truth conditions of generics, we do not (as Boutilier does) postulate one ordering of normal worlds relative to a given world w that must verify all the generics that are true in w. Instead, normality is captured by a set of worlds that is determined by the world of evaluation and the proposition expressed by the antecedent of the $>$ statement.

Recall, however, that in Sections 2.5 and 4.1 we made several criticisms of other modal approaches; we saw problems both in the very conception of normal worlds and in the way such proposals deal with problems of conflicting, true characterizing sentences. We noted above that some modal approaches, such as Delgrande's (1987), postulate a set of absolutely normal worlds. This leads to difficulties with (true) generic sentences like

(39) Turtles live to be a hundred years or more old,

because such a theory proposes that the most normal worlds are those where every single turtle lives to a grand old age, instead of dying while very young as the vast majority do in the actual world. Such worlds would be biologically very strange, and strike us as most implausible candidates for absolutely normal worlds. Such worlds might even violate general biological laws such as those about what would happen to turtles as a species if all of them were to be long-lived, and we think that all worlds which are normal with respect to a world w must respect the biological laws that true in w.

So let us reconsider the problems with the conception of normal worlds in the light of our new proposal for the semantics of generics. We first note that the semantics of $>$ we have sketched above does *not* presuppose *any* absolute normality order on possible worlds. In particular, we explicitly reject the idea that $^*(w, p)$ is to be identified with those most normal of all possible worlds where p holds. Let us see how this semantics evaluates sentence (39) in a world w.

(39′) $\forall x(\mathrm{T}x > \mathrm{H}x)$.

For each individual d, $*$ determines, together with w and the predicate 'T', a set of normal "d-turtle with respect to w" worlds, $*(w, [\mathrm{T}](d))$. In every member of this set of normal worlds, d must have the property of living to be a hundred years or more old, but other objects that are turtles in these worlds may die, as expected, when very young.

Another problem with modal approaches is brought out by such true characterizing sentences as *ducks lay eggs*. If we accept the constraint that only female ducks can lay eggs, then apparently we must conclude, using Weakening of the Consequent, the unintuitive generic *ducks are female*. Recall also the problem of interacting defaults. The potential problem for the semantics of generics that we raised was where two intuitively acceptable generics with the same antecedent but incompatible consequents held:

(40) a. Ducks have brightly colored feathers.

b. Ducks lay eggs.

Since only males have brightly colored feathers and only females lay eggs, we see that no normal duck can have both properties. Nevertheless, it is easy to verify that in the conditional semantics that we have assigned to generic sentences, the following property holds:

$$(\forall x(\mathrm{A} > \mathrm{B})\ \&\ \forall x(\mathrm{A} > \mathrm{C})) \rightarrow \forall x(\mathrm{A} > (\mathrm{B}\ \&\ \mathrm{C})).$$

Using this principle and the intuitive translations of the two generic sentences in (40), and the knowledge that only male ducks have bright feathers and only female ducks lay eggs, and that no duck is both a male and a female, would yield the following unintuitive result:

$$\forall x(\mathrm{duck}(x) > \perp).$$

This observation poses a prima facie difficulty for the proposed modal semantics for generics. But at first glance it would appear to be part of a much more general problem of determining quantificational structures in natural language. In Section 1 above we assumed that a natural language quantificational structure has three components: the quantificational relation itself, a restrictor, and a nuclear scope; and in general, determining the restrictor of a natural language quantifier is a nontrivial matter. The problem is most immediately apparent with universal quantification where the domain of quantification is most often contextually determined. For instance, imagine a group of school students on holiday with a chaperone who calls some apprehensive parents after an outing and says,

(41) Everyone had a good time.

The chaperone does not thereby assert that every human being in the universe had a good time, which is what would follow on the usual translation for such a sentence. The inadequacy of translation of (41) is that it fails to consider the contextually sensitive nature of the domain of quantification. One way to rectify this situation is to imagine plausible content, following Lewis's (1983) ideas on accommodation and von Fintel's

(1994) ideas about quantification. On such a view, in addition to the content derived from the subject noun phrase *everyone*, we add to the restrictor a formula of the form $x \in \mathrm{A}$, where x is the variable introduced by *everyone* and the value for A is contextually specified. In the particular context of (41) used above the value for A would be the set of children and the logical form for (41) would be:

(42) $\forall x((\text{Person}(x)\ \&\ x \in \mathrm{A}) \rightarrow$ had a good time$(x))$.

Is this strategy available to help solve the difficulty with (40a,b)? We believe that this strategy is *not* one that can in fact be used for the generic examples. First of all, (40a,b) may be perfectly felicitously be uttered in a context where there are no obvious values for a contextually specified variable. And both (40a,b) would be true in such a context. In contrast (41) in such a context (and thus without a determinate value for the set variable A in the logical form above) would truly be a universal quantification over all people. Second, attempting to create an appropriate contextual value for the parameter that is present in the logical form of (41) must not make false generic sentences magically become true ones. Consider for instance a case similar to the one in (41) where a lab technician observes fruitflies that have been genetically altered in an experiment. He observes that the altered fruitflies have five eyes. Eager to tell someone of the observation, he calls up the chief scientist. We suppose that were the researcher to recount what he had observed and then summarized with

(43) Fruitflies have five eyes,

he would have said something false. But if that is right, then the strategy of relativizing the quantification to some contextually specified set is irrelevant to accommodating generics. The problem of determining the restrictor for generics seems to be different from that for ordinary quantifiers.

We think that the proposed modal semantics in fact points the way to a proposed solution for this problem. The problem with *ducks lay eggs* is not our general understanding of the logical form of this sentence but rather our strategy for determining the antecedent of the $>$ that functions as a restrictor for the universal quantifier. So there is a sort of accommodation for generics. Our semantics, however, tells us that generic statements are modal statements, and so accommodation or further restriction of the quantification should not be done with respect to some extensionally defined parameter like the value of a set as with ordinary quantifiers, but rather with respect to some *intension*. Our view is that the appropriate restriction for a problematic generic like that in (40a) or (40b) is some subtype of the type given by the common noun that is formed into the bare plural. That is, what people really understand by (40a,b), at least when they are presented with the pair and the facts is that the universal quantifiers in each are relativized to a particular subtype. *Ducks lay eggs* is implicitly understood as *female ducks lay eggs*. On the other hand there is no natural subtype of fruitflies that would pick out just those in the experiment.

One way to test this hypothesis is to look at how speakers treat such problematic generics as (40a,b) when they exploit them in inferences. We have argued that an inference like Defeasible Modus Ponens is defeasibly valid and so should be part of a theory of generics. But many speakers express reluctance to draw even the defeasible conclusion that if Allie is a duck and ducks lay eggs, then Allie lays eggs. It would depend, they say, upon whether Allie is a male or a female duck. Similarly for (40.b) and for other cases of conflicting generics. If Defeasible Modus Ponens requires such an additional premise, then this strongly suggests that the logical form of (40.a) is not what the string itself would suggest but rather:

(44) $\forall x((\text{duck}(x)\ \&\ \text{female}(x)) > \text{lays eggs}(x))$.

We would hypothesize that such relativization of the quantifier by a type occurs only in cases where we have conflicting generic statements.

4.3. Nonmonotonic reasoning and epistemic semantics for generics

In addition to having truth conditions, generics also play an important role in plausible or nonmonotonic inference. Because we believe this role has something to do with the meaning of generics, we need to show how our theory of meaning for generics can be linked to, and help explain, the role of generics in nonmonotonic reasoning. We can exploit the truth conditional theory sketched in the previous sections to develop an epistemic semantics in which sets of possible worlds will represent information states. It is this epistemic semantics which models defeasible reasoning, and which comprises Asher and Morreau's system CE. The intuitive picture of what goes on when one reasons by defeasible modus ponens is this: first one assumes the premises *birds fly* and *Tweety is a bird*, and no more than this. Second, one assumes that Tweety is as normal a bird as is consistent with just these premises which have been assumed. Finally, one then looks to see whether one is thereby required to believe that Tweety flies or not, and finds that he does. Asher and Morreau (1995) argue that all of the plausible patterns of defeasible reasoning arise in essentially this way, from assuming just their premises, then assuming individuals to be as normal as is epistemically possible, and finally seeing whether one believes their conclusions. The details of this position can be found in (Asher and Morreau, 1995), also (Asher, 1995) and (Morreau, 1992).

To model defeasible reasoning, we will follow Asher and Morreau and build epistemic semantics using information states (which, as we said, are just sets of worlds), and two functions on these information states: updating and normalization.[33] The first of these functions is eliminative, simply removing from an information state all those possible worlds where the sentences with which one is updating are not true. Assuming just the premises of an argument can then be modeled as updating a distinguished informationally-minimal state called **0** (or the set of all the worlds in the canonical model) with those premises.

The second of these functions, normalization, codifies the notion of assuming individuals to be as normal as is consistent with premises. Individuals are assumed to be normal in certain respects, and these respects are represented by a set of propositions P.

[33] Alternatively, we could represent these information states as situations in the sense of Barwise and Perry (1983). Updating and normalization then would involve adding information to a given situation.

Normalization takes place after the just-mentioned updating function has applied. Normalizing the updating result (of updating **0** with a set of premises Γ) in every relevant respect yields a set of information states – which are the fixed points of the normalization process. The conclusions which may reasonably (though not in general validly) be drawn from premises Γ are then those sentences which are true at all the worlds in all of these fixed points. (In general, the order in which the relevant premises are considered for normalization will affect the resulting information state; and it is for this reason that we want the conclusion to hold in every fixed point of the normalization process.)

There are two algorithms one could use to generate the normalizations: a semantic method and a syntactic method. In the semantic method, one looks at each proposition defined by an instantiation of one of the antecedents of a $>$ statement that occurs in the updated premises by one of the constants that occurs in the premises (say this antecedent was instantiated to yield the proposition Fa). Given that the updating process has left us in state S, and that we want now to normalize with respect to this instantiation of a $>$ antecedent, we need to find the set of worlds in S that are normal with respect to Fa and with any $\sim$Fa worlds in S (if there are any). Normalization just throws out the abnormal Fa worlds from S, if doing so leaves us with some worlds – i.e., if the result is consistent. As mentioned, we need to do this with respect to *each* instantiation of *every* antecedent of a $>$ statement in the updated premises; and then we need to consider *every* order of normalization of premises before we can know what conclusion follows from the initial premises.

The second algorithm that we can use is the syntactic method. Normalization in effect converts $>$ statements into $\rightarrow$ statements. And the syntactic view of normalization makes this explicit. Take the theory T and consider each proposition defined by an instantiation of the antecedent (as in the semantic method). Let Fa be one of these. Normalization of T with respect to Fa is: $\mathrm{T} \cup \{(\mathrm{Fa} \rightarrow \mathrm{q}) \mid \mathrm{T} \vdash (\mathrm{Fa} > \mathrm{q})\}$, if the result is consistent. The normalization is simply T otherwise. We carry out an example of updating and normalization in the Appendix.

We shall write this notion of nonmonotonic consequence as $\Gamma \approx_P \mathrm{A}$. The subscript indicates the set of respects in which Γ has been normalized. $\approx$ is a defeasible consequence relation which generates the patterns of defeasible reasoning which we set out to capture. Also note that $\Gamma \models \mathrm{A} \Rightarrow \Gamma \approx \mathrm{A}$, and so commonsense entailment is supra-classical.[34]

4.4. *Patterns of reasonable and defeasible inference with generics*

In this section, we return briefly to the patterns of nonmonotonic reasoning which motivated some of the nonmonotonic formalisms for generics. Notice that these patterns are not to be thought of as completely general argument schemas, into which arbitrary generic sentences can be instantiated. Take for instance the pattern of defeasible modus ponens, which is schematically represented as: from premises $\forall x(\mathrm{A} > \mathrm{B})$ and $\mathrm{A}(t)$ follows $\mathrm{B}(t)$. But this conclusion is supposed no longer to follow if the premise $\neg\mathrm{B}(t)$ is

[34] For more details on the model theory of commonsense entailment see Asher and Morreau (1991, 1995) and Morreau (1992). For more concerning an accompanying proof theory, see Asher (1995).

added, which intuitively speaking amounts to adding the additional information that t, though an A, is not a normal A.

The following fact states conditions under which commonsense entailment captures the patterns of defeasible modus ponens:

FACT.

(i) Let $(\Gamma \cup \{\forall x(\mathrm{A} > \mathrm{B}),\ \mathrm{A}(t)\})$ be $\vdash$ consistent, and let $\{\mathrm{A}(t)\}$ be the set P of normalization respects. Then:

$$\Gamma,\ \forall x(\mathrm{A} > \mathrm{B}),\ \mathrm{A}(t) \mathrel{|\!\approx_P} \mathrm{B}(t).$$

(ii) If $\Gamma \cup \{\forall x(\mathrm{A} > \mathrm{B}),\ \mathrm{A}(t)\} \nvdash \mathrm{B}(t)$, then if $\{\mathrm{A}(t)\}$ is the set P of normalization respects, then:

$$\Gamma,\ \forall x(\mathrm{A} > \mathrm{B}),\ \mathrm{A}(t), \neg\mathrm{B}(t) \mathrel{|\!\not\approx_P} \mathrm{B}(t).$$

The restrictions of this FACT are sufficient to verify Defeasible Modus Ponens. A proof is contained in (Asher and Morreau, 1995), as well as in (Morreau, 1992). A similar fact establishes that commonsense entailment indeed predicts a skeptical conclusion from the premises of the Nixon Diamond but that it also predicts the conclusion that Nixon is politically active from the premises of the Extended Nixon Diamond.

By adding more constraints on *, it becomes possible to prove that other defeasible inferences can be consistently inferred within commonsense entailment.[35] The Penguin principle is captured by having the OR condition on * (see Section 4.1 above). Morreau (1992) shows that, when * is subject to this additional constraint, commonsense entailment captures the Penguin Principle in addition to Defeasible Modus Ponens and the Nixon Diamond. To capture the Weak Penguin Principle, we must add another constraint on *:

SPECIFICITY. If $^*(w,p) \subseteq q$, $^*(w,p) \cap {}^*(w,q) = \emptyset$, and $^*(w,p) \neq \emptyset$, then $^*(w,q) \cap p = \emptyset$.

Which corresponds to the axiom:

$$((\mathrm{A} > \mathrm{B})\ \&\ (\mathrm{B} > \mathrm{C})\ \&\ (\mathrm{A} > \neg\mathrm{C})) \rightarrow (\mathrm{B} > \neg\mathrm{A}).$$

Morreau (1992) shows that commonsense entailment then still preserves the inference patterns of Defeasible Modus Ponens. But it becomes difficult to verify the other inferences model-theoretically. As we place more constraints on * it becomes more difficult to verify that our candidates for "survivor" worlds actually exist in the canonical model. So one might wish to pursue alternative lines concerning defeasible reasoning, such as done in (Asher, 1995) or in (Asher and Morreau, 1995). The idea in the latter is of ordered normalization, which corresponds closely to the strategy in prioritized circumscription (McCarthy, 1986) for capturing inferences like the Penguin Principle or the Weak Penguin Principle, where certain minimizations of extensions of predicates are performed

[35] For details see Morreau (1992).

before others. Also in extensions of Reiter's default logic (for instance, Brewka, 1991), orderings on the application of default rules are used to construct extensions that will verify these principles.

Other approaches that use conditionals as the semantic treatment of generics are the nonmonotonic logics of Delgrande (1987), Pearl's (1990) system Z, Boutilier (1994), Bacchus, Grove, Halpern and Koller (1993). All these systems verify the Penguin Principle and the Weak Penguin Principle. Lehmann (1989) provides a system, called Rational Closure, equivalent to that of Pearl's system Z but treats generics more as akin to some sort of default rule. Benferhat, Dubois and Prade (1992) provide yet another system equivalent to Lehmann's Rational Closure in which generics are thought of as constraints on an appropriate ordering of models. Often these systems are simpler than commonsense entailment and yet appear to produce equivalent results. But this only holds for simple fragments of the language and for a restricted notion of what are the acceptable defeasible patterns of reasonable inference (Asher, 1995).

What many of these simpler systems cannot do is extend readily to handle more complex forms of defeasible reasoning in which generics appear as conclusions, as well as premises. Asher and Morreau (1991) attempt to capture defeasible reasoning that has generic conclusions by complicating the normalization function. More specifically, an extra clause was added which used the behavior of * at a world to remove possible worlds from information states. So in the case of defeasible transitivity for example, possible worlds would be removed where it does not hold that *sparrows fly*. These are worlds where something is a normal sparrow without being a normal bird, and that is a possibility which this stronger notion of normalization filtered out.

We think it an extremely powerful positive feature of commonsense entailment that, unlike the theories of nonmonotonic reasoning we surveyed in Section 3, it makes the behavior of generics within patterns of nonmonotonic reasoning be determined by their truth conditional meaning. Unlike circumscription (in which generics are treated as extensional) or probabilistic theories of nonmonotonic reasoning (in which generics are treated as probability statements or statements with probabilistic quantifiers), commonsense entailment has the *prima facie* philosophical virtue of basing the acceptability of nonmonotonic reasoning patterns on a truth conditional semantics for generics that is, if we are correct, more plausible than the other alternatives.

5. Summary and conclusions

Generics attract the attention of logically-oriented scholars because the semantic phenomena surrounding them have been implicated in "default reasoning". But this implication is, upon scrutiny, rather tenuous. This article reports why the implication is tenuous and makes some suggestions as to the direction such logical accounts should take in order to accommodate generics.

We first investigated the full range of linguistic and semantic phenomena involving generics. We distinguished generic reference from generic predication – settling on the latter as relevant to default reasoning. We argued that these "characterizing statements" are either literally true or literally false – and are not therefore either "strictly false but acceptable" nor "merely directives about what sort of inferences to draw and what to expect in our experience." We then investigated the range of "exceptions" that a true generic can tolerate while remaining true, and our conclusion was that there is no unique number or proportion that will account for all the various generic statements. It follows that there is no quantifier, not even a "vague" probabilistic one, like *in a significant number of cases*, which can be used to give an account of the truth-functional force of a characterizing statement.

Part of the difficulty is that generics come with a kind of nomic force; they express a "general truth" which is enforced by some sort of law. Generics about the biological world, for example, are backed by the laws of biology; those about individual people are backed by psycho-social laws that describe what such people do. A purported generic statement which does not have such a backing is simply not true, even if the objects of which the subject is true in fact all have the property designated by the predicate. Even were it the case that all the remaining koalas in the world were to lose a paw, it still would not be true that koalas are three-pawed. For a generic statement to be true, it simply cannot report a mere accidentally true generalization.

There are two aspects to generic statements: their nomic force and their extensional consequences. In a large number of cases the extensional consequence of a generic is that most, or most of the important or noticeable or significant, instances of the subject term manifest the property indicated by the predicate. But we also noted that we cannot validly infer that any particular individual falling under the subject term has the property indicated by the predicate. As a we have indicated, not all generics have this extensional consequence; but perhaps there is an identifiable and distinguishable subset of them which do. For these types of generic statement, what would an appropriate analysis of just the extensional consequences be? We surveyed seven possibilities: relevant quantification, abstract entities, prototype theory, stereotypes, modal conditionals, situation semantics, and default theories. We found that none of the seven fared very well, but that the modal conditional and situation semantics approaches seemed to hold some promise especially as accounts of (some of) the nomic force of generic statements. The default approach fared best as an account of some of the extensional inferences we may wish to draw in the face of "generic information", and so we investigated four different approaches to default reasoning: default logic, autoepistemic logic, circumscription, and probabilistic logic.

In the end we found each of these inadequate. The problems we found ranged from incorrect prediction of conclusions to be drawn from generics, to denying truth conditional status to generics, to making dramatically incorrect predictions concerning the truth conditions of certain generics, to being unable to integrate the account with epistemic and modal logics necessary to account for the nomic force of generics.

Finally we presented Commonsense Entailment as a theory which could address many of these issues. It adopts a modal conditional analysis of generic statements, and locates the nomic force of generics in this area. But it does not posit "absolutely normal" possible worlds, nor does it locate "nonmonotonicity" there. Rather it adopts the situation-semantic strategy of judging "how much information" is in the premises of an argument, and allows that an increase here can make us wish to retract previous conclusions. Note that the generic premises continue to be true – nonmonotonicity is not to be found in denying generic statements.

After presenting the formal structure of Commonsense Entailment, we offered an evaluation of it in terms of the sort of argumentation we used against the other accounts of default reasoning as theories of generics, and we conclude that it fares well indeed. We also, however, note that there are some objections of a more general nature concerning generic statements that CE does not answer. The answer to these objections will require more investigation into the logical structure of generics.

Appendix

In Section 4.3 of the main text we presented two equivalent ways of calculating the defeasible entailments of a theory containing generics, a semantic method and a syntactic method. As an example of how our theory works, we present an extended example. Consider the theory T_0:

1a. Isis is a cat.
2a. Cats meow.
3a. Cats like to eat fish.
4a. Leonore is a cat who does not like to eat fish.
5a. Cats that like to eat fish are not fat.
6a. Cats are necessarily mammals.
7a. Mammals do not meow.

We will now see what follows from T_0 using our nonmonotonic formalism. Our theory will make defeasible predictions about particular cats like Isis and Leonore. Recall that on the syntactic theory what we do is add the instantiations of our generic statements with constants mentioned in the theory or added as witnesses for existential claims. In T_0, we have two constants: Isis and Leonore, but no existentials (hence no need for witnesses). The theory makes no links between Isis's properties and Leonore's so we may consider each of them separately. Let's consider Isis first, and the instantiations of the various conditionals to Isis.

1b. Cat(isis).
2b. $\text{Cat(isis)} > \text{Meows(isis)}$.
3b. $\text{Cat(isis)} > (\forall x)(\text{Fish}(x) > \text{Likes-to-eat(isis}, x))$.
5b. $\text{Cat(isis)} > ((\forall x)(\text{Fish}(x) > \text{Likes\_to\_eat(isis}, x)) > {\sim}\text{Fat(isis)})$.
6b. $\Box\ (\text{Cat(isis)} \rightarrow \text{Mammal(isis)})$.
7b. $\text{Mammal(isis)} > \text{Meows(isis)}$.

We now consider the conditionals (2) and (3), involving the antecedent Cat(isis) and consider the result of changing the main occurrences of the $>$ connectives in them into material conditionals. This theory, the $T_{\text{Cat(isis)}\rightarrow}$ theory, looks like this:

$$\text{Cat(isis)} \rightarrow \text{Meows(isis)}.$$

$$\text{Cat(isis)} \rightarrow (\forall x)(\text{Fish}(x) > \text{Likes-to-eat(isis}, x)).$$

Now the question is, are the consequents and antecedents of these conditionals consistent with our original theory? The answer is yes (which we could show by constructing a model of the original theory together with the antecedents and consequents of $T_{\text{Cat(isis)}\rightarrow}$ but we will forego that here). So according to our "syntactic" procedure for calculating defeasible entailments, we will add $T_{\text{Cat(isis)}\rightarrow}$ to our original theory, calling the result T_1. From T_1, we infer by Modus Ponens:

$$\text{Meows(isis)}$$
$$(\forall x)(\text{Fish}(x) > \text{Likes-to-eat}(\text{isis}, x)).$$

We now consider another $>$ antecedent of our theory, say

$$(\forall x)\ (\text{Fish}(x) > \text{Likes\_to\_eat}(\text{isis}, x)),$$

and collect all the conditionals entailed by T_1 with this antecedent. Beyond the instantiations of the axioms, which don't add any new information to the information state, we have just the conditional:

$$(\forall x)\ (\text{Fish}(x) > \text{Likes-to-eat}(\text{isis}, x)) > \text{Fat(isis)},$$

which we now convert into a material conditional

$$(\forall x)\ (\text{Fish}(x) > \text{Likes-to-eat}(\text{isis}, x)) \rightarrow {\sim}\text{Fat(isis)}. \qquad (*)$$

We now again test to see whether the consequent and antecedent of this conditional are consistent with T_1. Inspection shows that they are. So we add $(*)$ to T_1, forming the theory T_2, in which we infer by Modus Ponens that Isis is not fat.

The last antecedent that we must consider for Isis is Mammal(isis). It follows monotonically from T_1 that Mammal(isis), since 'cats are necessarily mammals' is a strict conditional. But it also follows in the monotonic logic of commonsense entailment (for details see Morreau, 1992) that:

$$\text{Mammal(isis)} > {\sim}\text{Cat(isis)}.$$

This conditional is a theorem of T_1 because cats and mammals have in T_1 (and in our original theory as well!) incompatible generic properties – cats meow and mammals don't. Let's now form the $\rightarrow$ theory in which all the $>$ conditionals with Mammal(Isis) as an antecedent are translated into material conditionals. This theory looks like this:

$$\text{Mammal(isis)} \rightarrow {\sim}\text{Cat(isis)}.$$
$$\text{Mammal(isis)} \rightarrow {\sim}\text{Meows(isis)}.$$

The consequents of this theory are evidently not consistent with T_2 (or even with our original theory T_0!). So we do *not* add $T_{\text{Mammal(isis)}\rightarrow}$ to T_2.

We claim that had we pursued these normalizations in a different order, we would still have arrived at T_2 as our final theory concerning Isis. One might wonder about how on a different ordering of normalizations with respect to the antecedents of $>$ conditionals in T_0 we could be assured of arriving at the defeasible conclusion that Isis is not fat. Suppose we had first considered adding to T_0 the material conditional

$$(\forall x)\ (\mathrm{Fish}(x) > \text{Likes-to-eat}(\mathrm{isis}, x)) \rightarrow {\sim}\mathrm{Fat}(\mathrm{isis}). \tag{$*$}$$

$(*)$ is evidently consistent with T_0 and so could be added, but we could not yet deduce the conclusion that Isis was not fat. But once of course we add $T_{\mathrm{Cat(isis)}\rightarrow}$ we would get the defeasible entailment. Thus, we have the following set of defeasible entailments concerning Isis – namely that she meows, she eats fish, and she is not fat.

Of course we have not yet computed all the defeasible consequences of our original theory T_0, since we have not examined any inferences concerning Leonore. So let us continue and consider the conditionals with antecedents Cat(leonore). Once again there are two:

$$\mathrm{Cat}(\mathrm{leonore}) \rightarrow \mathrm{Meows}(\mathrm{leonore}).$$
$$\mathrm{Cat}(\mathrm{leonore}) \rightarrow (\forall x)(\mathrm{Fish}(x) > \text{Likes-to-eat}(\mathrm{isis}, x)).$$

Because Leonore, unlike Isis, does not like to eat fish, according to T_0 (and hence T_2) we cannot add this set of conditionals to T_2. Normalization with the proposition that Leonore is a cat yields no new inferences. Similarly we cannot infer anything about Leonore's girth, because the antecedent of that conditional, namely

$$(\forall x)\ (\mathrm{Fish}(x) > \text{Likes-to-eat}(\mathrm{leonore}, x)),$$

is inconsistent with our information. But neither can we infer from Mammal(leonore) that she does not meow, because after all, she is a cat. So we end up knowing very little about Leonore. Since we could have performed these normalizations in any order while still eventually getting the same results as in T_2, T_2 represents all the defeasible entailments of the original theory T_0.

References

Adams, E. (1975), *The Logic of Conditionals*, Reidel, Dordrecht.

Asher, N. (1984), *Linguistic understanding and non-monotonic reasoning*, Proceedings of the 1984 Nonmonotonic Reasoning Workshop, American Association for Artificial Intelligence, New Paltz, NY, 1–20.

Asher, N. (1995), *Commonsense entailment: A logic for some generics*, Conditionals and Artificial Intelligence, G. Crocco, L. Farinas del Cerro and A. Herzig, eds, Oxford Univ. Press, Oxford, 115–156.

Asher, N. and Morreau, M. (1991), *Commonsense entailment: A modal theory of nonmonotonic reasoning*, Proceedings of the 12th IJCAI, Morgan Kaufmann, Los Altos, CA, 387–392.

Asher, N. and Morreau, M. (1995), *What some generics mean*, The Generic Book, G. Carlson and F.J. Pelletier, eds, Chicago Univ. Press, Chicago, IL, 224–237.

Bacchus, F., Grove, A., Halpern, J. and Koller, D. (1993), *Statistical foundations for default reasoning*, Proceedings of the 13th IJCAI, Morgan Kaufmann, Los Altos, CA, 563–569.

Barth, E. (1974), *The Logic of the Articles in Traditional Philosophy*, Reidel, Dordrecht.

Barwise, J. (1986), *Conditionals and conditional information*, On Conditionals, E.C. Traugott, A. ter Meulen, J. Snitzer Reilly and Ch.A. Ferguson, eds, Cambridge Univ. Press, Cambridge, 21–54.

Barwise, J. and Perry, J. (1983), *Situations and Attitudes*, MIT Press, Cambridge, MA.

Benferhat, S., Dubois, D. and Prade, H. (1992), *Representing default rules in possibilistic logic*, Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning, J. Allen, R. Fikes and E. Sandewall, eds, Morgan Kaufmann, Los Altos, CA, 673–684.

Benferhat, S., Cayrol, C., Dubois, D., Lang, J. and Prade, H. (1993), *Inconsistency Management and Prioritized Syntax-Based Entailment*, IJCAI-93 (Chambery), 640–645.

Berkeley, G. (1710), *A treatise concerning the principles of human knowledge*, Berkeley's Philosophical Writings, D. Armstrong, ed., Reprinted, with many others of his works, together with a penetrating commentary, Macmillan, New York, 1965.

Bosanquet, B. (1888), *Logic, Or the Morphology of Knowledge*, Oxford Univ. Press, Oxford.

Boutilier, C. (1992), *Conditional logics of normality as modal systems*, AAAI, MIT Press, Cambridge, MA, 594–599.

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

Brewka, G. (1991), *Cumulative default logic: In defense of nonmonotonic inference rules*, Artif. Intell. **50**, 183–205.

Carlson, G.N. (1977), *Reference to Kinds in English*, PhD Dissertation, University of Massachusetts. Published 1980, Garland, New York.

Carlson, G.N. (1979), *Generics and atemporal when*, Ling. and Philos. **3**, 49–98.

Chellas, B. (1980), *Modal Logic: An Introduction*, Cambridge Univ. Press, Cambridge, UK.

Declerck, R. (1991), *The origins of genericity*, Linguistics **29**, 79–101.

Delgrande, J. (1987), *A semantics for defaults using conditional logic*, Artif. Intell. **33**, 105–130.

Elio, R. and Pelletier, F.J. (1993), *Human benchmarks on AI's benchmark problems*, Proceedings of the 15th Congress of the Cognitive Science Society, Boulder, 406–411.

Elio, R. and Pelletier, F.J. (1996), *On reasoning with default rules abd exceptions*, Proceedings of the 18th Congress of the Cognitive Science Society, San Diego.

Fine, K. (1985), *Reasoning with Arbitrary Objects*, Basil Blackwell, Oxford.

Frege, G. (1904), *Was ist eine Funktion?*, Festschrift Ludwig Boltzmann gewidmet zum sechzigsten Geurtstag, 20. Februar 1904, S. Meyer, ed., Ambrosius Barth, Leipzig, 656–666. Translation and reprint in P. Geach and M. Black, *Translations from the Philosophical Writings of Gottlob Frege*, Oxford, Blackwell, 107–116.

Gabbay, D. (1985), *Theoretical foundations for non-monotonic reasoning in expert systems*, Logics and Models of Concurrent Systems, K. Apt, ed., Springer, Berlin, 439–457.

Geach, P. (1962), *Reference and Generality*, Cornell Univ. Press, Ithaca.

Goldszmidt, M. and Pearl, J. (1990), *A maximum entropy approach to nonmonotonic reasoning*, Proceedings of AAAI-90, Boston, 646–652.

Gottlob, G. (1993), *The power of beliefs, or translating default logic into standard autoepistemic logic*, Proceedings of the 13th IJCAI, Morgan Kaufmann, Los Altos, CA, 570–575.

Heim, I. (1982), *The semantics of definite and indefinite noun phrases*, PhD Dissertation, University of Massachusetts, Amherst.

Heyer, G. (1987), *Generische Kennzeichnungen*, Philosophia Verlag, Munich.

Hilbert, D. and Bernays, P. (1934), *Grundlagen der Mathematik*, Springer, Berlin.

Jespersen, O. (1924), *The Philosophy of Grammar*, Republished 1965 by Norton, New York.

Konolige, K. (1988), *On the relationship between default and autoepistemic logic*, Artif. Intell. **35**, 343–354.

Kratzer, A. (1981), *The notional category of modality*, Words, Worlds and Contexts: New Approaches to Word Semantics, H.J. Eikmayer and H. Rieser, eds, Gruyter, Berlin.

Kraus, S., Lehmann, D. and Magidor, M. (1990), *Nonmonotonic reasoning, preferential models and cumulative logics*, Artif. Intell. **44**, 167–207.

Krifka, M., Pelletier, F.J., Carlson, G., Ter Meulen, A., Chierchia, G. and Link, G. (1995), *Genericity: An introduction*, The Generic Book, G.N. Carlson and F.J. Pelletier, eds, Chicago Univ. Press, Chicago, 1–124.

Lawler, J. (1973), *Studies in English Generics*, University of Michigan Papers in Linguistics vol. 1, Ann Arbor.

Lehmann, D. (1989), *What does a conditional knowledge base entail?*, Principles of Knowledge Representation and Reasoning: Proceedings of the First International Conference, R. Brachman and H. Levesque, eds, Morgan Kaufmann, 212–222.

Lewis, D. (1972), *General semantics*, Semantics of Natural Language, D. Davidson and G. Harman, eds, Reidel, Dordrecht, 169–219.

Lewis, D. (1973), *Counterfactuals*, Basil Blackwell, Oxford.

Lewis, D. (1975), *Adverbs of quantification*, Formal Semantics of Natural Language, Ed. Keenen, ed., Cambridge Univ. Press, Cambridge, MA.

Lewis, D. (1983), *Scorekeeping in a language game*, Reprinted in his Philosophical Papers, vol. 1, Oxford Univ. Press, Oxford.

Lifschitz, V. (1985), *Computing circumscription*, Proceedings IJCAI-85, Los Angeles, 121–127.

Lifschitz, V. (1989), *Benchmark problems for formal nonmonotonic reasoning, v. 2.0*, Proceedings of the Second International Workshop on Non-Monotonic Reasoning, M. Reinfrank, J. de Kleer, M. Ginsberg and E. Sandewall, eds, Lecture Notes in Artif. Intell. vol. 346, 202–219.

Locke, J. (1690), *An Essay Concerning Human Understanding.* The current edition of choice is edited by P. Nidditch, Oxford Univ. Press, 1975.

McCarthy, J. (1980), *Circumscription – A form of non-monotonic reasoning*, Artif. Intell. **13**, 27–39.

McCarthy, J. (1986), *Applications of circumscription to formalize common sense knowledge*, Artif. Intell. **28**, 89–116.

McDermott, D. and Doyle, J. (1980), *Non-monotonic logic I*, Artif. Intell. **13**, 41–72.

Meinong, A. (1904), *Über Gegenstandstheorie*. Reprinted as *The theory of objects*, Realism and the Background of Phenomenology, R. Chisholm, ed., Free Press, Glencoe, 1960, 76–117.

Milsark, G. (1974), *Existential sentences in English*, PhD Dissertation, MIT, Distributed by Indiana University Linguistics Club.

Moore, R. (1984), *Possible-world semantics for autoepistemic logic*, Proceedings of the 1984 Nonmonotonic Reasoning Workshop, American Association for Artificial Intelligence, New Paltz, NY, 334–354.

Morreau, M. (1992), *Three applications of imaging*, Thesis, University of Amsterdam.

Nute, D. (1980), *Topics in Conditional Logic*, Reidel, Dordrecht.

Osherson, D. and Smith, E. (1981), *On the adequacy of prototype theory as a theory of concepts*, Cognition **9**, 35–58.

Pearl, J. (1988), *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo.

Pearl, J. (1990), *System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning*, Theoretical Aspects of Reasoning about Knowledge (TARK III), R. Parikh, ed., Morgan Kaufmann, 121–135.

Putnam, H. (1975), *The Meaning of 'Meaning'*, Language, Mind and Knowledge, K. Gunderson, ed., Univ. of Minnesota Press, Minneapolis.

Reiter, R. (1980), *A logic for default reasoning*, Artif. Intell. **13**, 81–132.

Rosch, E. (1978), *Principles of Categorization*, Cognition and Categorization, E. Rosch and B.B. Lloyds, eds, Lawrence Erlbaum Assoc., New York, 27–48.

Routley, R. (1967), *A simple natural deduction system*, Logique et Anal. **12**, 129–152.

Russell, B. (1905), *On denoting*, Mind **14**, 479–493.

Schubert, L. and Pelletier, F.J. (1987), *Problems in the representation of the logical form of generics, bare plurals, and mass terms*, New Directions in Semantics, E. LePore, ed., Academic Press, New York, 385–451.

Schubert, L. and Pelletier, F.J. (1989), *Generically speaking, or, using discourse representation theory to interpret generics*, Property Theory, Type Theory, and Semantics vol. II: Semantic Issues, G. Chierchia, B. Partee and R. Turner, eds, Reidel, Dordrecht, 193–267.

Scribner, S. (1977), *Modes of thinking and ways of speaking: Culture and logic reconsidered*, Thinking: Readings in Cognitive Science, P.N. Johnson-Laird and P.C. Wason, eds, Cambridge Univ. Press, Cambridge, MA.

Smirnov, V. (1976), *The theory of quantification and the ε-calculus*, Essays on Mathematical and Philosophical Logic, J. Hintikka, ed., Reidel, Dordrecht.

Smith, E. and Osherson, D. (1984), *Conceptual combination with prototypical concepts*, Cogn. Science **8**, 357–361.

Stalnaker, R. (1968), *A theory of conditionals*, Studies in Logical Theory, N. Rescher, ed., Basil Blackwell, Oxford, 98–112.

Stalnaker, R. and Thomason, R. (1967), *A Semantic Analysis of Conditional Logic*, Mimeograph, Yale University.

Touretzky, D. (1986), *The Mathematics of Inheritance Systems*, Morgan Kaufmann, Los Altos, CA.

Veltman, F. (1995), *Defaults in update semantics*, J. Philos. Logic, forthcoming.

von Fintel, K. (1994), *Restrictions on quantifier domains*, PhD Dissertation, University of Massachusetts, Dept. Linguistics.