Default Reasoning and Generics

Ariel Cohen
arikc@bgumail.bgu.ac.il

Appeared in *Computational Intelligence* 13:4, 1997

Abstract

There are numerous logical formalisms capable of drawing conclusions using default rules. Such systems, however, do not normally determine where the default rules come from; i.e. what it is which makes *Birds fly* a good rule, but *Birds drive trucks*—a bad one.

Generic sentences such as *Birds fly* are often used informally to describe default rules. I propose to take this characterization seriously, and claim that a default rule is *adequate* iff the corresponding generic sentence is true. Thus, if we know that Tweety is a bird, we may conclude, by default, that Tweety flies, just in case *Birds fly* is a true sentence.

In this paper, a quantificational account of the semantics of generic sentences will be presented. It will be argued that a generic sentence is evaluated not in isolation, but with respect to a set of relevant alternatives. For example, *Mammals bear live young* is true because among mammals which bear live young, lay eggs, undergo mitosis or engage in some alternative form of procreation, the majority bear live young. Since male mammals do not procreate in any form, they do not count. Some properties of alternatives will be presented, and their interactions with the phenomena of focus and presupposition will be investigated.

It will be shown how this account of generics can be used to characterize adequate default reasoning systems, and several desirable properties of such systems will be proved. The problems of the automatic acquisition of rules from natural language will be discussed. Since rules are often explicitly expressed as generics, it will be argued that the interpretation of generic sentences plays a crucial role in this endeavor, and it will be shown how the theory presented here can facilitate such interpretation.

Key words: generics, default reasoning, semantics, automatic knowledge acquisition.

1 Introduction

A reasoner does not normally have complete knowledge about a problem it attempts to solve; it must, therefore, be able to draw conclusions that are plausible, though not certain. The reasoner must also be able to retract these conclusions if new knowledge becomes available which contradicts them. In other words, such reasoning is *nonmonotonic*; unlike deductive reasoning, adding more knowledge to a nonmonotonic system may result in the retraction of previously derived conclusions. Considerable amount of work has been conducted on default reasoning, and many formalisms have been proposed to facilitate it. The problem which those studies attempt to solve may be described as follows: given a set of default rules (possibly in addition to deductive rules, and/or ordered by priority), and a set of propositions known to be true, which conclusions should be drawn, and how are they to be derived?

Systems of default reasoning have been proposed to model common-sense reasoning. However, for this goal to be achieved, drawing appropriate conclusions from rules is not sufficient; one has to make sure the "right" rules are used. For example, the canonical example of a default rule is the one stating that, by default, birds fly. Any default reasoning system worth its salt should be able to infer that, if Tweety is a bird, and nothing further is known about Tweety, it flies. This conclusion may be retracted if it later becomes known that Tweety does not fly, say because it is, in fact, a penguin. But note that we might, instead, use a rule which states that birds do not fly, and only infer that they do if we have evidence for it. We might even have a rule stating that birds drive trucks, write dissertations or are married bachelors. Such rules would be harmless, in the sense that the conclusions drawn from them will never be inconsistent. For example, if we know that all truck drivers are human, and that no birds are human, we will be prevented from concluding that Tweety drives a truck, even though Tweety is a bird. Still, having such a rule seems intuitively wrong, whereas having a rule stating that birds fly seems intuitively right. The question is, then, how to capture formally this intuition; how to determine which rules are desirable and which are not.[1]

In the literature concerning default reasoning, default rules are usually described informally using generic sentences, of which (1) is a typical example:

(1) Birds fly.

I propose to take this representation seriously. That is to say, I propose that a default rule be *adequate* just in case the corresponding generic sentence is true.

The question of what makes a generic sentence true is by no means an easy one; some would even claim it has no general answer. Yet we are in need of such an answer. Generic sentences are prevalent in language; take the previous sentence, for example. Indeed, one needs only to look at a newspaper article, not to mention an encyclopedia, to find numerous instances of generics. Neither a theory of language meaning nor a natural language processing system can easily

[1] For more detailed comments in a similar vein, see Israel (1980); Brachman (1985); Loui (1986); Neufeld (1989).

afford to ignore them. In the following section I am going to propose a novel truth-conditional theory (i.e. a theory of generics as functions from states of affairs to truth values) of the meaning of generic expressions; then, I will present and prove some desirable properties of default reasoning systems whose rules correspond to true generic sentences; finally, I will discuss some implications of the research presented here to the problem of automatic acquisition of rules from natural language.

The goal of this paper, then, is three-fold: one is a formal representation of generic sentences which captures their truth conditions correctly; another aim is to provide means by which rules used in default reasoning can be evaluated; yet a third goal is to facilitate automatic acquisition of desirable rules given a natural language text.

2 The meaning of generic sentences

A significant portion of our knowledge about the world is expressed by sentences such as the following:

(2) a. Dogs are mammals.
b. Birds fly.
c. Mammals give birth to live young.
d. The Frenchman eats horsemeat.
e. Bulgarians are good weightlifters.

Such statements are usually referred to as *generics*. This name implies that they are used to describe generalizations. But what is the meaning of such generalizations? How do we know whether a generalization is true or false? The answer is not immediately clear. Sentence (2.a) seems to hold for all dogs, (2.b) for most birds, (2.c) for most female mammals (presumably less than half the total number of mammals), (2.d) for few Frenchmen and (2.e) for very few Bulgarians. While most researchers believe that generics do, indeed, have truth conditions, observations such as the ones above have led many of them (most notably Carlson 1977) to the conclusion that no quantificational account of generics is possible. In other words, it is impossible to consider instances of dogs, birds etc., determine how many of them satisfy a certain property, and on that basis conclude whether a given generic sentence is true or false. Restricting the domain of quantification (by context or other factors), it is claimed, will not help; there would still be no principled way to relate the truth or falsity of a generic to the truth or falsity of statements about individual instances. Thus the truth conditions of a generic are held by these workers to be somewhat mysterious; they are related not to properties of individuals, but to some ontologically irreducible "rules and regulations" (Carlson 95), e.g. primitive properties of kinds (Carlson 77) or rules determining what is "normal" (Delgrande 1987). In this paper I intend to demonstrate that a quantificational theory is not only possible, but is, in fact, advantageous.

2.1 Alternative-based generics

Before describing my proposed solution, let me aggravate the problem. Sentences such as (3) (from Thomason 1988) pose a difficult problem for any theory of generics:

(3) People have a hard time finding Carnegie Mellon University.

The vast majority of people do *not* have a hard time finding CMU, for the simple reason that they never try to find it. Only a handful of people, namely some significant proportion of first-time visitors to CMU, do have a hard time finding CMU.

To find out what makes (3) true, it is instructive to try to see what sort of people count as exceptions to it. As Thomason points out, people who do not try to find CMU, either because they are not interested in finding it, or because they already know where it is (e.g. a professor who goes to CMU every morning), do not count as exceptions. Only people who do try to find CMU, but do not have a hard time finding it (presumably having an easy time), count as exceptions. I suggest that when we evaluate the truth of (3), we are comparing people who find CMU with difficulty to people who find CMU with ease: if more people experience difficulties than those who do not, (3) is true, otherwise it is false. All other people, such as those who have never heard of CMU or those who work there, neither have a hard time finding CMU, nor an easy time finding it, nor, indeed, any kind of time finding CMU. These people are irrelevant to the truth of (3)—they are simply not counted. In other words, we consider two alternative levels of difficulty for the task of finding CMU, and if a high level of difficulty is the most common one, we conclude that (3) is true.

Let us try to generalize from this. Suppose we are given a generic sentence S, i.e. a sentence which predicates some property ϕ of a kind κ. Let us assume that the property ϕ is a member of a set A of alternative properties. Then if there are more instances of κ which satisfy ϕ than those which satisfy some alternative $\phi' \in A$ but do not satisfy ϕ, we conclude that S is true. Put another way, S is true just in case the majority of instances of κ which satisfy one of the alternatives in A—satisfy ϕ.

It will be instructive to see how this proposal deals with sentence (2.c), which belongs to a type of sentences which are notoriously problematic for quantificational accounts of generics. Why is it that (2.c) is true although most mammals do not, in fact, give birth to live young? Suppose the property *give birth to live young* is a member of the set of alternative means to produce offspring, say {*give birth to live young, lay eggs, undergo mitosis*}. Although less than half of all mammals give birth to live young, it is true that more mammals give birth to live young than those which lay eggs or undergo mitosis, and this is why (2.c) is true.

2.2 Some observations

Since we speak of alternatives, we may wish to enquire whether they are mutually exclusive. It seems that the answer is negative. Consider the true sentence (4):

(4) Israelis speak Hebrew.

The set of alternatives with respect to which (4) is evaluated will, presumably, be of the form $\{speak\ L | L\ is\ a\ language\}$. Since a given individual may be multilingual, the alternatives cannot be mutually exclusive.

Another question is whether majority is not too strong a requirement, and whether plurality may not be more appropriate. That is to say, is it sufficient for ϕ to be more common than any single $\phi' \in A$, or does it have to be more common than the disjunction of all alternatives? In other words, does the winner take all, or does the winner have to be satisfied by a majority of the instances? It would seem that the latter is the case. Given that the most common language in India is Hindi, but that only 30% of the population speak it,[2] (5) is false:

(5) Citizens of India speak Hindi.

Indeed, for every single language L, there are more Indians who speak Hindi than those who speak L; however, it is not the case that there are more speakers of Hindi than speakers of some language or other. Plurality, then, is not sufficient, and the majority requirement is, indeed, necessary.[3]

There is still a lingering problem for our definition, and this is the observation that generics seem to refer to an unbounded number of instances. Even though the number of birds is finite, *Birds fly* seems to say something not just about current actual birds, but about future and possible birds as well. Indeed, even if a generalization does happen to hold for a finite domain, but is not expected to hold in the future, the generic sentence is not judged true. Suppose it happened to be the case that all current Supreme Court judges had an odd Social Security number; sentence (6), nevertheless, would not be true:

(6) Supreme Court judges have an odd Social Security number.

It may be noted that there are strong similarities between generics and statements of probability judgments. Thus, sentence (7) seems false for the same reasons that (6) is:

[2] This information was taken from the Academic American Encyclopedia.

[3] But cf. Shastri (1989) who, in a somewhat different context, takes the view that the winner does take all.

(7) The probability for a Supreme Court judge to have an odd Social Security number is high.[4]

The situations in which (6) and (7) would become true are similar too: if the property of having an odd Social Security were not *accidental* for a Supreme Court judge—say if judges were chosen by a lottery, and the computer program performing the lottery contained a bug which caused it to choose odd numbers only—both would then be true. We will, therefore, use probabilities in our definition, rather than simple counting.[5]

We will, then, formally define the truth conditions of a generic sentence as follows:

Definition 1 (Generic truth conditions) *Let* $\mathbf{gen}_A(C(\kappa))(\phi)$ *be a proposition, where* ϕ *is a property,* κ *is a kind and* $C(k)$ *is a function which maps a kind onto its instances. Let* A *be a set of alternatives to* ϕ*. Then* $\mathbf{gen}_A(C(\kappa))(\phi)$ *is true iff*

$$P(\phi \wedge (\bigvee A)|C(\kappa)) > P((\neg\phi) \wedge (\bigvee A)|C(\kappa)),$$

or, equivalently (assuming $P(C(\kappa) \wedge (\bigvee A)) > 0$*),*

$$P(\phi|C(\kappa) \wedge (\bigvee A)) > 0.5$$

(where $P(A|B)$ *indicates the conditional probability of* A *given* B*, and* $\bigvee A$ *is the disjunction of the members of* A*).*[6]

2.3 Determining the alternatives

The theory of generics which I am proposing here relies rather crucially on the notion of a relevant set of alternatives. Alternatives and related notions play an important role in linguistics: *P-sets* (Rooth 1985) and *C-sets* (Rooth 1992) in the theory of focus, *scales* in theories of scalar implicature (Horn 1972; Hirschberg 1985), *comparison classes* in theories of adjectives (Klein 1980; Ludlow 1989) are just a few examples. However, very little work has been done on the formal

[4] There is a reading of (7) under which it is true, namely that if we picked at random one of the *current* Supreme Court judges, he or she would be likely to have an odd Social Security number. This is what Pollock (1990) refers to as *material probability*, which only describes what is actually the case. We are not concerned with this reading here.

[5] In their treatment of frequency adverbs, Åquist *et al* (1980) also use probabilities; however, they define conditional probabilities to be simple ratios, and consequently they cannot account for the puzzle of sentences such as (6). See Cohen (1995) for an interpretation of probability which is more appropriate for the analysis of the meaning of generics.

[6] This definition is not quite complete; it can be easily extended to handle the case of sets of alternative kinds, rather than properties, which are needed to account for sentences such as (2.e) and (2.d) (see Cohen 1996 for the details). However, since these cases do not give rise to default inferences (e.g. even if we know that (2.e) is true, and that Boris is a Bulgarian, we are still not justified in concluding, from these facts alone, that Boris is a good weightlifter), we will not deal with this issue further here.

properties of such alternatives, or how the alternatives can be determined in any given case.[7]

An important question is whether the set of alternatives is dependent on the particular language used; Blok and Eberle (1994) suggest that it is. As an example, they discuss the set of alternative kinds of beer, and claim that

> the alternatives *of which the native speaker is aware* in [a German sentence] are those kinds of beer that have a name in German, and correspondingly for [an English sentence]. Normally, neither *Kölsch* is an alternative of *lager*, nor is *ale* an alternative of *Pils* (Blok and Eberle 1994:240, my emphasis).[8]

It may very well be true that a hearer would normally be aware only of alternatives which have a name in his or her language, but it does not follow that these are, indeed, the alternatives with respect to which the sentence needs to be evaluated. Suppose that 80% of German beer drinkers drink Kölsch, 20% drink lager, and Germans drink no other beer. If, as Blok and Eberle (1994) claim, *Kölsch* were not an alternative of *lager*, (8) would be true, but I think it would clearly be false in the situation described:

(8) Germans drink lager.

I would, therefore, propose that alternatives are language independent;[9] they depend on the meaning of a sentence, on the context and world knowledge, but not on the language used. The question, then, is how the alternatives are determined in a given case.

Sometimes the solution is easy—the alternatives are explicitly stated. For example, if (9.b) is uttered in response to (9.a), the alternatives are overtly given:

(9) a. Do birds lay eggs, give birth to live young or undergo mitosis?
    b. Birds lay eggs.

In most cases, however, the alternatives have to be inferred somehow. There are, in fact, two different problems to be addressed: what the alternatives induced by a given sentence are, and which part of the sentence induces these alternatives. In other words, when evaluating a sentence like *Birds fly*, we may consider either alternative means of movement (*fly, walk, swim* etc.) or alternative classes of animals (*birds, mammals, fish* etc.). One question is whether *birds* or *fly* induces the alternatives; a second question is which alternatives are induced by a given property. We will address the second question first, a consideration of which will suggest an answer to the first question.

---

[7] But see Gabbay and Moravcsik (1978); Blok (1994); Blok and Eberle (1994).

[8] *Kölsch* and *Pils* are kinds of German beer. Interestingly, although the etymology of the English *lager* is German, the German word *lager* does not, in fact, name a kind of beer, and its meaning is *camp*.

[9] Or, more precisely, that they depend on a particular language only to the extent that the language affects the ontology of the language user, if at all.

2.3.1 Determinables and determinates

The problem of alternatives is similar to an issue which has been discussed by philosophers under the heading of the *determinable-determinate* relation. In its modern form this notion was introduced by Johnson (1921). He writes:

> I propose to call such terms as colour and shape *determinables* in relation to such terms as red and circular which will be called *determinates*... [A]ny one determinable such as colour is distinctly other than such a determinable as shape or tone; i.e. colour... is, metaphorically speaking, that from which the specific determinates, red, yellow, green, etc., emanate; while from shape emanate another completely different series of determinates such as triangular, square, octagonal, etc... Further, what have been assumed to be determinables—e.g. colour, pitch, etc.—are ultimately *different*, in the important sense that they cannot be subsumed under some one higher determinable, with the result that they are incomparable with one another; while it is the essential nature of determinates under any one determinable to be comparable with one another... [T]he ground for grouping determinates under one and the same determinable is not any partial agreement between them that could be revealed by analysis, but the unique and peculiar kind of difference that subsists between the several determinates under the same determinable, and which does not subsist between any one of them and an adjective under some other determinable (Johnson 1921:174–176, original emphases; page numbers are from the 1964 edition).

The idea that alternatives are determinates under a common determinable has some intuitive appeal. To use Johnson's terms, the alternatives seem to "emanate" from some common property; for example, laying eggs, giving birth to live young and undergoing mitosis all emanate from the property of procreating. These properties are different from each other,[10] but they are intuitively comparable; whereas they are not comparable to other properties such as flying. Unfortunately, Johnson's definition was not formal, and more recent definitions fare no better.[11] Most work subsequent to Johnson's (e.g. Searle 1959; Rosenberg 1966) concentrated on attempts to analyze the relation into more basic terms, and, in particular, to distinguish between the determinable-determinate relation on the one hand, and the genus-species relation on the other hand. Unfortunately, those attempts met with less than clear success, which may cause one to suspect (with Thomason 1969) that such an analysis would be fruitless and that an abstract, algebraic characterization is all that one can hope for. Let

[10] In fact, Johnson and others have claimed that determinates under the same determinable are mutually exclusive, but this claim has been questioned by e.g. Armstrong (1978).

[11] The determinable-determinate relation has also been mentioned in the linguistic literature (Lyons 1977), under the name of *quasi-hyponymy*, but although Lyons claims that the relation "can easily be made precise within the framework of a reasonably comprehensive transformational grammar of English" (Lyons 1977:299), neither he nor anyone else, to my knowledge, has ever attempted to do this.

us assume, then, that the relation is somehow given, as part of the language user's pragmatic knowledge. If p and q are properties, let $p \prec q$ indicate the fact that p is a determinate of q. Several properties of this relation seem desirable:[12]

- If $p \prec q$, then $p \rightarrow q$.
- Antireflexivity: for no p, $p \prec p$.
- Antisymmetry: for no p, q, $p \prec q$ & $q \prec p$.
- Transitivity: for all p, q, r, $p \prec q$ & $q \prec r \Longrightarrow p \prec r$. For example, **red**$\prec$ **color**, and, in turn, **color** $\prec$ **physical-property**. Then it seems correct to say that **red**$\prec$ **physical-property**.[13]

Since the relation is transitive, it makes sense to talk about a *minimal determinable* of a property:

Definition 2 *q is a minimal determinable of p, written $p \prec_m q$, iff $p \prec q$ and there is no r such that $p \prec r \prec q$.*

Note that a minimal determinable need not be unique, since not every two properties are in the $\prec$ relation. It is not clear that a minimal determinable exists for every property. If a property does have some determinable, it seems plausible that it has a minimal one, since an infinite descending chain of determinables does not seem to make any immediate sense for properties referred to in natural language—though judgment should be reserved, as there might be some pathological cases.

We wish to say that the alternatives induced by a property p are the properties which share a determinable with p. However, this will not do; *red* is a determinable of *scarlet*, yet, plausibly, we want *scarlet* to induce the set of alternatives consisting of all colors, not just shades of red. We will assume, then, that every determinate has a determinable which is somehow distinguished. A useful idea here is Searle's (1959) notion of *absolute determinables*. Searle notes that both *red* and *colored* are determinables of *scarlet*, yet *colored* seems to be more fundamental:

> The more fundamental position which "colour" occupies *vis à vis* both "red" and "scarlet" is shown by the fact that the predication of "red", "not

[12]It should be reiterated that these properties do not constitute a definition of the relation; indeed, one can conceive of other relations which share similar properties.

[13]It is not clear whether Johnson himself would have agreed with the characterization of the relation as transitive, but subsequent authors certainly have taken it to be so. Searle (1967), for example, considers

> color terminology as providing us with a hierarchy of terms many of which will stand in the determinable relation to each other as the specification of shades progresses from the less precise to the more precise (p. 358).

Be that as it may, the transitivity of the determinable-determinate relation turns out to be highly useful in determining the set of alternatives, as we will see shortly.

> red", "scarlet" or "not scarlet" of any object *presupposes* that "coloured" is true of the object. A term A presupposes a term B if and only if it is a necessary condition of A's being true *or false* of an object x, that B must be true of x. For example, as we commonly use these words, in order for it to be either true or false of something that it is red, it must be coloured. Both "red" and "scarlet" then presuppose their common determinable "coloured". But "scarlet" does not presuppose its determinable "red", and we may generalise this point as a criterion: B is an *absolute determinable* of A if and only if A is a determinate of B, and A presupposes B. Thus "coloured" is an absolute determinable of "red", but "red" is not an absolute determinable of scarlet [sic].
>
> ...
>
> The notion of an absolute determinable is relevant to the traditional problem of categories: every predicate carries with it the notion of a kind or category of entities of which it can be sensibly affirmed or denied. For example, "red" is sensibly affirmed or denied only of objects which are coloured—this is part of what is meant by saying that "red" presupposes "coloured". Absolute determinables then provide us with a set of category terms (Searle 1959:149–150, original emphases).

We can combine the notions of absolute and minimal determinables, and talk about a minimal absolute determinable of a property:

Definition 3 *q is a minimal absolute determinable of p, written $p \prec_A q$, iff q is an absolute determinable of p, and there is no $r \prec q$ such that r is also an absolute determinable of p.*

We can now use the notion of a minimal absolute determinable to define the alternatives induced by a property p as the determinates under a minimal absolute determinable of p. In other words, all the alternatives to p presuppose the same minimal absolute determinable of p:

Definition 4 *Let p and q be properties such that $p \prec_A q$. Then a set of alternatives to p (given q) is $ALT(p) =_{\text{def}} \{p' | p' \prec_A q\}$*

Note that the definition refers to *a* set of alternatives, rather than *the* set, because a property may have more than one minimal absolute determinable, and, consequently, may induce more than one set of alternatives.

In his definition of absolute determinables, Searle uses semantic presupposition, i.e. a definition which is not dependent on context. This, however, seems too strong. Alternatives are heavily influenced by context, as can be seen by example (9) above. As an additional example, consider the context of a discussion of different forms of movement of animals. In this context, the alternatives under consideration would presuppose movement, i.e. {*walk, fly, swim, ride unicycles,...*}. Sentence (10), then, would be false, since the majority of bears do not ride unicycles:

(10) Bears ride unicycles.

However, in the context of a discussion of the acts performed in the Great Russian Circus, the alternatives would presuppose some circus act, i.e. {*juggle, walk the tightrope, ride unicycles,...*}. Given this context, (10) would be true, since, presumably, the majority of bears which perform in the Great Russian Circus ride unicycles.

We will, then, use a pragmatic definition of presupposition, according to which the property p presupposes the property q iff a felicitous application of p to an individual i requires that the proposition $q(i)$ be in the common ground.[14] Thus the common ground determines the absolute determinables in a given context.[15]

Suppose p and q are properties, and ALT(p) and ALT(q) are given. What about ALT($p \wedge q$), ALT($p \vee q$), and ALT($\neg p$)?

Let us look at negated properties first. Consider (11):

(11) Mammals don't bear live young.

The majority of mammals (including males, mammals which are too young or too old to procreate, etc.) do *not* bear live young, and yet (11) is false. This suggests that (11) is evaluated with respect to the same set of alternatives as (2.c), i.e. different forms of procreation. Since it is false that the majority of procreating mammals don't bear live young, (11) is, indeed, false. I propose that, in general, for a property p, ALT($\neg p$) =ALT(p).

What, then, about conjunctions and disjunctions?

There are three intuitively plausible candidates for the alternatives induced by $p \wedge q$:
$\mathrm{ALT}(p \wedge q) =_{\mathrm{def}}$

1. $\{p \wedge q' | q' \in \mathrm{ALT}(q)\}$
2. $\{p' \wedge q | p' \in \mathrm{ALT}(p)\}$
3. $\{p' \wedge q' | p' \in \mathrm{ALT}(p) \;\&\; q' \in \mathrm{ALT}(q)\}$

It would seem that all three are legitimate; they simply represent an ambiguity, to be resolved by, perhaps, the focus of the sentence.[16] Consider the sentences in (12):[17]

(12) a. Piranhas live in FRESHWATER tanks.

b. Piranhas live in freshwater TANKS.

[14] Of course, as with other definitions of pragmatic presupposition, the notions of *felicity* and *common ground* will be left vague here.

[15] cf. Blok and Eberle's (1994) discussion of how context can "frame" nodes in an ontology—in our terms, how context helps to determine absolute determinables.

[16] For more on the role of focus see section 2.3.3 below.

[17] Here, and henceforth, small capitals indicate a focused expression.

c. Piranhas live in FRESHWATER TANKS.

Sentence (12.a) induces alternative types of tanks; presumably, freshwater tanks and saltwater tanks. Since the majority of piranhas which live in tanks, live in freshwater tanks, (12.a) is true. The alternatives under consideration in (12.b), on the other hand, seem to be various kinds of bodies of freshwater: tanks, rivers, lakes, ponds. Since the majority of piranhas live in rivers, and not in tanks, (12.b) is false. Finally, (12.c) induces a Cartesian product of various types of water and various bodies of water. Since, again, it is not true that the majority of piranhas live in freshwater tanks, (12.c) is false.

It should be noted that there is an additional level of ambiguity, as a conjunction (or disjunction) may be interpreted as conjoining sentences rather than properties, the former being strongly preferred. Consider (13), for example:

(13) Peacocks have a magnificent blue tail and lay whitish eggs.

Taken as a conjunctive property, (13) is false, since no peacock both has a magnificent tail and lays whitish eggs. This ambiguity is absent, or almost absent, in the admittedly awkward (14), which is not readily interpretable as sentential conjunction:

(14) Peacocks are blue-tailed egg-layers.

Sentence (14), therefore, is unambiguously false.

An application of de Morgan's laws will yield the following three options for the alternatives of a disjunction:
$\mathrm{ALT}(p \vee q) =_{\mathrm{def}}$

1. $\{(\neg p) \wedge q' | q' \in \mathrm{ALT}(q)\}$
2. $\{p' \wedge (\neg q) | p' \in \mathrm{ALT}(p)\}$
3. $\{p' \wedge q' | p' \in \mathrm{ALT}(p)\ \&\ q' \in \mathrm{ALT}(q)\}$

Note that the alternatives of a disjunction are in the form of conjunctions, not disjunctions. This may seem counterintuitive at first, but is, in fact, quite desirable. Consider (15):

(15) a. Sinners repent or GO TO HELL.

b. Sinners REPENT or go to hell.

c. Sinners REPENT OR GO TO HELL.

Sentence (15.a) is evaluated with respect to sinners who do not repent; it claims that among all their alternative destinations, hell is the likeliest. Sentence (15.a), on the other hand, is evaluated with respect to sinners who are not sent to hell; it states that among all the alternatives reasons for that, the likeliest is that they repented. Sentence (15.c) is evaluated with respect to all sinners who are sent somewhere after their death and who have done something to cause their destination; the majority of those, according to (15.c), either repent or go to hell. Clearly, some of these sentences might be judged true and others false, depending, of course, on one's religious beliefs.

An important special case is where $\mathrm{ALT}(p) = \mathrm{ALT}(q)$. We have said above that there are three legitimate options for the set of alternatives to $p \wedge q$, and three options for the set of alternatives to $p \vee q$. The Cartesian product option is identical for both:

$$\{p' \wedge q' | p' \in \mathrm{ALT}(p) \;\&\; q' \in \mathrm{ALT}(q)\}.$$

Now, if $\mathrm{ALT}(p) = \mathrm{ALT}(q)$,

$$\bigvee \{p' \wedge q' | p' \in \mathrm{ALT}(p) \;\&\; q' \in \mathrm{ALT}(q)\} = \bigvee \mathrm{ALT}(p) = \bigvee \mathrm{ALT}(q).$$

In words, the disjunction of the alternatives to the conjunction (or disjunction) will be the same as the disjunction of the alternatives to either conjunct (or disjunct), though the actual set of alternatives may differ. Thus, for example, (16.a) will plausibly be evaluated with respect to alternative numbers of wheels, and (16.b)—with respect to alternative languages:

(16) a. Motor vehicles have two or four wheels.

b. American ambassadors to France speak English and French.

Let us now turn to complex properties, exemplified by (17):

(17) Healthy people eat apples.

What is the set of alternatives which (17) is evaluated with respect to? It should be the set induced by the complex property *eat apples*, but how can this set be deduced?

Blok (1994) suggests that the alternatives to *eat* are {*drink, chew,...*} and the alternatives to *apple* are {*banana, lime,...*}. He then claims that the set of alternatives to *eat apples* does not include *drink a banana*, and concludes that the problem of determining the set of alternatives of a property "seems unsolvable...from a logical or linguistic point of view" (Blok 1994:8). It is not at all clear to me what makes him reach this pessimistic conclusion. Indeed, drinking a banana may be just too bizarre a property to be normally considered; but *chew a lime* is certainly an acceptable alternative. Blok notes that the relevant alternatives to *eat apples* "may be *work* or *kiss a woman*, who knows" (Blok 1994:8). In that I believe he is right—*eat apples* may, on its own, be a determinate of some determinable such as *actions performed in the morning* or *things which John likes to do*, etc. This will be determined by the context. Surely, however, there are contexts where decomposing *eat apples* into the alternatives of *eat* and *apples* is preferred; it certainly seems to be the best we can do in the null context.

In such cases, again, there are three plausible sets of alternatives. One is the set of alternative foods, the second is the set of alternative actions performed with apples, and the third is the Cartesian product of both, i.e. the set composed of various foods and various actions performed with them. Assuming that the meaning of *eat apples* is

$$\lambda x.\exists y.\mathbf{apples}(y) \wedge \mathbf{eat}(x, y),$$

the three possible values for its set of alternatives would be:

1. $\{\lambda x.\exists y : (P(y) \wedge \mathbf{eat}(x, y))|P \in \mathrm{ALT}(\mathbf{apples})\}$
2. $\{\lambda x.\exists y : (\mathbf{apples}(y) \wedge R(x, y))|R \in \mathrm{ALT}(\mathbf{eat})\}$
3. $\{\lambda x.\exists y : (P(y) \wedge R(x, y))|P \in \mathrm{ALT}(\mathbf{apples}) \ \&\ R \in \mathrm{ALT}(\mathbf{eat})\}$

2.3.2 Presupposition

Since sets of alternatives share an absolute determinable, and, hence, a presupposition, it is reasonable to expect that a presupposition of the property inducing the alternatives would be shared by the alternatives.[18] For example, the verb *manage* is an *implicative* verb (Karttunen 1971); roughly put, saying that x managed p, presupposes that x attempted to accomplish p. The alternatives induced by the property *manage p*, then, would plausibly be possible outcomes of the attempt to accomplish p—success or failure.

Now consider (18):

(18) People manage to survive a week without food.

Sentence (18) is true if people are more likely to survive a week without food than to die in such circumstances. Note that many people do not survive a week without food, for the simple reason that they are never put in this predicament. In other words, these people do not satisfy the presupposition, hence they do not satisfy any of the alternatives (since those entail the presupposition) and, therefore, do not affect the truth or falsity of the sentence.

2.3.3 Focus

An idea common to many theories of focus is that the focused element is "new", whereas the unfocused part of the sentence is "old" or "presupposed."[19] Since the alternatives share a presupposition, it is reasonable to expect that what they would differ on would be the focused element. In other words, the focus would be predicted to induce the set of alternatives.

[18] cf. Schubert and Pelletier's (1987) claim that presuppositions affect the "ensembles of cases" which a generic sentence quantifies over.

[19] This is a rather crude and inaccurate way of putting the issue, but it will do for our purposes here. For more refined views on this aspect of focus, as well as for an overview of work on this topic, see Partee (1991); Vallduví (1992).

This prediction is, in fact, borne out. Rooth (1985) has suggested that the focus structure of a sentence indicates a set of alternatives under consideration. Consider (19), for example:

(19) a. Who did John introduce Bill to?

b. John introduced Bill to SUE.

c. John introduced BILL to Sue.

With respect to (19), Rooth suggests that

> a question introduces a set of alternatives into a discourse; the alternatives introduced by [(19.a)] are propositions of the form introduce'(j,b,y), where y is some individual.[20] The function of the focus in the answer [(19.b)], I suggest, is to signal that alternatives of this form are indeed under consideration... [(19.c)] as a reply to [(19.a)] would incorrectly suggest that alternatives of the form introduce'(j,y,s) are under consideration (Rooth 1985:13).

The truth conditions of (19.b) and (19.c) are the same, although their focus structures are different. Rooth observes, however, that there are cases where different focus structures result in different truth conditions. For example, "suppose I introduced Bill and Tom to Sue, and performed no other introductions. Then [(20.a)] is false and [(20.b)] is true" (Rooth 1985:2–3):

(20) a. I only introduced BILL to Sue.

b. I only introduced Bill to SUE.

I suggest that generics are another case where focus contributes to truth conditions. The set of alternatives with respect to which a sentence is evaluated would be induced by the focused element in the sentence. For example:

(21) a. Criminals are executed in accordance with the LAW.

b. Criminals are EXECUTED in accordance with the law.

Sentence (21.a) is true, whereas (21.b) is false. The reason is that (21.a) and (21.b) induce different sets of alternatives. Sentence (21.a) would be evaluated with respect to alternative factors potentially determining a person's execution: the law, the whims of the judge, the weather etc. Since criminals

[20] More recently, Rooth (1992) suggests that deriving the set of alternatives may be more complicated than simply replacing the focused element with a variable. He proposes

> an interpretation principle which introduces a variable, thought of as a contrasting element or set of contrasting elements. This variable can be anaphoric to a variety of pragmatic and semantic objects, resulting in a variety of focus-sensitive effects, including both discourse effects and sentence-internal association with focus effects (Rooth 1992:113).

This corresponds rather well to the theory of generics presented here, where alternatives are assumed to be derived using both semantic and pragmatic processes.

are more likely to be executed according to the law than as a consequence of some other factor (in democratic countries, at any rate), (21.a) is true. Sentence (21.b), on the other hand, would be evaluated with respect to alternative punishments criminals may be subjected to: execution, jail, fine, 30 lashes etc. Since an arbitrary criminal is rather unlikely to be punished by death, (21.b) is false.

2.3.4 Syntactic constructions

Certain syntactic constructions seem to influence the sets of alternatives induced. Consider Chomsky's (1975) well known example:

(22) a. Beavers build dams.

b. Dams are built by beavers.

It seems that (22.a) is true, whereas (22.b) is false. When the sentences in (22) are uttered with normal intonation, the subject is the topic, or, to put it roughly, the old, presupposed information, and the other noun phrase is the focus. Hence different constituents are focused in (22.a) and (22.b), and, consequently, different alternatives are induced. Note that the above is not dependent on any particular fact about the peculiarity of the passive construction, but rather on the topic-focus structure it induces. Other constructions which also influence the topic-focus structure have similar effects:

(23) a. As for beavers, they build dams.

b. As for dams, beavers build them.

(24) a. It is dams which beavers build.

b. It is beavers which build dams.

Plausibly, (22.a), (23.a), and (24.a) would be evaluated with respect to alternative constructions beavers might build. Since beavers are more likely to build dams than other artificial constructions, (22.a), (23.a) and (24.a), are all true. Sentences (22.b), (23.b), and (24.b), on the other hand, introduce alternative dam builders. Since dams are more likely to be built by humans than by beavers,[21] (22.b), (23.b), and (24.b) are all false.

[21] I am not sure this is entirely accurate. Perhaps if we count all dams, including man-made and beaver-made ones, we will find that most of them are built by beavers. If that is correct, then (22.b) is judged false by mistake. Still, the mistake is understandable, since it is certainly true that the vast majority of dams I have personally encountered or heard about were built by humans, and not by beavers.

2.4 Refutation statements

Consider the following mini-dialogues:[22]

(25) **A:** Nobody in India eats beef.

B: That's not true! Indians do eat beef.

(26) **A:** Many husbands beat their wives.

B: Well, women beat their husbands too!

B's statements in sentences (25) and (26) are made as refutations; they are meant to refute the overt claim that nobody in India eats beef and the implied claim that no woman beats her husband, respectively.[23] Refutation statements seem to constitute counterexamples to the theory proposed here, since they are interpreted existentially, rather than generically: for B's response in (25) to be true, it is sufficient that *some* Indians eat beef, not that Indians, in general, do so; similarly, for B's response in (26) to be true, only *some* women have to beat their husbands, and not necessarily women in general.

It is possible, nonetheless, to account for these sentences using the theory of generics proposed in this paper. Suppose the set of alternatives for both are singletons: {*eat beef*} and {*beat one's husband*}, respectively. Then, when evaluating (25), we would be comparing the probability of an Indian to eat beef to his or her probability to eat beef and not eat beef—which is zero. Similarly, when evaluating (26), we would compare the probability of a woman to beat her husband to her probability to beat her husband and not beat him—which is, again, zero. Formally:

1. $P(\textbf{eat-beef}|C(\textbf{Indian})) >$
   $P((\neg\textbf{eat-beef}) \wedge \textbf{eat-beef}|C(\textbf{Indian})) = 0$

2. $P(\textbf{beat-husband}|C(\textbf{woman})) >$
   $P((\neg\textbf{beat-husband}) \wedge \textbf{beat-husband}|C(\textbf{woman})) = 0$

Thus (25) and (26) would be true just in case there is a nonzero probability for women to beat their husbands and for Indians to eat beef, respectively—which is what brings about the existential readings.

What sort of evidence do we have that these are, indeed, the appropriate sets of alternatives? If, as claimed in section 2.3.3, the focused constituent is the one inducing the alternatives, focusing an element should change the truth conditions of a sentence. This does, indeed, seem to be the case:

(27) a. It is beef which Indians eat.

[22] Examples (25) and (26) were suggested to me by Clark Glymour and Lori Levin, respectively.

[23] Of course, B's statement in (26) does not refute A's *overt* claim about husbands beating their wives.

b. It is eating beef which Indians engage in.

c. It is their husbands whom women beat.

d. It is beating their husbands which women engage in.

In a situation where only few Indians eat beef, and few women beat their husbands, the sentences in (27) would be false.[24]

3 Reasoning with generics

A default reasoning system contains rules which determine the conclusions to be drawn in case of insufficient knowledge. Various methods have been proposed to formalize and draw conclusions using such rules; I will refer to all such mechanisms, rather loosely, as *default rules*. A default rule will be written as $\alpha \rightharpoonup \beta$, where α and β are properties.

The same rule will be represented differently in different nonmonotonic formalisms. Consider, for example, **bird** $\rightharpoonup$ **fly**:

Default Logic (Reiter 1980). If x is a bird, and it is consistent to assume that x flies, assume so:

$$\frac{\mathbf{bird}(x) : \mathbf{fly}(x)}{\mathbf{fly}(x)}$$

Circumscription (McCarthy 1980). For all x, if x is a normal bird, assume x flies:

$$\forall x.\mathbf{bird}(x) \land \neg\mathbf{abnormal}(x) \rightarrow \mathbf{fly}(x)$$

Autoepistemic Logic (Moore 1985). If x is a bird, and the agent does not believe it does not fly, assume it flies:

$$\mathbf{bird}(x) \land \neg L \neg\mathbf{fly}(x) \rightarrow \mathbf{fly}(x)$$

In the literature concerning nonmonotonic logics, default rules are usually described informally using generic sentences, such as *Birds fly* and the like. Presumably, a rule like **bird** $\rightharpoonup$ **fly** seems intuitively a desirable rule to have, because birds do, indeed, fly. I propose to take this informal characterization seriously. That is to say, I propose that a default rule be *adequate* just in case the corresponding generic sentence is true. The rule I propose, then, is as follows:

[24] Sentence (27.c) may still be true if, whenever women beat somebody, it is mostly their husbands. But if it turns out that women are more likely to beat, say, their dogs, (27.c) would be false, yet (26) would still be true.

Definition 5 (Adequacy) *Let α and β be properties, and A a set of properties. Then a default rule $\alpha \rightharpoonup \beta$ will be adequate with respect to A iff $\mathbf{gen}_A(\alpha)(\beta)$ is true, i.e. iff*

$$P(\beta|\alpha \wedge (\bigvee A)) > 0.5$$

It will be instructive to see how our definition of adequacy relates to an actual nonmonotonic system. The system I have chosen is Reiter's (1980) Default Logic, but this is done simply for the sake of convenience, and it should be possible to suggest a similar treatment using a different formalism. Default Logic employs rules of the following form:

$$\frac{\kappa(x) : \beta_1(x), \ldots, \beta_m(x)}{\alpha(x)}$$

The intuitive meaning of this rule is the following: if, for some instantiation of x, $\kappa(x)$ is derivable, and none of $\neg\beta_1(x), \ldots, \neg\beta_m(x)$ are derivable, then derive $\alpha(x)$.[25] An adequate Default Logic rule will then be defined as follows:

Definition 6 (Adequate Default Logic rule)

$$\frac{\kappa(x) \wedge (\alpha_1(x) \vee \alpha_2(x) \vee \cdots \vee \alpha_n(x)) : \beta_1(x), \ldots, \beta_m(x)}{\alpha(x)}$$

is an adequate default rule iff

$$P(\alpha|\kappa \wedge (\alpha_1 \vee \alpha_2 \vee \cdots \vee \alpha_n)) > 0.5$$

3.1 Specificity and relevance

When default rules, even adequate ones, are applied to a given instance, a problem often arises: two or more rules are applicable, but the statements derived using them are incompatible. It is widely agreed that, in general, a more specific rule should have precedence over a less specific one. For example, let $A = \{\mathbf{bear\text{-}live\text{-}young}(x), \mathbf{lay\text{-}eggs}(x), \mathbf{mitosis}(x)\}$, and suppose we have the following two rules:

1. $$\frac{\mathbf{mammal}(x) \wedge \bigvee A : \mathbf{bear\text{-}live\text{-}young}(x)}{\mathbf{bear\text{-}live\text{-}young}(x)}$$

2. $$\frac{\mathbf{platypus}(x) \wedge \bigvee A : \mathbf{lay\text{-}eggs}(x)}{\mathbf{lay\text{-}eggs}(x)}$$

[25] Usually, on the right side of the colon there is only one formula, which is equivalent to $\alpha(x)$. In this case, the rule is said to be *normal*. The definition I will propose for adequate Default Logic rules would work just as well for both normal and non-normal rules, though in the examples to follow I will use normal rules.

Now if Pat is a platypus, she is also a mammal. Depending on the order in which the rules apply, we may conclude either that Pat gives birth to live young or that she lays eggs. But since every platypus is a mammal, rule 2 is more specific, and therefore should have precedence, i.e. we should conclude that Pat lays eggs.

The specificity constraint is particularly important in our system for the following reason. Suppose we have two adequate default rules:

$$\frac{\kappa(x) \wedge (\alpha_1(x) \vee \alpha_2(x) \vee \cdots \vee \alpha_n(x)) : \beta(x)}{\alpha(x)}$$

$$\frac{\kappa(x) \wedge (\alpha_1(x) \vee \alpha_2(x) \vee \cdots \alpha_n(x) \vee \alpha_{n+1}(x)) : \beta'(x)}{\alpha'(x)}$$

If $\alpha_1(x) \vee \alpha_2(x) \vee \cdots \vee \alpha_n(x)$ is derivable, then surely $\alpha_1(x) \vee \alpha_2(x) \vee \cdots \vee \alpha_n(x) \vee \alpha_{n+1}(x)$ is derivable too, since the latter is entailed by the former. Yet the specificity constraint will ensure that the more specific rule will defeat the less specific one, so that $\alpha(x)$, and not $\alpha'(x)$, will be concluded.

The specificity constraint is not part of Default Logic (or most other non-monotonic formalisms). If one wishes to enforce it, one needs to encode its consequences for each rule separately (Etherington and Reiter 1983). One simple way to ensure globally that more specific rules supersede less specific ones would be to require that inferences be made based on *all* available information. Returning to Pat, we know that Pat is a platypus, and also that she is a mammal, and we are wondering what her form of offspring is. We want a rule, then, whose antecedent is

$$\mathbf{platypus}(x) \wedge \mathbf{mammal}(x) \wedge \bigvee A.$$

Since all platypuses are mammals,

$$\mathbf{platypus}(x) \wedge \mathbf{mammal}(x) \leftrightarrow \mathbf{platypus}(x).$$

Rule 2, then, but not rule 1, has the desired antecedent, and only it will be applied; hence we will conclude that Pat lays eggs, as desired.

Requiring inference to be made from all available information will also provide a simple way to handle cases of conflicting defaults, exemplified by the *Nixon diamond.* Assume that we have two rules, one stating that Quakers are pacifists, and the other—that Republicans are not pacifists. Suppose we learn that Nixon is both a Quaker and a Republican. What can be said about whether or not Nixon is a pacifist? Since our information about Nixon contains the facts about his being both a Quaker and a Republican, we need a rule concerning Republican Quakers. If we have such an adequate rule, we can use it to draw a conclusion about Nixon; but if we do not, we can draw no conclusion. In other words, the theory proposed here adopts a *skeptical* approach to such cases.

There is a problem with the demand that inference be made from all available information, though. Suppose we also know that a certain mammal has a cold, and we are wondering about this mammal's form of offspring. We would, then, need a rule whose antecedent is

$$\mathbf{mammal}(x) \wedge \mathbf{has\text{-}a\text{-}cold}(x) \wedge \bigvee A.$$

We do not, however, have such a rule; nor is

$$\mathbf{mammal}(x) \wedge \mathbf{have\text{-}a\text{-}cold}(x)$$

equivalent to

$$\mathbf{mammal}(x).$$

It seems that some information need not be considered; it is irrelevant. Intuitively, whether or not a mammal has a cold should have little effect on its form of offspring, and the rule stating that mammals give birth to live young should apply all the same. In other words, the rule

$$\frac{\mathbf{mammal}(x) \wedge \mathbf{has\text{-}a\text{-}cold}(x) \wedge \bigvee A : \mathbf{bear\text{-}live\text{-}young}(x)}{\mathbf{bear\text{-}live\text{-}young}(x)}$$

should be adequate iff

$$\frac{\mathbf{mammal}(x) \wedge \bigvee A : \mathbf{bear\text{-}live\text{-}young}(x)}{\mathbf{bear\text{-}live\text{-}young}(x)}$$

is adequate. This would be true just in case

$$P(\mathbf{bear\text{-}live\text{-}young} | \mathbf{mammal} \wedge \mathbf{have\text{-}a\text{-}cold} \wedge \bigvee A) = P(\mathbf{bear\text{-}live\text{-}young} | \mathbf{mammal} \wedge \bigvee A).$$

This, in turn, would be true if **bear-live-young** and **have-a-cold** are conditionally independent given $\mathbf{mammal} \wedge \bigvee A$, according to the usual definition of conditional independence:

Definition 7 (Conditional independence) *Properties α and β are conditionally independent given γ iff*

$$P(\alpha \wedge \beta | \gamma) = P(\alpha|\gamma) P(\beta|\gamma)$$

This will give us the desired results.

It should be emphasized that such assumptions of independence are not specified in our knowledge base; they are assumptions made in order to be able to make inferences. Several ways to arrive at independence statements have been proposed in the literature. Shastri (1989) proposes a way to infer when two properties are irrelevant; Bacchus (1990) suggests a system where such assumptions are made explicitly and nonmonotonically, as a special case of default inference; and in more recent work (Bacchus *et al* 1993;1994), assumptions of independence follow from general properties of the system.

3.2 Acceptance rules

The system of adequate default reasoning proposed above can be used to define what philosophers of science refer to as an *acceptance rule*. An acceptance rule is a rule which determines when beliefs are to be accepted. Thus, for example, a good acceptance rule should let us accept the belief that an arbitrary bird flies, but not that it drives a truck.[26] The problem of evaluating acceptance rules is a difficult one, and numerous acceptance rules have been proposed in the literature.[27]

Definition 5 above cannot be immediately applied, since it deals with *properties*, whereas an acceptance rule determines when a *proposition* is to be accepted. Therefore,we should be careful about what we mean when we talk about the probability $P(A|B)$. If A and B are properties, such as being a bird or flying, as has been heretofore assumed, then we are discussing *indefinite* probabilities. However, when we want to discuss the probabilities of statements, such as Tweety's being a bird or flying, A and B are propositions, and we are discussing what has been called *definite* probabilities.[28]

The problem of inferring definite probabilities from indefinite ones has been called the problem of *direct inference*, and it is not an easy one. There is a considerable body of work on the matter,[29] primarily focusing on the problem of identifying the relevant information, mentioned in the previous section. A discussion of the issue is, however, beyond the scope of this paper.

Taking the probability measure P to represent definite probabilities, we can define our acceptance rule as follows:

Definition 8 *Let κ be a set of propositions representing the background knowledge, i.e. the propositions considered to be true, and let*

$$A = \{\alpha_1, \alpha_2, \ldots, \alpha_n\}$$

be a set of propositions, and α a proposition. Then α is accepted from κ with respect to A, written $\kappa \overset{A}{\rhd} \alpha$, iff $P(\alpha|\kappa \wedge \bigvee A) > 0.5$

Note that if the alternatives are exhaustive, the acceptance rule reduces to simple majority:

Proposition 1 *If $\kappa \models \bigvee A$, then α is accepted iff $P(\alpha|\kappa) > 0.5$.*

Proof. *If $\kappa \models \bigvee A$ then $P(\alpha|\kappa \wedge \bigvee A) = P(\alpha|\kappa)$.* ■

[26] Of course, accepting a belief does not necessarily mean that it is true, or that the reasoner is certain that it is true; rather it means that the reasoner thinks it is *plausible*, and accepts the belief as true, pending additional information to the contrary.

[27] See Kyburg (1970) for a good overview.

[28] I am using Pollock's (1990) terminology in distinguishing between definite and indefinite probabilities.

[29] See Bacchus *et al* (1994) for a brief overview.

The acceptance rule proposed here may be seen as combining two previous suggestions: Levi (1967) and Kyburg (1961). Levi suggests that a statement is believed with respect to a set of *potential answers.* Given a different set of potential answers, different statements will be accepted. The device of potential answers is similar to our alternatives, but not identical. According to Levi, the potential answers form a partition, i.e. κ entails that exactly one of the potential answers is true. This contrasts with our approach, where more than one alternative, or none at all, may be true. His acceptance rule is more complicated than ours, and the accepted statement is a disjunction of one or more *likely* statements among the potential answers.[30] Kyburg (1961) suggests a high probability rule, i.e. a statement is accepted, according to him, iff it has a probability greater than some parameter r. This is similar to our approach, if we take $r = 0.5$. However, Kyburg considers a statement to be accepted in either all circumstances or none; there is no room in his theory for a set of alternatives under consideration.

A well known test for a proposed acceptance rule is how well it deals with the *lottery paradox* (Kyburg 1961). Imagine a fair lottery with one million tickets, and suppose the rules of the lottery are such that exactly one ticket will win. Will ticket #1 win? Not very likely—the odds are a million to 1 against its winning. It is therefore reasonable to accept the proposition that ticket #1 will lose. Similarly, it is reasonable to believe that ticket #2 will lose, and so on, for each one of the one million tickets. We are assured that one ticket will, in fact, win, so we accept this statement too. But the set of a million and one statements which we have just accepted is inconsistent!

Levi's system proposes a solution; if the question we are interested in is whether ticket #i will win or lose, we will conclude that ticket #i loses; if we are interested to know whether *any* ticket will win, we will believe that some ticket will, in fact, win. If we ask which ticket will win, Levi's rule will lead us to accept only the following disjunction: *ticket #1 will win or ticket #2 will win or ... or ticket #1,000,000 will win.* In Levi's system, we cannot even consider the question which gives rise to the paradox, because the potential answers are not mutually exclusive, i.e. it is possible (indeed, necessary) for one ticket to win and for that ticket to be ticket #i, for some i. This solution may seem to make sense in the case of the lottery paradox; however, there are many cases where the potential answers are not mutually exclusive. We cannot, for example, ask (28), since John may be multilingual:

(28) Which language does John speak?

Hence we cannot accept the plausible belief that John speaks English on the basis that John is an American, and that the vast majority of Americans speak

[30] I will not go into the details of the criteria for likelihood here. It should be noted, though, that in general Levi's approach will accept fewer statements than ours. The reason is that one of Levi's goals is to ensure that the set of accepted statements is closed under conjunction, an issue to which I will return below.

English.

Kyburg has a different solution. His rule will accept all million and one statements about the lottery. However, he does not require that the set of accepted statements be deductively closed; in particular, if p and q are accepted, $p \wedge q$ is not necessarily accepted. Hence the set of accepted statements will, indeed, be inconsistent, but harmlessly so. That is to say, it will not be possible to conclude all sentences of the language from it. Instead of closure under conjunction, Kyburg (1970) suggests the following two constraints:[31]

The Weak Deduction Principle: If $\kappa \overset{A}{\rhd} \alpha$ and $\models \alpha \rightarrow \beta$ then $\kappa \overset{A}{\rhd} \beta$.

The Weak Consistency Principle: There is no statement α such that $\kappa \overset{A}{\rhd} \alpha$ and for every statement β, $\alpha \rightarrow \beta$.

What about our acceptance rule? It seems to enjoy the best of both worlds. If the set of alternatives is {*ticket #i wins, ticket #i loses*}, we will conclude that ticket #i loses; if the set of alternatives is {*some ticket wins, no ticket wins*}, we will conclude that some ticket wins. But suppose that the set of alternatives is the one which gives rise to the paradox, namely

$$A = \{\textit{some ticket will win, no ticket will win}\} \cup \{\textit{ticket \#i will win} \text{---} 1 \leq i \leq 1,000,000\} \cup \{\textit{ticket \#i will lose} \text{---} 1 \leq i \leq 1,000,000\}.$$

In that case, all statements will be accepted, but no harm will be done; our system, like Kyburg's, is not closed under conjunction. It does, however, satisfy the following constraints:

Proposition 2

1. *If* $\kappa \overset{A}{\rhd} \alpha$ *and* $\kappa \wedge \bigvee A \models \alpha \rightarrow \beta$, *then* $\kappa \overset{A}{\rhd} \beta$.

2. *There is no* α *s.t.* $\kappa \overset{A}{\rhd} \alpha$ *and* $\kappa \overset{A}{\rhd} \neg\alpha$.

Proof. 1. *If* $\kappa \wedge \bigvee A \models \alpha \rightarrow \beta$, *then* $P(\beta|\kappa \wedge \bigvee A) \geq P(\alpha|\kappa \wedge \bigvee A) > 0.5$

2. *Suppose* $K \overset{A}{\rhd} \alpha$ *and* $K \overset{A}{\rhd} \neg\alpha$. *Then* $P(\alpha|\kappa \wedge \bigvee A) > 0.5$ *and* $P(\neg\alpha|\kappa \wedge \bigvee A) > 0.5$, *which is impossible.*

■

Note that the constraints above are stronger than Kyburg's. The Weak Deduction principle only applies when $\alpha \rightarrow \beta$ is a tautology, whereas our constraint (1) applies whenever $\alpha \rightarrow \beta$ is entailed by $\kappa \wedge \bigvee A$. The Weak Consistency principle merely requires that it be impossible to accept all statements in the language, whereas our constraint (2) forbids the acceptance of any statement and its negation.

[31] I am using my own notation rather than Kyburg's.

3.3 Properties of adequate nonmonotonic reasoning

In his work concerning the semantics of nonmonotonic reasoning, Pearl (1988) suggests several properties that any desirable nonmonotonic system needs to satisfy. Two of those, our system satisfies as it stands:[32]

Proposition 3

1. *Propositions which are entailed by the background knowledge are accepted: if* $\kappa \models \alpha$ *and* $P(\kappa \wedge \bigvee A) > 0$ *then* $\kappa \overset{A}{\rhd} \alpha$.
2. *If* $\kappa \wedge \beta \overset{A}{\rhd} \alpha$ *and* $\kappa \wedge \neg\beta \overset{A}{\rhd} \alpha$ *then* $\kappa \overset{A}{\rhd} \alpha$.

Proof. 1. *If* $\kappa \models \alpha$ *and* $P(\kappa \wedge \bigvee A) > 0$ *then* $P(\alpha|\kappa \wedge \bigvee A) = 1 > 0.5$

2. $P(\alpha|\kappa \wedge \bigvee A) \ = P(\alpha|\kappa \wedge \beta \wedge \bigvee A) \times P(\beta|\kappa \wedge \underline{i}gveeA) +$
$P(\alpha|\kappa \wedge (\neg\beta) \wedge \bigvee A) \times P((\neg\beta)|\kappa \wedge \bigvee A) \ >$
$0.5 \times (P(\beta|\kappa \wedge \bigvee A) + P(\neg\beta|\kappa \wedge \bigvee A)) \ = 0.5$

■

The first property is clearly highly desirable, and the second aims at capturing the intuition that if big mammals generally bear live young and small mammals generally bear live young, then mammals generally bear live young.

Pearl proposes two additional properties:

1. If $\kappa \overset{A}{\rhd} \beta$ and $\kappa \overset{A}{\rhd} \alpha$ then $\kappa \wedge \beta \overset{A}{\rhd} \alpha$.
2. If $\kappa \overset{A}{\rhd} \beta$ and $\kappa \wedge \beta \overset{A}{\rhd} \alpha$ then $\kappa \overset{A}{\rhd} \alpha$

Pearl claims that common-sense reasoning seems to agree with these rules, whereas theories such as the one proposed here (which he refers to as "majority" logics) do not. However, in a later paper he writes:

> I have speculated that this agreement is more in line with the rules of [Pearl's theory] than with those of "support" or "majority" logics. I am now in the opinion that this agreement is more reflective of tacit assumptions of independence (Pearl 1991:180n).

Let us, then, make this tacit assumption of independence explicit. What is needed here is a straightforward extension of conditional independence:

Definition 9 (κ-independence) 1. *Let* α, β *and* κ *be formulas, and let* A *and* B *be sets of formulas. Then* α *and* β *are* κ*-independent given* $\langle A, B\rangle$ *iff*

$$P(\alpha \wedge \beta|\kappa \wedge \bigvee(A \wedge B)) = P(\alpha|\kappa \wedge (\bigvee A))P(\beta|\kappa \wedge (\bigvee B))$$

[32] Again, I am using my own notation rather than Pearl's.

2. *Let A and B be sets of formulas. Then A and B are κ-independent iff for every $\alpha \in A, \beta \in B$, α and β are κ-independent with respect to $\langle A, B\rangle$.*

The definition of the conjunction $A \wedge B$ referred to above is as follows:

Definition 10 *Let A and B be two sets of formulas. The conjunction of A and B is*

$$A \wedge B =_{\text{def}} \{\alpha \wedge \beta | \alpha \in A \quad \& \quad \beta \in B\}$$

Using κ-independence, we can now prove a property even stronger than the one Pearl proposes:

Proposition 4 *If α and β are κ-independent with respect to $\langle A, A\rangle$, and $P(\kappa \wedge \beta \wedge \bigvee A) > 0$, then*

$$\kappa \overset{A}{\rhd} \alpha \Leftrightarrow \kappa \wedge \beta \overset{A}{\rhd} \alpha.$$

Proof. *By the definition of κ-independence,*

$$P(\alpha \wedge \beta | \kappa \wedge \bigvee A) \;=\; P(\alpha | \kappa \wedge \bigvee A) \times P(\beta | \kappa \wedge \bigvee B).$$

Hence,

$$P(\alpha \wedge \beta \wedge \kappa \wedge \bigvee A) \;=\; \frac{P(\alpha \wedge \kappa \wedge \bigvee A) \times P(\beta \wedge \kappa \wedge \bigvee A)}{P(\kappa \wedge \bigvee A)}.$$

Therefore,

$$P(\alpha | \beta \wedge \kappa \wedge \bigvee A) \;=\; \frac{P(\alpha \wedge \beta \wedge \kappa \wedge \bigvee A)}{P(\kappa \wedge \beta \wedge (\bigvee A))} \;=$$

$$\frac{P(\alpha \wedge \kappa \wedge \bigvee A)}{P(\kappa \wedge \bigvee A)} \;=\; P(\alpha | \kappa \wedge \bigvee A).$$

■

This property captures the intuition that mammals generally bear live young iff mammals which have a cold generally bear live young, given that having a cold is irrelevant to the form of one's offspring.

3.4 Conjunctive properties

Suppose α and β are properties evaluated with respect the sets of alternatives A and$B$, respectively. It is natural to evaluate $\alpha \wedge \beta$ with respect to the set composed of conjunctions of formulas from A and B. We can now prove the following properties:

Proposition 5

1. *If A and $\{\beta\}$ are κ-independent, and $P(\kappa \wedge \beta) > 0$, then for every $\alpha \in A$,*

$$\kappa \overset{A}{\rhd} \alpha \Leftrightarrow \kappa \overset{A \wedge \{\beta\}}{\rhd} \alpha \wedge \beta$$

2. *If A and B are κ-independent, then for every $\alpha \in A$, $\beta \in B$,*

$$\kappa \overset{A \wedge B}{\rhd} \alpha \wedge \beta \Rightarrow \kappa \overset{A}{\rhd} \alpha \quad \& \quad \kappa \overset{B}{\rhd} \beta$$

Proof. 1. *Since A and $\{\beta\}$ are κ-independent,*

$$P(\alpha \wedge \beta | \kappa \wedge \bigvee (A \wedge \{\beta\})) \ =$$

$$P(\alpha | \kappa \wedge \bigvee A) \times P(\beta | \kappa \wedge \beta) \ =$$

$$P(\alpha | \kappa \wedge \bigvee A).$$

2. *Observe that*

$$P(\alpha \wedge \beta | \kappa \wedge \bigvee (A \wedge B)) \ = \ P(\alpha \wedge \beta | \kappa \wedge \bigvee A \wedge \bigvee B).$$

By the definition of κ-independence, this is equal to

$$P(\alpha | \kappa \wedge \bigvee A) \times P(\beta | \kappa \wedge \bigvee B).$$

Since probabilities are numbers between 0 and 1, if

$$P(\alpha | \kappa \wedge \bigvee A) \times P(\beta | \kappa \wedge \bigvee B) > 0.5$$

then

$$P(\alpha | \kappa \wedge (\bigvee A)) > 0.5$$

and

$$P(\beta | \kappa \wedge (\bigvee B)) > 0.5.$$

■

The first property corresponds to the intuition that procreating mammals generally bear live young iff procreating mammals which have a cold generally bear live young—and have a cold. The second is aimed at capturing the idea that if unhealthy procreating mammals, in general, bear live young and have a cold, then procreating mammals generally bear live young and unhealthy mammals generally have a cold.

3.5 The conjunction fallacy

In light of the above, it is interesting to consider a well known psychological phenomenon, the *conjunction fallacy* (Tverksy and Kahneman 1983). An obvious fact about probabilities is that the probability of a conjunction cannot be greater than the probability of each one of the conjuncts separately, i.e. for every A and B, necessarily $P(A \wedge B) \leq P(A)$. Tversky and Kahneman have discovered that, in certain cases, humans consistently violate this rule. In one of their experiments, subjects were given a description of an individual, and then were given a list of occupations, and asked to rank them according to how likely it was, for each occupation, that the individual engaged in it. The following is a typical description:

> Linda is thirty-one years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations (Tversky and Kahneman 1983:297).

The test subjects were asked to rank a list of occupations with regard to how likely they were to apply to Linda. Among others, the list included the following:

(29) a. Linda is active in the feminist movement.

b. Linda is a bank teller.

c. Linda is a bank teller and is active in the feminist movement.

Tversky and Kahneman found that the vast majority of subjects ranked the conjunction (29.c) as more likely than (29.b), even though the latter is one of the conjuncts of the former. Remarkably, many subjects committed this error even when they had prior knowledge of statistics or when the conjunction rule was explicitly stated.

Tversky and Kahneman explain their findings by proposing that subjects do not use probability (or logic) when they make their probability judgments; instead, they report measures of resemblance. Thus (29.c) is judged more likely than (29.b) because Linda's character resembles a feminist bank teller more than it resembles a bank teller. This interpretation takes rather a dim view of human rationality. If Tversky and Kahneman are correct, human reasoning, at least in these cases, is not guided by reason, but rather by the ill-defined notion of resemblance; implausibly, this is claimed to be true even of people who are well versed in probability theory.

I am going to suggest an alternative account of this phenomenon. Consider (30), a slightly changed version of (29.c):

(30) Linda is a feminist bank teller.

Presumably, sentence (30) would elicit similar results to (29.c). Now consider what happens if we change the focus of the sentence:

(31) a. Linda is a FEMINIST bank teller.

b. Linda is a feminist BANK TELLER.

I have not, admittedly, conducted experiments based on these sentences, but let us consider our intuitive judgments regarding them. It seems that, whereas there is still a strong urge to consider (31.a) as probable, there is no such urge in the case of (31.b). Tversky and Kahneman's (1983) theory does not seem capable of accounting for this difference, since both sentences are presumed to compare Linda with a feminist bank teller. Assuming, following section 2.3.3, that different focus structures induce different sets of alternatives, the differences between (31.a) and (31.b) are easily accounted for. Sentence (31.a) induces alternatives of the form {*Linda is a feminist bank teller, Linda is a non-feminist bank teller*}. Since presumably Linda's description is more consistent with being a feminist than with the alternative, (31.a) is accepted. Sentence (31.b), on the other hand, induces a different set of alternatives, of the form {*Linda is a feminist bank teller, Linda is a feminist journalist, Linda is a feminist philosopher, Linda is a feminist social worker,...*}. Since presumably Linda's description is not as well correlated with a bank teller as it is with some other occupations, (31.b) is not perceived to be likely. The conjunction fallacy, then, is not a fallacy of probabilistic or logical reasoning; the error, in fact, is much more subtle. If I am correct, humans do employ sound reasoning in these experiments, but they fail to identify the relevant set of alternatives. In general, this is indeed, as we have seen, not an easy task; it is certainly more difficult than realizing that a conjunction entails its conjuncts.[33]

4 Automatic acquisition of adequate rules

In the previous section it was argued that adequate default rules are desirable. Once we have a set of rules, we can determine whether or not they are adequate. Rules which are found to be inadequate will then have to be dropped and replaced with adequate ones. This process, however, would be time consuming, expensive, and prone to errors. Figuring out which set of adequate rules is equivalent to a set of inadequate ones is by no means an easy task. It would be desirable, then, to apply the adequacy criterion already at the acquisition stage. That is to say, we would like to make sure that only adequate rules are acquired in the first place.

As the research on knowledge intensive rule-based systems progresses, it has become increasingly clear that the major stumbling block in the development of such systems is the acquisition of rules. Getting human experts to provide

[33] I am not unaware of the fact that in order to decide between Tversky and Kahneman's theory regarding the conjunction fallacy and mine, additional psychological experiments need to be carried out. Yet I believe the theory proposed here can account for the available data at least as well Tversky and Kahneman's approach can.

a detailed, explicit and useful list of the rules they use in their work is extremely difficult, time consuming, expensive, and—what is, perhaps, worst of all—inaccurate. No wonder, then, that there is growing interest in efforts to automate the acquisition process, wholly or in part.

The fact of the matter is that much of the necessary knowledge is out there, available for anyone to use: in manuals, guidebooks, almanacs, encyclopedias. This information, however, is provided in natural language, and cannot be readily used by a computer. A number of researchers (see, among others, Sager *et al.* 1987; Rinaldo 89; Moulin and Rousseau 1990; Gomez and Segami 1990; Castell and Verdejo 1991; Zarry 1992; Graziadio *et al.* 1992; Hodges and Cordova 1993; Jensen 1993; Sykes *et al.* 1994) have applied Natural Language Processing techniques to elicit rules automatically from natural language texts, with varying degrees of success.

When extracting rules from natural language texts, the most useful expressions in these texts are, of course, those which explicitly express rules. As it turns out, a great many of these expressions are generic sentences. To exemplify, let us consider a typical system, the one described in Hodges and Cordova (1993).

The system's purpose is to extract knowledge from texts in the domain of veterinary medicine.[34] In Hodges and Cordova's system, rules are represented as relations between objects. Each relation is represented as a frame, whose slots indicate the various arguments of the relation, their role and whether they are optional or obligatory. For example, the **cause** relation has two mandatory roles: the cause, which is either a disease agent (e.g. a bacterium or a virus) or a disease, and an effect, which is a disease. It has two optional arguments as well: a species, which may be any potential patient in the domain of veterinary medicine, and a location, which is some body part.

Hodges and Cordova describe how their system handles a number of example sentences occurring in their corpus (the Merck Veterinary Manual). Their first example (p. 925), for instance, contains no fewer than six generic terms, which I have italicized:

(32) *Herpesviruses* produce conjunctivitis in *the cat, cow, horse*, and *pig*, and transiently in *the dog.*

It seems, then, that an understanding of the semantics of generics would greatly facilitate the interpretation of such sentences, and, therefore, the automated acquisition of rules. However, none of the work on automatic acquisition I am aware of employs such a theory, or even acknowledges the need for it.

It may be argued that for the task we are discussing here, namely automatic acquisition of rules from natural language texts, a detailed theory of generics is

[34] The authors emphasize, though, that the algorithm is largely domain independent, and while the veterinary domain is just used as the primary example of its application, the system has been applied to other domains as well.

not required, and some rough approximation (e.g. that generics express universal quantification) would do just as well. However, I believe that this view is incorrect; ignoring the special nature of generics—in particular, their dependence on alternatives—may lead to unfortunate results. Hodges and Cordova's system, for example, will extract from (32) a causal relation between herpesviruses and conjunctivitis in the animals indicated. Yet it would be wrong to infer, on that basis, that any herpesvirus would produce conjunctivitis in any cat, cow, horse, pig and dog. The virus may not infect the animal in the first place; and even if the animal does get infected, it may not suffer any symptoms, because its body managed to fight off the disease, or the number of infecting viruses was not sufficient to make it sick. What (32) is saying is that *if* a herpesvirus produces a disease in one of the aforementioned animals, this disease is likely to be conjunctivitis. That is to say, (32) is evaluated with respect to a set of alternative diseases; only herpesviruses which cause some disease or other are relevant to the truth of the generic, and the rule extracted from it must represent this fact.

Consider another of Hodges and Cordova's examples:

(33) Healthy adult females in this species typically produce several offspring each year.

Sentence (33) does not actually appear in their corpus; the authors have made it up in order to demonstrate a specific point (the ambiguity of the word *produce*). As a matter of fact, (33) is not very likely to occur in a natural language text; something like (34) is much more likely.

(34) Cats typically produce several offspring each year.

That is to say, the writer of the text will not normally bother to specify that only healthy, adult female cats are relevant; this is something the readers are expected to figure out on their own. Again, the notion of a set of alternatives plays a prominent role in the interpretation of (34); this sentence seems to be evaluated with respect to alternative frequencies of producing offspring: only those cats which produce offspring with *some* frequency are considered, and (34) would be true just in case the majority of those produce more than one kitten a year. Any rule extracted from (34) would have to take this into account, otherwise we would draw incorrect conclusions about male cats or neutered females.

In fact, the implicit set of alternatives plays a role even in the interpretation of the very explicit (33), since some assumptions are left implicit nonetheless. For example, given a healthy adult female of this species which happens to be infertile, we should not expect it to produce any number of offspring, and this without invalidating (33). The rule extracted from (33), then, just like the rule extracted from (34), needs to take into account alternative frequencies of producing offspring, or it would lead to undesirable conclusions.

Any system, then, which takes seriously the task of extracting knowledge from natural language, must, sooner or later, come to terms with the meaning of generics. If the theory presented in section 2 is correct, such a system would need to be able to determine the set of alternatives with respect to which a given generic sentence is to be evaluated. This, in turn, requires a representation of determinables and determinates. A knowledge representation system based on these concepts has, in fact, been developed for independent reasons (Way 1991).

Way's main concern is the interpretation of metaphor. Her Dynamic Type Hierarchy theory states that the interpretation of metaphor involves the dynamic creation of new concepts. For example, in order to interpret (35), one needs to find a concept which subsumes both the concept **car** (the argument of *thirsty*) and the concept **animal** (the normal, nonmetaphorical selectional restriction on the argument of *thirsty*).

(35) The car is thirsty.

This concept, Way claims, would be **mobile entity**. The interpretation of the metaphor requires the dynamic construction of a specialization of the common concept, which, in this case, would be **mobile entity which requires liquid**. From here it is a short road (though by no means a trivial one; see Way 1991 for the details) to the correct interpretation of the metaphor, namely that the car requires some liquid, probably gasoline.

Way claims that, for the purpose of interpreting metaphor, a hierarchy based on the determinable-determinate relation is preferable to one based on the usual IS-A relation, and she goes on to develop such a hierarchy. It remains to be seen to what extent her knowledge representation system can be applied to the interpretation of generics, but it is encouraging to note that two unrelated problems in the interpretation of natural language seem to call for a representation based on the determinable-determinate relation.

5 Conclusion

In this paper I suggested a criterion by which to judge the adequacy of default inference rules. I believe that the criterion proposed here captures fairly well our intuitions about which rules are adequate and which are not. Moreover, systems of adequate default rules satisfy a number of properties deemed desirable by both philosophers and researchers in artificial intelligence.

The adequacy intuitions, and hence the criterion, turn out to hinge on a linguistic phenomenon—the meaning of generic sentences. While not necessarily subscribing to Barwise and Cooper's (1981) conclusion that "the traditional logical notions of validity and inference are a part of linguistics" (Barwise and Cooper 1981:203), I believe I have shown that when formalizing our commonsense intuitions, it is beneficial to look closely at the language we use to express them.

Acknowledgements

This paper has benefited greatly from the advice and comments of many people, especially Bob Carpenter, Jonathan Ginzburg, Clark Glymour, Hector Levesque, Lori Levin, Daniel Marcu, Allan Ramsay, Raymond Reiter, Carolyn Penstein Rosé, Oliver Schulte and Rich Thomason. Thanks also to audiences at Carnegie Mellon University, The Technion—Israel Institute of Technology, University of Toronto and The International Conference on Computational Semantics for helpful discussions.

References

[1] Åquist, L., Hoepelman, J., and Rohrer, C. 1980. Adverbs of frequency. In *Time, Tense and Quantifiers*, ed. by C. Rohrer. Tübingen: Verlag.

[2] Armstrong, D. M. 1978. *Universals and Scientific Realism* volume II: A Theory of Universals. Cambridge: Cambridge University Press.

[3] Bacchus, F. 1990. *Representing and Reasoning with Probabilistic Knowledge*. Cambridge, Massachusetts: MIT Press.

[4] Bacchus, F., Grove, A. J., Halpern, J. Y., and Koller, D. 1993. Statistical foundations for default reasoning. In *Proceedings of the International Joint Conference on Artificial Intelligence 1993 (IJCAI-93)* 563–569.

[5] Bacchus, F., Grove, A. J., Halpern, J. Y., and Koller, D. 1994. From statistical knowledge bases to degrees of belief. Unpublished manuscript.

[6] Barwise, J., and Cooper, R. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* 4.159–219.

[7] Blok, P. I. 1994. On the contribution of contextual information to the semantics and pragmatics of focus. In *Modeling Epistemic Propositions*, ed. by H. J. Bürchert and W. Nutt 5–17. German AI Society.

[8] Blok, P. I., and Eberle, K. 1994. What is the alternative? In *Focus & Natural Language Processing*, ed. by P. Bosch and R. van der Sandt volume 2: Semantics 235–246. Heidelberg, Germany: IBM Deutschland Informationssysteme GmbH, Scientific Centre, Institute for Logic and Linguistics.

[9] Brachman, R. 1985. "I lied about the trees" or, defaults and definitions in knowledge representation. *AI Magazine* 6.80–93.

[10] Carlson, G. 1977. *Reference to Kinds in English*. University of Massachusetts at Amherst dissertation. Also published 1980, New York: Garland, to which edition page numbers refer.

[11] Carlson, G. 1995. Truth-conditions of generic sentences: Two contrasting views. In *The Generic Book*, ed. by G. Carlson and F. J. Pelletier 31–51. University of Chicago Press.

[12] Castell, N., and Verdejo, M. 1991. Automatic extraction of factual information from text and its integration in a knowledge base. In *Proceedings of Recherche d'Informaion assiste par Ordinateur* 718–737 Barcelona: Universitat Autonama de Barcelona.

[13] Chomsky, N. 1975. Questions of form and interpretation. *Linguistic Analysis* 1.75–109.

[14] Cohen, A. 1995. Generics and frequency adverbs as probability judgments. In *Proceedings of the 10th Amsterdam Colloquium*, ed. by P. Dekker and M. Stokhof 163–182.

[15] Cohen, A. 1996. *Think Generic: The Meaning and Use of Generic Sentences*. Carnegie Mellon University dissertation.

[16] Delgrande, J. P. 1987. A first-order conditional logic for prototypical properties. *Artificial Intelligence* 33.105–130.

[17] Etherington, D. W., and Reiter, R. 1983. On inheritance hierarchies with exceptions. In *AAAI 83*.

[18] Gabbay, D. M., and Moravcsik, J. M. 1978. Negation and denial. In *Studies in Formal Semantics*, ed. by F. Guenthner and C. Rohrer 251–265. Amsterdam: North Holland.

[19] Gomez, F., and Segami, C. 1990. Knowledge acquisition from natural language for expert systems based on classification problem-solving methods. *Knowledge Acquisition* 2.107–128.

[20] Graziadio, B., Mussetto, P., Pietrosanti, E., and Pesce, E. 1992. A multi-layered architecture for automatic knowledge acquisition from legal texts. In *Proceedings of the 12th International Conference on Artificial Intelligence, Expert Systems, Natural Language* volume 3 165–177.

[21] Hirschberg, J. 1985. *A Theory of Scalar Implicature*. University of Pennsylvania dissertation.

[22] Hodges, J. E., and Cordova, J. L. 1993. Automatically building a knowledge base through natural language text analysis. *International Journal of Intelligent Systems* 8.921–938.

[23] Horn, L. R. 1972. *On the Semantic Properties of Logical Operators in English*. University of California, Los Angeles dissertation.

[24] Israel, D. 1980. What's wrong with non-monotonic logic? In *Proceedings of the 1st Annual National Conference on Artificial Intelligence* 99–101 Stanford.

[25] Jensen, O. 1993. Knowledge elicitation from natural language. In *Proceedings of the 6th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems* 331–440.

[26] Johnson, W. E. 1921. *Logic*. Cambridge: Cambridge University Press. Page numbers refer to the 1964 edition, published by Dover, New York.

[27] Karttunen, L. 1971. Implicative verbs. *Language* 47.340–358.

[28] Klein, E. 1980. A semantics for positive and comparative adjectives. *Linguistics and Philosophy* 4.1–45.

[29] Kyburg, H. 1961. *Probability and the Logic of Rational Belief*. Middletown, Conn.: Wesleyan University Press.

[30] Kyburg, H. 1970. Conjunctivitis. In *Induction, Acceptance and Rational Belief*, ed. by M. Swain 55–82. Dordrecht: Reidel.

[31] Levi, I. 1967. *Gambling with Truth; an Essay on Induction and the Aims of Science*. New York: Knopf.

[32] Loui, R. P. 1986. Real rules of inference: Acceptance and non-monotonicity in AI. Technical Report 191 Department of Computer Science, University of Rochester.

[33] Ludlow, P. 1989. Implicit comparison classes. *Linguistics and Philosophy* 12.519–533.

[34] Lyons, J. 1977. *Semantics*. Cambridge: Cambridge University Press.

[35] McCarthy, J. 1980. Circumscription—a form of nonmonotonic reasoning. *Artificial Intelligence* 13.

[36] Moore, R. C. 1985. Semantical considerations on nonmonotonic logic. *Artificial Intelligence* 25.

[37] Moulin, B., and Rousseau, D. 1990. Structuring a knowledge base from the content of prescriptive texts. In *Proceedings of the 10th International Workshop on Expert Systems and their Applications* 179–173.

[38] Neufeld, E. 1989. Defaults and probabilities; extensions and coherence. In *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, ed. by R. J. Brachman, H. J. Levesque, and R. Reiter 312–323.

[39] Partee, B. H. 1991. Topic, focus and quantification. In *Proceedings of the First Conference on Semantics and Linguistic Theory*, ed. by S. Moore ad A. Z. Wyner number 10 in Working Papers in Linguistics 159–187. Cornell University.

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

[41] Pearl, J. 1991. Probabilistic semantics for nonmonotonic reasoning. In *Philosophy and AI*, ed. by R. Cummins and J. Pollock 157–187. Cambridge, Massachusetts: MIT Press.

[42] Pollock, J. 1990. *Nomic Probability and the Foundations of Induction*. New York: Oxford University Press.

[43] Reiter, R. 1980. A logic for default reasoning. *Artificial Intelligence* 13.81–132.

[44] Rinaldo, F. 1989. *Deriving Rules for Medical Expert Systems Using Natural Language Parsing and Discourse Analysis*. Illinois Institute of Technology dissertation.

[45] Rooth, M. E. 1985. *Association with Focus*. University of Massachusetts at Amherst dissertation.

[46] Rooth, M. E. 1992. A theory of focus interpretation. *Natural Language Semantics* 1.75–116.

[47] Rosenberg, J. F. 1966. *A Study of the Determinable-Determinate Relation*. University of Pittsburgh dissertation.

[48] Sager, N., Friedman, C., and Lyman, M. 1987. *Medical Language Processing: Computer Management of Narrative Data*. Reading, Massachusetts: Addison-Wesley.

[49] Schubert, L. K., and Pelletier, F. J. 1987. Problems in the representation of the logical form of generics, plurals, and mass nouns. In *New Directions in Semantics*, ed. by E. LePore 385–451. London: Academic Press.

[50] Searle, J. 1959. On determinables and the notion of resemblance. In *Proceedings of the Aristotelian Society, supplementary volume 33* 141–158 London. Harrison and Sons.

[51] Searle, J. 1967. Determinables and determinates. In *The Encyclopedia of Philosophy*, ed. by P. Edwards volume 4 357–359. New York: Macmillan Publishing Co. and The Free Press.

[52] Shastri, L. 1989. Default reasoning in semantic networks: a formalization of recognition and inheritance. *Artificial Intelligence* 39.283–355.

[53] SYKES, J. T., KONSTANTINOU, V., and MORSE, P. L. R. 1994. Extracting explicit and implicit knowledge from natural language texts. In *Conceptual Structures: Current Practices. Proceedings of the 1st International Conference on Conceptual Structures* 97–112.

[54] THOMASON, R. 1969. Species, determinates, and natural kinds. *Nous* 3.95–101.

[55] THOMASON, R. 1988. Theories of nonmonotonicity and natural language generics. In *Genericity in Natural Language: Proceedings of the 1988 Tübingen Conference*, ed. by M. Krifka 395–406. SNS-Bericht 88–42, Seminar für natürlich-sprachliche Systeme, Tübingen University, Germany.

[56] TVERSKY, A., and KAHNEMAN, D. 1983. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*.

[57] VALLDUVÍ, E. 1992. *The Informational Component*. New York: Garland.

[58] WAY, E. C. 1991. *Knowledge Representation and Metaphor*. Dordrecht: Kluwer.

[59] ZARRY, G. P. 1992. Semantic modeling of the content of (normative) natural language documents. In *Proceedings of the 12th International Conference on Artificial Intelligence, Expert Systems, Natural Language* volume 3 181–193.