Genericity

Ariel Cohen

Appeared in *Linguistische Berichte*, 10, 2002

1 Introduction

We often express our knowledge about the world in sentences such as the following:

(1) a. Ravens are black.
b. Tigers have stripes.
c. Mary jogs in the park.

We refer to such sentences as *generics*. They appear to express some sort of generalization: about ravens, about tigers, and about Mary, respectively. Yet it is far from clear exactly what they mean. What does it mean to say that some generalization holds?

It turns out that there are a great many theories trying to answer this question. This, in itself, is a fact that is in need of explanation. Why are generics so puzzling? What is it about them that forces researchers to come up with one theory after another, with no clear agreement on what the correct theory is? And, if they are so strange, why are generics so prevalent?

In this article we will consider some of the puzzles concerning generics, why they are so hard, and the various solutions proposed. Let me say at the outset that readers who expect to find definitive answers to these puzzles will unfortunately be disappointed. But if not the answers, I hope this article will at least convey the depth and significance of the problems.

2 Are generics quantificational?

Possibly the first interpretation that comes to mind is that generics expresses a quantification of some sort. Perhaps a sentence such as (1.a) is really just a different way to say something like

(2) Every raven is black.

Things are not that simple, unfortunately. First, note that generics do not express universal quantification: while (1.a) is true, (2) is false, because there are some albino ravens. Still, even if the quantifier is not the universal one, perhaps generics use some other quantifier. If this is the case, our role is to figure out what this quantifier is.

This, however, is far from an easy task. Consider the following examples of generics:

(3) a. Dogs are mammals.
b. Birds fly.
c. Mammals bear live young.
d. The Frenchman eats horsemeat.
e. Bulgarians are good weightlifters.
f. The giant panda is an endangered species.
g. Primary school teachers are female.
h. People are over three years old.
i. Members of this club help each other in emergencies.
j. Supreme Court judges have a prime Social Security number.
k. **A:** Nobody in India eats beef.
B: That's not true! Indians do eat beef.

Sentences (3.a)–(3.f) are all presumably true, but what is it that makes them true? Sentence (3.a) seems to hold for all dogs, (3.b) for most birds, (3.c) for most female mammals (presumably less than half the total number of mammals), (3.d) for rather few Frenchmen, (3.e) for very few Bulgarians, and (3.f) for no individual giant panda. On the other hand, the majority of primary school teachers are female, and the majority of people are over three years old, and yet (3.g) and (3.h) are not true. Even if no emergencies ever occurred, (3.i) may be true, and even if all Supreme Court judges happened to have a prime Social Security number, (3.j) may be false. The truth of B's answer in (3.k) requires only that *some* Indians eat beef.

The diversity of interpretations of generics, as exemplified by the sentences in (3), poses severe problems for any theory that attempts to relate the truth or falsity of a generic to properties of individual instances; e.g. any theory that relates the truth of (1.a) to properties of individual ravens. Given this difficulty, there are two approaches one may take.

One, which Carlson (1995) calls the *rules and regulations* approach, is to deny that any semantic relation exists between generics and properties of individuals; generics, according to this view, are evaluated with respect to rules and regulations, which are basic, irreducible entities in the world. Each generic sentence denotes a rule; if the rule is *in effect*, in some sense (different theories construe differently what it means for a rule to be in effect), the sentence is true, otherwise it is false. The rule denoted by a generic may be physical, biological, social, moral, etc. The paradigmatic cases for which this view seems readily applicable are sentences that refer to conventions, i.e. man-made, explicit rules and regulations, such as the following example (Carlson 1995):

(4) Bishops move diagonally.

According to the rules and regulations view, (4) is not about the properties of individual bishop moves, but refers directly to a rule of chess; it is true just in case one of the rules of chess is that bishops move diagonally. It is important to note that, according to the rules and regulations view, *all* generics are so analyzed: for example, (1.a) is true not because of the properties of individual ravens, but because there is a rule in the world (presumably a rule of genetics) that states that ravens are black.

An alternative approach, which Carlson (1995) calls the *inductivist* view, is to accept the existence of a semantic relation between generics and properties of individuals. Theories that take this view attempt to define this relation in such a way that its nature (possibly in conjunction with facts about context, intonation and world knowledge) may account for the diversity of readings of generics, exemplified in (3).

The rules and regulation view and the inductivist view are each a cover term for a number of specific proposals. Let us briefly consider some of them.

2.1 Rules and regulations theories

Carlson (1977) proposes that a generic expresses simple predication of a property of a kind. Thus, (1.a) has a very similar logical form to that of

(5) Nevermore is black.

Both sentences express simple predication rather than quantification. The difference is only that whereas (5) predicates a property (being black) of an object (the individual raven Nevermore), (1.a) predicates this property directly of the kind *raven*. Thus, Carlson's approach can, in principle, account for all the examples in (3). This, however, is done at a price: no explanation is given for why, say, eating horsemeat is a property of the kind *Frenchman*, or why being female is not a property of the kind *primary school teacher*. Moreover, Carlson's theory cannot account for scope ambiguities of generics, exemplified by the following sentences (from Schubert and Pelletier 1987):

(6) a. Canadian academics are supported by a single granting agency.
b. Storks have a favorite nesting area.
c. Sheep are black or white.
d. Whales are mammals or fish.

Krifka (1987) proposes that generics express a *default rule*. This is a type of inference rule that allows for exceptions. For example, we may assume that any raven, by default, is black, but we are ready to retract this conclusion if we learn more information about the raven—that it is an albino, that it fell into a bucket of whitewash, etc. According to Krifka, then, (1.a) is true just in case every raven is black, unless its being black is not consistent with the facts assumed so far. One challenge that Krifka's approach has to face is to determine which rules are in effect and which are not. For example, a default rule that states that a primary school teacher is female is presumably a useful one, since, if we know that someone is a primary school teacher, we can reasonably assume that she is a woman, unless we learn something to the contrary. Yet this rule is not in effect, since (3.g) is false. On the other hand, a default rule stating that a given Bulgarian is a good weightlifter is probably not very useful—if we know that someone is Bulgarian, we will be reluctant to conclude, solely on the basis of this information, that he or she is a good weightlifter. Yet this rule *is* in effect, since (3.e) is true.

An alternative theory is that generics express not rules in the world, but rules of conversation (McCarthy 1986; Reiter 1987). Thus, the truth of (7.a) implies that there is a convention of verbal behavior, according to which, whenever a speaker says (7.b), the hearer is expected to infer (7.c).

(7) a. Birds fly.

b. Tweety is a bird.

c. Tweety flies.

This view can be cast in terms of *conversational implicature* (Grice 1975): (7.a) is true just in case (7.b) conversationally implicates (7.c). Just like the case of conversational implicature, the inference from (7.b) to (7.c) is cancelable, e.g. by the additional information that Tweety is a penguin. Some sentences in (3) appear better suited for this approach than others. For example, it is hardly a language convention that if we hear that Charles is a Frenchman, we are supposed to infer that he eats horsemeat, and yet (3.d) is true. On the other hand, there may very well be a convention that if we talk about a person, we may safely assume that he or she is over three years old; and yet (3.h) is not true.

Rather than a verbal convention, other researchers (Geurts 1985; Declerk 1986) have suggested that a true generic sentence corresponds to a cultural convention, a stereotype. Thus (1.a) is true because it corresponds to stereotypical beliefs about ravens in our culture—the stereotypical raven is black. Not all sentences are as amenable to the treatment proposed by this theory: for example, it may very well be that the stereotypical primary school teacher is female, and yet (3.g) is not true. Another problem with this theory is that it takes a generic to be not a statement of fact about the world, but rather a statement about the stereotypical beliefs prevailing in one's culture. But it appears that this is not the way we interpret generics. For example, while (8.a) (after Krifka *et al.* 1995) is a coherent sentence, (8.b) is not:

(8) a. Snakes are stereotypically believed to be slimy, but in fact they are not.

b. *Snakes are slimy, but in fact they are not.

Working within the framework of Situation Semantics (Barwise and Perry 1983), several researchers (ter Meulen 1986; Cavedon and Glasbey 1994) have proposed that generic sentences express *constraints* on situations. Roughly speaking, (1.a) expresses the constraint that every situation involving a raven involves a black raven.

Cavedon and Glasbey (1994) treat constraints as part of the natural order of the world; in particular, they are not reducible to properties of individual instances. This property of constraints enables them to tolerate exceptions, so that (1.a) is true even if some ravens are not black. Crucial to Cavedon and Glasbey's account is the notion of a *channel* (Barwise and Seligman 1992). Roughly speaking, the role of channels is to relativize the interpretation of a generic sentence to a given context. For example, (9.a) is evaluated relative to a channel which is concerned with female peacocks, and (9.b), with respect to a channel which is concerned with male peacocks.

(9) a. Peacocks lay eggs.

b. Peacocks have brightly colored tail-feathers.

Just like other theories that follow the rules and regulations approach, both the strength and the weakness of Cavedon and Glasbey's account lies in this separation between the meaning of a generic and the properties of individual instances. The theory implies that the truth of a generic cannot be observed directly; as such, all the sentences in (3), as well as many others, can be accounted for: the true ones correspond to a constraint that is in effect, the false ones do not. However, this is also a weakness, since while the theory cannot be refuted, it is not clear what it would take corroborate it—there is no clear prediction about the way things ought to be in the world for a generic sentence to be true or false. Some researchers have considered this to be an undesirable situation for a truth-conditional semantics, which defines the meaning of a sentence as the states of affairs that would make it true. Instead, they proposed versions of the inductivist view, in the hope of providing some relation between the facts obtaining in the world and the truth of a generic sentence. We will now turn to some of these theories.

2.2 Inductivist theories

The idea underlying the inductivist approach is rather simple. A generic sentence is true just in case sufficiently many relevant individuals in the domain of the generic satisfy the predicated property. This idea is, of course, vague on at least two issues: which instances count as "relevant," and how many is "sufficiently many"? Various inductivist approaches offer different answers to these questions. Let us briefly discuss some of them.

Farkas and Sugioka (1983) suggest that the quantifier is *significantly many*. For example, (1.a) is true because significantly many ravens are black. *Significantly many* is, of course, a vague quantifier, so for many generic sentences it could be argued that this quantifier is applied correctly. It is not clear, however, that all generics can be accounted for in this way. For example, significantly many people are over three years old, and yet (3.h) is not true.

Another possibility is that the appropriate quantifier is *most*. For example, (1.a) is true because most ravens are black. This approach has not been explicitly proposed in any source of which I am aware, but it appears to be alluded to in Parsons (1970) and Nunberg and Pan (1975), as well as in many theories proposed within the Artificial Intelligence community, as discussed by Pelletier and Asher (1997). The problem with this proposal is that in order for a generic to be true, it does not need to be the case that the majority of individuals satisfy the predicated property: sentences (3.c) through (3.e) are good counterexamples.

A more sophisticated version of this theory has been proposed by Schubert and Pelletier (1989). According to them, generics do not quantify over actual individuals, but possible ones. Thus, for example, if most, or even all actual Supreme Court judges had a prime Social Security number, (3.j) would not be true: if we consider all possible judges, it is not true that most of them have a prime Social Security number. Schubert and Pelletier suggest that *most* is defined relative to a measure function on possible worlds, which favors worlds that are close to the real one in terms of the essential or inherent nature of things.

What is meant by terms such as "inherent" or "essential" is candidly left open by Schubert and Pelletier. Apparently, it is a modal notion, but it is clearly not the same as logical necessity: there is no logical necessity for birds to fly or for mammals to bear live young. The problematic nature of these notions becomes even more apparent when we consider sentences such as the following:

(10) a. A cheetah outruns any other animal.

b. Spices are affordable.

c. Gold cubes are smaller than 10 cubic meters (after Koningsveld 1973).

d. Dogs annoy Sam.

Perhaps running fast is an inherent property of cheetahs, but certainly not the property of running faster than any other animal, since some other animal could have been faster. Affordability is not a necessary property of spices; in fact, throughout much of history, spices were extremely expensive; yet (10.b) is true nonetheless. Similarly, we would be hard-pressed to claim that gold cubes are inherently smaller than 10 cubic meters, or that annoying Sam is an essential property of dogs.

If *most* is problematic, perhaps the universal quantifier will work better. In fact, a number of scholars (e.g. Quine 1960; Bartsch 1972; Bacon 1974; Bennett 1974) have assumed that a generic sentence expresses universal quantification over actual individuals. It should be stressed, however, that these researchers do not propose a theory of generics as such; the researcher has some different goal in mind, and the precise interpretation of generics is not important for that goal. A moment's reflection, of course, shows that this suggestion cannot stand. If a counterexample is required, our well worn example, (1.a), will suffice. This sentence is true, despite the fact that not all ravens are black.

Alternatively, we can take the quantifier to be a restricted universal (Declerk 1991; Chierchia 1995). Context, according to this view, provides a restriction for the domain of the quantifier. For example, (3.c) says that all relevant mammals bear live young. Which are the relevant mammals? Declerk and Chierchia do not provide a principled account of how this restriction is obtained. Presumably, male mammals are irrelevant, as are females that are too young or too old to bear live young, etc. Strange mammals, such as the platypus, which lays eggs, are also somehow left outside the domain of the quantifier. The remaining mammals do lay eggs, hence the truth of (3.c).

For Declerk, the universal quantifier ranges over actual individuals; for Chierchia, it ranges over possible individuals. Hence, Chierchia, unlike Declerk, can explain why (3.j) is not true, but his account suffers from similar problems to those of Schubert and Pelletier (1989).

Schubert and Pelletier (1987) offer a more detailed discussion of how the restriction to relevant individuals is provided. It could be induced by the presupposition of the VP, as in (11), by focus, as in (12), by the linguistic context, as in (13), or by an explicit *when* clause, as in (14).

(11) Cats land on their feet.
= Cats that drop to the ground land on their feet.

(12) a. Leopards attack monkeys IN TREES.
= Monkeys that are attacked by leopards are in trees.

b. Leopards attack MONKEYS in trees.
= Animals in trees that are attacked by leopards are monkeys.

c. LEOPARDS attack monkeys in trees.
= Animals that attack monkeys in trees are leopards.

(13) Most monkeys flee when leopards approach. Baboons form a protective circle with males on the outside.
= Baboons approached by a leopard form a protective circle with males on the outside.

(14) When cats drop to the ground, they land on their feet.
= Cats that drop to the ground land on their feet.

Later works (Partee 1991; Krifka 1995; Cohen 1996) combine this approach with theories of focus, claiming that generics are *associated* with focus, in the sense of Rooth (1985): focus provides a set of alternatives that restricts the domain of the generic quantifier.

For example, the following sentences have different truth conditions:

(15) a. In Saint Petersburg, ballerinas escorted OFFICERS.
b. In Saint Petersburg, BALLERINAS escorted officers.

Sentence (15.a) is true just in case, whenever a ballerina accompanied someone, it was generally an officer (but officers may have had other companions as well); sentence (15.b), on the other hand, conveys the statement that, whenever someone escorted an officer, it was generally a ballerina (but ballerinas may have accompanied other people as well).

Rooth suggests that the union of the set of alternatives induced by focus determines the domain of the quantifier. With respect to (15.a), this union is the set of ballerinas who escorted someone. The generic quantifier then quantifies over this set, conveying that such ballerinas generally escorted officers.

With respect to (15.b), the union of the alternatives would be the set of officers escorted by somebody. Hence (15.b) quantifies over such officers, stating that such companions were generally ballerinas.

This type of approach is quite powerful, in providing empirically testable predictions about the interpretations of many generics—see, for example, the effect of focus in (12). It can even be explained why, in cases such as (3.k), generics get quasi-existential readings—B's response only requires that some Indians eat beef. In such cases, it has been proposed (Cohen 1996), the role of the contrastive focus is to restrict the domain to only those Indians who eat beef; if this domain is not empty, the sentence is true.

Yet it is not clear that such approaches can account for the full range of readings of generics. For example, it is hard to see what sort of restriction of the domain of Frenchmen would yield the truth of (3.d), when the sentence is not uttered in a contrastive context.

Yet another view of generics as expressions of universal quantification is that the quantifier quantifies over *normal* individuals (Delgrande 1987; Morreau 1992; Asher and Morreau 1995; Krifka 1995; Pelletier and Asher 1997, among others). Sentence (1.a) is true, according to this view, because all normal ravens are black—albino ravens are abnormal ravens. Normality is taken to be a modal notion. Following Kratzer (1981), a partial ordering relation is assumed to be defined on possible worlds. This relation orders worlds according to their normality. Then, a generic sentence such as (1.a) is true just in case in all worlds that are most normal, all ravens are black.

Thus, we can account for sentences such as (3.g) and (3.h): although male teachers are in the minority, they are still normal teachers; and although most people are over three years old, babies are still normal people. On the other hand, we can account for the truth of (3.i) even when no emergencies occurred in the actual world, provided that in those most normal worlds where emergencies do occur, all members of the club help each other.

One problem with these approaches is that the ordering source of normality is not given an independent definition. Why is a black raven normal, and a white raven abnormal? Note that the interpretation of normality seems to change from sentence to sentence, as the following sentences (from Krifka *et al.* 1995) indicate:

(16) a. Two and two equals four (normal = the rules of mathematics hold).

b. A spinster is an old, never-married woman (normal = the rules of English hold).

c. This machine crushes oranges (normal = machines perform as intended).

d. Mary smokes cigarettes (normal = Mary shows her typical behavior).

e. Bob jumps 8.90 meters (normal = Bob performs as well as he can).

f. A lion has a mane (normal = stereotypical properties hold).

g. Six apples cost one dollar (normal = the actual world).

h. A turtle is long-lived (normal = ?).

i. A pheasant lays speckled eggs (normal = ?).

There is some debate over what the standard of normality would be for (16.h) and (16.i), since worlds in which all turtles reach an old age (no predators?) or where all pheasants lay eggs (no males?) do not, on the face it, appear to be normal. But perhaps this problem could be solved by adding a restriction to the domain of the generic quantifier (Krifka 1995; Pelletier and Asher 1997), thus, in a sense, combining the normality approach with a domain-restriction theory such as Schubert and Pelletier (1987).

Other than these skeptical doubts, quantification over normal individuals runs into some empirical problems as well. It is not clear how it would account for (3.d) and (3.e): it is hardly the case that all normal Frenchmen eat horsemeat or that all normal Bulgarians are good weightlifters. Moreover, sentences that express relations pose a particular problem for this approach. Sentence (17.a), for example, clearly does not mean (17.b).

(17) a. Women live longer than men.

b. Every normal woman lives longer than every normal man.

A somewhat similar idea is to regard generics as expressions of universal quantification over a set of *typical individuals*, rather than normal ones (Heyer 1985, 1990; see also Link 1995). Unlike normal individuals, typical individuals are usually not defined in terms of possible worlds. A distinction is drawn between characteristic and noncharacteristic properties of kinds; those individuals that possess the characteristic properties are considered to be typical representatives. A generic sentence is true, then, to the extent that it is true of all typical individuals: since all typical ravens are black, (1.a) is true. This approach shares many of the strengths of the normality approach, but also its weaknesses. In particular, it provides no independent, non-circular definition of typicality.

Given the difficulty of deciding what the meaning of the generic quantifier is, some people have proposed that it is, in fact, ambiguous.

Strzalkowski (1988) takes a generic such as (18.a) to be ambiguous between the senses paraphrased by (18.b) and (18.c).

(18) a. Birds fly.
b. All except for a negligible number of birds fly.
c. A non-negligible number of flying animals are birds.

In this way he is able to account for sentences such as (3.d) and (3.e), assuming that a non-negligible number of horsemeat eaters are French, and that a non-negligible number of good weightlifters are Bulgarian. However, his theory predicts, worngly, that (19) is true, since a non-negligible number of birds are grey.

(19) Grey animals are birds.

It should be noted that under both readings of the generic quantifier, Strzalkowski takes it to quantify over actual individuals. Hence, his theory is subject to the problems with sentences such as (3.j) and (3.i).

In contrast, Dahl (1975) interprets the generic quantifier as quantifying over possible worlds. According to him, the quantifier is ambiguous between (restricted) universal and existential quantification over worlds, i.e. between the modal notions of necessity and possibility. Thus, (3.a) states that all dogs are necessarily mammals. Dahl can account for (3.d): it means that if we pick an arbitrary Frenchman, it is possible that he would eat horsemeat. This approach, however, would predict no difference between (3.d) and (20), since it is also possible that an arbitrary American would eat horsemeat.

(20) The American eats horsemeat

Dahl's approach can handle with ease the cases that are difficult for Strzalkowski's theory, such as (3.j) and (3.i). However, just like Schubert and Pelletier's (1989) theory, it runs into difficulties with cases of contingent generics, such as those in (10).

Cohen (1996; 2000) has a different account of the ambiguity of generics. Generics, according to this proposal, express probability judgments. Thus, (3.b) is about the probability that an arbitrarily chosen bird flies, and (3.d) is about the probability that an arbitrarily chosen Frenchman eats horsemeat. However, generics are ambiguous with respect to the requirement that this probability needs to satisfy in order for the sentence to be true: the most plausible interpretation of (3.b) is that the probability is higher than some constant (specifically, 0.5); the most plausible interpretation of (3.d) is that the probability is greater than the probability that some arbitrary person eats horsemeat. Thus, (3.d) is true just in case, if we pick an arbitrary Frenchman, however unlikely this person is to eat horsemeat, he would still be likelier to do so than a person of an arbitrary nationality.

2.3 Combining the two types of theory

It appears that there are some generics, e.g. (4) that are better explained by rules and regulations theories, and others, e.g. (1.a), that are better explained by inductivist theories. One may wish to consider, then, whether the two types of theory can somehow be combined.

This possibility is rejected by Carlson (1995). He describes the two approaches as a dichotomy: one has to choose one or the other, but not both. How can we decide which? One way is to consider a case where the behavior of observed instances conflicts with an explicit rule. For example, Carlson describes a supermarket where bananas sell for \$.49/lb, so that (21.a) is true. One day, the manager decides to raise the price to \$1.00/lb. Immediately after the price has changed, claims Carlson, sentence (21.a) becomes false and sentence (21.b) becomes true, although all sold bananas were sold for \$.49/lb.

(21) a. Bananas sell for \$.49/lb.

b. Bananas sell for \$1.00/lb.

Consequently, Carlson reaches the conclusion that the rules and regulations approach is the superior one.

This conclusion has been challenged by Greenberg (1998) and Cohen (forthcoming). Suppose the price has, indeed, changed, but the supermarket employs incompetent cashiers who consistently use the old price by mistake, so that customers are still charged \$.49/lb. In this case, there seems to be a reading of (21.a) which is true, and a reading of (21.b) which is false. These readings are more salient if the sentence is modified by expressions such as *actually* or *in fact*:

(22) a. Bananas actually sell for \$.49/lb.

b. In fact, bananas sell for \$1.00/lb.

Consequently, Greenberg and Cohen claim that generics are ambiguous: on one reading they express a descriptive generalization, stating the way things are. Under the other reading, they carry a normative force, and require that things be a certain way. When they are used in the former sense, they should be analyzed by some sort of inductivist account; when they are used in the latter sense, they ought to be analyzed as referring to a rule or a regulation. The respective logical forms of the two readings are different; whereas the former reading involves, in some form or another, quantification, the latter has a simple predicate-argument structure: the argument is the rule or regulation, and the predicate holds of it just in case the rule is "in effect."

A language that makes an explicit distinction between these two types of reading is French. In this language, generically interpreted plural nouns are preceded by the definite determiner *les*, whereas the indefinite determiner *des* usually induces existential readings. However, *des* may also be used to make a normative statement, i.e. to express some rule or regulation.

(23) a. Des agents de police ne se comportent pas ainsi dans une situation d'alarme.

'INDEF-PL police officers do not behave like that in an emergency situation.'

b. Les agents de police ne se comportent pas ainsi dans une situation d'alarme.

'DEF-PL police officers do not behave like that in an emergency situation.'

An observation which de Swart (1996) ascribes to Carlier (1989) is that (23.a) "would be uttered to reproach a subordinate with his behavior. [(23.b)] does not have the same normative value, but gives us a descriptive generalization which could possibly be refuted by providing a counterexample."

3 Lawlikeness and intensionality

3.1 Generics are lawlike

Perhaps one of the reasons why it is so hard to determine whether generics are quantificational, and, if so, what the quantifier is, is that generics are *lawlike*. The distinction between lawlike and nonlawlike statements is a well known in philosophy, and is easily demonstrated using universally quantified sentences. For example (24.a) intuitively expresses a law of nature; (24.b), in contrast, expresses an accidental fact.

(24) a. All copper wires conduct electricity.

b. All coins in my pocket are made of copper.

One way to characterize the difference between lawlike and nonlawlike statements is that only the former, not the latter, support counterfactuals. Thus, (24.a) entails (25.a), but (24.b) does not entails (25.b).

(25) a. If this were a copper wire, it would conduct electricity.

b. If this coin were in my pocket, it would be made of copper.

Note that we can turn (24.a), but not (24.b), to a felicitous generic; (26.a) is fine (and true) but (26.b), under its generic interpretation, is odd (cf. (3.j) above).

(26) a. Copper wires conduct electricity.

b. Coins in my pocket are made of copper.

Generics, in general, support counterfactuals; the truth of (27.a) entails (27.b).

(27) a. Birds fly.

b. If Dumbo were a bird, he would probably fly.

It is tempting to think that rules and regulations theories are particularly well suited to handle this aspect of generics: it seems that all we need to require is that the rule or regulation a generic denotes be nonaccidental. Things are not that simple, however: rules and regulations approaches have difficulties accounting for the fact that generics support counterfactuals. If there is no relation between the truth of (27.a) and the flying abilities of actual birds, why should there be such a relation between its truth and the flying abilities of hypothetical birds?

Inductivist theories face difficulties too. If generics involve a quantifier, it has rather special properties: this quantifier must be sensitive not only to the number of individuals satisfying a certain property, but also to whether the statement is lawlike or not. It is for this reason that, as we have seen above, many researchers proposed modal treatments of generics; the hope is that the notion of lawlikeness is similar enough to the notion of necessity to be formalizable within a possible worlds framework. If, indeed, generics can be captured by a theory that is based on possible worlds, it follows that they must be intensional. Let us now turn to the issue of intensionality.

4 Are generics intensional?

Suppose ψ_1 and ψ_2 are two extensionally equivalent properties, i.e. at this moment in time and in the actual world, the respective sets of individuals that satisfy ψ_1 and ψ_2 are equal. If generics behave extensionally, we would expect the following sentences to have the same truth conditions for every property ϕ:

(28) a. ψ_1s are ϕ.

b. ψ_2s are ϕ.

This does not hold in general. Consider (29), from Carlson (1989).

(29) A computer computes the daily weather forecast.

Carlson observes that

> "the daily weather forecast" requires an *intensional* interpretation, where its meaning cannot be taken as rigidly referring to the present weather forecast, e.g. the one appearing in today's copy of the *Times* predicting light rain and highs in the upper thirties (p. 179, emphasis added).

For example, if today's weather forecast predicts a blizzard, this may well be the main news item. Yet, (29) does not entail

(30) A computer computes the main news item.

While a computer may have computed today something that turned out to be the main news item, this does not hold in general; on most days, the main news item will not be computed by a computer, hence (30) is false.

Intensionality, it is important to note, does not come in one form only. In particular, a construction may exhibit intensionality with respect to the time index, but not with respect to possible worlds, or vice versa. For example, Landman (1989), in his discussion of groups, draws the following distinction:

> The intensionality that I am concerned with here concerns... the fact that committees **at the same moment of time** can have the same members, without being the same committee. Another form of intensionality concerns the well known observation that... in the course of time, they may change their members, while staying the same committee. I do not think that this kind of intensionality has the same source as the 'atemporal' intensionality that is the topic of this paper (pp. 726–727, original emphasis).

Generics and frequency statements, it turns out, behave intensionally with respect to the time index, but not with respect to possible worlds. Suppose that the weather report is Mary's favorite newspaper column. Then (31) would have the same truth conditions as (29), although there are any number of worlds where Mary has no interest in the daily weather forecast:

(31) A computer computes Mary's favorite newspaper column.

To give other examples, it is true in the actual world that the whale is the largest animal on earth, and the quetzal is Guatemala's national bird, but there are any number of possible worlds where this is not the case. Yet (32.a) and (33.a) have the same respective truth conditions as (32.b) and (33.b).

(32) a. The whale suckles its young.
b. The largest animal on earth suckles its young.

(33) a. The quetzal has a magnificent, golden-green tail.
b. Guatemala's national bird has a magnificent, golden-green tail.

Generics, then, are parametric on time, but not on possible worlds; if two properties have the same extension throughout time, they can be freely interchanged in a generic sentence *salva veritate*. In other words, the truth conditions of the generic

(34) ψs are ϕ

do not depend on the extensions of ψ and ϕ in any other world but the actual one, though the truth conditions do depend on the extensions of these properties at different times.

How can a theory of generics account for this behavior? Clearly, a fully extensional theory, such as Declerk (1991) or Strzalkowski (1988), will not do justice to this phenomenon; according to such theories, generics ought not to be parametric on either time or possible worlds, which is not the case. On the other hand, a fully intensional theory would not do either, since it would predict that generics are parametric on possible worlds, which they are not.

Theories that make use of possible worlds, but restrict them to worlds that are normal, or that are close to the actual world in terms of its essential properties, fare better. They do, however, have to face the problem of defining normality or essence in such a way, that a world where Mary is not interested in the weather, or where the quetzal is not Guatemala's national bird, is somehow abnormal, or violates essential principles holding in the actual world.

An alternative way to explain the behavior of generics with respect to intensionality has been proposed by Cohen (1999), who uses a *branching* model of time. That is to say, for any given time there is more than one possible future: there is a future where it is going to rain tomorrow, and one where it is not. The generic (34) is evaluated with respect to all those futures where the frequency of ϕ among ψs is more or less the same as during an interval of time containing the reference time of the sentence. For example, (29) is true just in case in the extended present the daily weather forecast is computed by a computer, but (30) is false because in the extended present, the weather is rarely the main news item. On the other hand, (31) is true just in case (29) is true, given that in the extended present Mary's preference for the weather forecast remains unchanged.

4.1 Frequency adverbs

It is often pointed out that generics are similar to sentences involving an overt adverb of quantification. Consider the sentences in (1), when modified by an overt adverb.

(35) a. Ravens are usually black.
b. Tigers always have stripes.
c. Mary sometimes jogs in the park.

Just like the generics in (1), these sentences express some generalization about ravens, tigers, and Mary, respectively. The difference is that, unlike generics, which have no quantifier (according to the rules and regulations view), or an implicit quantifier with some special properties (according to the inductivist view), here we have an explicit quantifier. Thus, the sentences in (35) are also similar to overtly quantified sentence such as the following:

(36) a. Most ravens are black.
b. All tigers have stripes.
c. Some occasions of Mary's jogging are in the park.

Some researchers (de Swart 1991; Chierchia 1992) have proposed that frequency statements, are, in fact, equivalent to sentences such as (36); they express simple quantification. However, a problem with this approach is that frequency statements, like generics but unlike the sentences in (36), are lawlike. For example, the truth of (37), just like the generic (3.j), requires more than simply that the current Supreme Court judges have a prime Social Security number.

(37) Supreme Court judges always have a prime Social Security number.

Moreover, frequency statement, just like generics, support counterfactuals. The truth of (35.b), for example, entails the counterfactual

(38) If Simba were a tiger, he would have stripes.

An alternative is to treat frequency statements as just another kind of generic. As Carlson (1995) points out, this is problematic for the rules and regulations approach. While we may expect that there is a (genetic) rule making ravens black, it is hard to accept a rule that states that *most* of them are; while there may be a rule of Mary's behavior that makes her jog in the park, it is hard to imagine a rule that says, in effect: "Mary, jog in the park sometimes!"

Not all versions of the inductivist view fare better. As we have seen, some of them, being extensional, fail to account for the lawlike nature of generics, and hence cannot account for the lawlikeness of frequency adverbs either.

The normality approach, if applied to generics, faces a different problem. If frequency adverbs, just like generics, quantify over normal individuals only, (39) would be (wrongly) predicted false, since, by hypothesis, all normal ravens are black.

(39) Ravens are sometimes white.

Other inductivist approaches, which take generics to express some quantification over possible individuals, appear to have better prospects for a uniform account of generics and frequency adverbs. The generic quantifier can be taken to be just another frequency adverb, with the semantics of *generally*, *usually*, or something of the sort.

The situation is more complicated, however. There is a difference between generics and frequency adverbs that needs to be commented upon. Sentences (3.g) and (3.h), although bad as generics, become perfectly fine (and true) if the frequency adverb *generally* (or *usually* and the like) is inserted:

(40) a. Primary school teachers are generally female.

b. People are generally over three years old.

Therefore, the interpretation of generics, though similar to that of some adverbs of quantification, cannot be identical to it.

Cohen (1999) proposes that the generics presuppose their domain to be homogeneous, in the following sense. The generic (34) requires that the property ϕ hold not only for ψs, but over every psychologically salient subset of ψ. For example, assuming that it is salient to partition the domain of teachers according to sex, (3.g) requires that both male and female teachers be female—a requirement that is clearly violated. Similarly, assuming that a partition of people according to age is salient, (3.h) requires that people of all ages be over three years old, hence it is not true.

In contrast, frequency adverbs do not require homogeneity. Sentence (40.a) only requires that the property of being female hold of the domain of teachers as a whole, which it does, since the vast majority of primary school teachers *are* female. Similarly, (40.b) requires merely that the property of being over three years old hold, in general, of people as a whole, which it does.

5 Manifestations of generics

No known language contains a specific construction which is exclusively devoted to the expression of genericity (Dahl 1995). Yet there is no language that does not express genericity in some form or another. It follows that expressions used for generics have a double nature: they have generic as well as nongeneric uses. Of particular interest are the forms of noun phrases that may be given generic interpretation. In English, generic noun phrases are may be bare plurals, definite singulars or indefinite singulars (and in some marked cases, definite plurals). It turns out that there are differences in the generic interpretations of these constructions; let us look at each one of them in turn.

5.1 Bare plurals

The most common way to express a generic sentences in English is with a bare plural, i.e. a plural noun preceded by no determiner. It is well known that bare plurals may receive not only a generic reading, but an existential one as well. Thus, while (41.a) makes a generalization about plumbers in general, (41.b) states that there are some plumbers who are available.

(41) a. Plumbers are intelligent.

b. Plumbers are available.

There has been much research on the conditions that determine when a bare plural is interpreted generically, and when existentially

(Carlson 1977; Diesing 1992; Chierchia 1995; Kratzer 1995; Dobrovie-Sorin and Laca 1996; Cohen and Erteschik-Shir 1997; de Smet 1997; Kiss 1998; McNally 1998; Jäger 1999, among others). In this section we will concentrate on the generic interpretation only.

What is the denotation of a generically interpreted bare plural? There are cases where the answer appears to be simple. Consider this typical example:

(42) Dinosaurs are extinct.

There is no individual dinosaur that is extinct; individual dinosaurs are just not the sort of thing that can be extinct—only the kind *dinosaur* can have this property. A natural account for (42) is that it predicates the property of being extinct directly of the kind *dinosaur*. It follows, then, that the bare plural *dinosaurs* denotes this kind in (42).

Krifka *et al.* (1995) refer to such sentences, which predicate a property directly of a kind, as cases of *direct kind predication*. They distinguish between them and sentences such as (1.a), which predicate a property of instances of a kind, and not of the kind as a whole; these are named *characterizing generics*.

One test for cases of direct kind predication is to verify that it is impossible to modify the sentence by an overt adverb of quantification. For example, (43) is bad, confirming that (42) is a case of direct kind predication:

(43) $$\text{*Dinosaurs are} \left\{ \begin{array}{l} \text{always} \\ \text{usually} \\ \text{sometimes} \\ \text{never} \end{array} \right\} \text{extinct.}$$

On the other hand, (44) is fine, indicating that (1.a) is, indeed, a characterizing generic.

(44) $$\text{Ravens are} \left\{ \begin{array}{l} \text{always} \\ \text{usually} \\ \text{sometimes} \\ \text{never} \end{array} \right\} \text{black.}$$

Another test involves scope: characterizing generics, but not direct kind predication, display scope ambiguities. For example, the characterizing generic (45.a) may mean either that each stork has a (possibly different) favorite nesting area, or that there is one nesting area favored by storks. In contrast, (45.b) can only mean that there is one predator endangering the species.

(45) a. Storks have a favorite nesting area (Schubert and Pelletier 1987).

b. Storks are in danger of being exterminated by a predator.

What is the denotation of a bare plural in a characterizing generic? Some researchers (e.g. Wilkinson 1991; Diesing 1992; Kratzer 1995) claim that bare plurals are ambiguous: they may denote kinds, in which case we get direct kind predication, or they may be interpreted as indefinites, i.e. as variables ready to be bound by the generic quantifier, resulting in characterizing generics.

There are, however, reasons to believe that generic bare plurals uniformly refer to kinds, in characterizing generics as well as in cases of direct kind predication. Consider the case of a bare plural that serves as the subject of two clauses: one a characterzing generic and one exprsssing direct kind predication:

(46) a. Dodos lived in Mauritius and (they) became extinct in the 18th century (after Heyer 1990).

b. Elephants are killed for their tusks and are therefore an endangered species.

c. Dinosaurs, which are now extinct, were very large.

The most straightforward explanation for the phenomena exemplified by the sentences in (46) is that a generic bare plural unambiguously refers to kinds.

Moreover, Carlson (1977; 1982) points out that generic bare plurals behave in a way that is similar to referring expressions, rather than indefinites or quantifiers. His arguments apply equally well to characterizing generics and direct kind predication. For example, he note that if the antecedent of a pronoun is a name, it can replace the pronoun without a change in meaning; not so, in general, when the antecedent is an indefinite. Generics seem to behave like names, rather than indefinite, in this regard:

(47) a. Fred walked into the room. *He* smiled (= *Fred* smiled).

b. A man walked into the room. *He* smiled ($\neq$ *A man* smiled).

c. Dogs are intelligent mammals. *They* are also man's best friend (= *Dogs* are man's best friend).

An additional observation, ascribed by Carlson to Postal (1969), is that names and generics, and only names and generics, can participate in *so-called* constructions:

(48) a. Giorgione is so-called because of his size.

b. Machine guns are so-called because they fire automatically.

c. *A machine gun is so-called because it fires automatically.

Krifka *et al.* (1995) agree that bare plurals refer to kinds in characterizing generics too, but restrict this only to "well-established kinds." We will discuss this issue further in the next section.

If bare plurals in characterizing generics denote kinds, a natural question arises: how is a characterizing generic obtained from a kind-denoting bare plural? In order to answer this question, Carlson (1977) proposes a *realization* relation between an instance and a kind. Thus, for example, $R(x, \mathbf{dog})$ indicates that x is an instance of the kind *dog*, i.e. x is a dog.

Ter Meulen (1995) proposes a type-shifting operator, which transforms a kind into the property of being an instance of the kind. The application of this type-shifting operator is optional. When it is applied, the result is a characterizing generic; when it is not—direct kind predication. Thus every generic sentence is ambiguous between characterizing and kind interpretations; but one of these readings is ruled out as semantically anomalous. For example, (1.a) has a reading where the property of being black is predicated directly of the kind *ravens*. But this reading is ruled out, because a kind is not the sort of thing that can have a color. Similarly, (42) has a reading where individual dinosaurs are extinct. This time, the characterizing interpretation will be ruled out, since individual dinosaurs cannot be extinct.

When interpreted generically, bare plurals may receive collective readings, e.g. they may be the arguments of predicates such as the intransitive *meet* and *gather*. Consider the following example (attributed by Krifka *et al.* to Gerstner 1979):

(49) Lions gather near acacia trees when they are tired.

To account for this fact, Krifka *et al.* (1995) propose that groups of individuals are also individuals in their own right (see e.g. Link 1983; Ojeda 1993), and that, therefore, they can be instances of a kind just like single individuals can. Thus, (49) predicates the property of gathering on groups of lions, rather than individual lions. If sufficiently many such groups gather, (49) is true.

There is reason to believe, however, that groups of individuals are not always considered individuals. Consider generics with a distributive property, e.g.

(50) Lions have a bushy tail.

According to Krifka *et al.* (1995), (50) is true just in case sufficiently many groups of lions have a bushy tail. The problem is that when a distributive property such as *have a bushy tail* is applied to a group, it needs to hold of all members of a group. For example, (51) means that each one of the lions in the cage has a bushy tail.

(51) The lions in the cage have a bushy tail.

Now suppose that only two of all lions lack a bushy tail. Given that the number of lions is fairly large, sentence (50) ought to be true. However, it can be easily seen that only a quarter of all possible groups of lions contain neither of the “deficient” lions. If we grant that a quarter of all groups of lions is not sufficiently many, (50) would be predicted false.

Sentence (50), for example, would cease to be a problem if we assume that groups are simply not allowed as instances in this case, but only simple individuals are. And, in general, when the predicated property is distributive, group instances are not considered in evaluating the truth of the generic sentence.

Mixed predicates, i.e. predicates that allow both distributive and collective readings, pose a problem too. Consider the following example (from Krifka 1987):

(52) The German customer bought 83,000 BMWs last year.

The most plausible interpretation of (52) is that the total number of BMWs bought by Germans last year is 83,000; in other words, the number of BMWs bought collectively by the group of all German customers is 83,000. No other group of German customers bought 83,000 BMWs, yet it seems that only the maximal group of German customers is relevant to the truth or falsity of (52). Suppose the people living in the former West Germany bought 83,000 BMWs, and the customers living in what used to be East Germany bought 83,000 BMWs too. In this case, there would be two groups of German customers, each of which satisfies the predicate denoted by the VP, yet (52) would be false. It seems, then, that not all possible individuals and groups of individuals should be taken into account as instances of a kind. Moreover, what counts as an instance of a kind may vary across sentences, and is probably affected by the context (Cohen 1996).

5.2 Definite singulars

Just like bare plurals, definite singular generics may occur in cases of direct kind predication as well as characterizing generics, as exemplified by the following:

(53) a. The giant panda eats bamboo shoots.

b. The giant panda is an endangered species.

Sentence (53.a) is about individual pandas, whereas (53.b) is about the kind *giant panda* as a whole.

We therefore expect definite singulars to have collective readings, as the following examples indicate:

(54) a. The lion gathers near acacia trees when it is tired (Gerstner 1979; Krifka *et al.* 1995)

b. The antelope gathers near water holes (Heyer 1990).

Are the sentences in (54) indeed acceptable, as the cited sources maintain? Some informants agree, but others judge them to be marginal. This marginality, however, may be due to other reasons, perhaps the number feature of the verb; intransitive *gather* is not normally used in the singular. When the collective verb is predicated of a conjunction of definite singular generics, so that the number of the verb is plural, the acceptability of the sentence improves markedly:

(55) Two species of cats, the lion and the leopard, gather near acacia trees when they are tired.

A noncontrived, naturally occurring example is the following sentence, taken from the entry for *shark* in the American Academic Encyclopedia:

(56) Some sharks, such as the tiger shark and the great white, are loners and seemingly swim at random, although they sometimes gather to feed.

We mentioned above that Krifka *et al.* (1995) claim that bare plurals may only refer to well established kinds. They reach this conclusion by comparing the distribution of generic definite singulars with that of bare plurals: they find that the distribution of the former is much more restricted. Compare the acceptability of (57.a) with the oddness of (57.b) (an example which Carlson 1977 ascribes to Barbara Partee).

(57) a. The Coke bottle has a narrow neck.

b. ?The green bottle has a narrow neck.

Krifka *et al.*'s account of this fact is as follows. Definite singulars must refer to a kind in order to be interpreted generically. The kind *Coke bottle* is well established in our culture, hence the reference succeeds and (57.a) is interpreted generically. The kind *green bottle*, on the other hand, is not well established, hence the reference fails and (57.b) cannot be interpreted generically (it is, of course, fine under the nongeneric reading). In contrast, both (58.a) and (58.b) are fine.

(58) a. Coke bottles have narrow necks.
  b. Green bottles have narrow necks.

The acceptability of the sentences in (58) is explained by the claim that bare plurals do not always refer to kinds. The subject of (58.a) denotes the kind *Coke bottle*, but the subject of (58.b) does not refer to any kind—it is interpreted as a variable.

While the distribution of definite generics is, indeed, restricted, it is not clear that the facts about this distribution can be explained in terms of well established kinds. The acceptability of the definite generic seems to depend on a variety of factors (see Vendler 1971; Carlson 1977; Bolinger 1980; Dayal 1992, among others). For example, the definite generic is often more acceptable when the descriptive content of the common noun is richer. Contrast the oddness of (59.a) (under the generic reading) with the acceptability of (59.b).

(59) a. ?The politician never misses a photo opportunity.
  b. The successful politician never misses a photo opportunity.

Yet one would be hard pressed to argue that *successful politician* is a well-established kind, whereas *politician* is not.

There are additional, poorly understood factors affecting the productivity of the definite generic, which appear idiosyncratic and language dependent. Contrast (60.a), which is fine, with (60.b), which is odd (under the generic reading).

(60) a. The tiger lives in the jungle.
  b. ?The dog barks.

Yet there is no reason to suppose that the kind *tiger* is better established than the kind *dog*. The distinction seems to be an idiosyncratic property of English; indeed, there are languages where the equivalent of (60.b) is perfectly acceptable, e.g. German:

(61) Der Hund bellt (Heyer 1990).

5.3 Indefinite singulars

Unlike bare plurals and definite singulars, indefinite singulars may not refer to kinds, as the unacceptability of the following examples indicate:

(62) a. *A giant panda is an endangered species.
  b. *A dinosaur is extinct.

There is, in fact, a reading under which these sentences are acceptable, the *taxonomic* reading, according to which some subspecies of giant panda is endangered, or some species of dinosaurs is extinct. Under this reading, however, the subject is interpreted existentially, rather than generically, with the existential quantifier ranging over kinds. Therefore, this reading need not concern us here.

If indefinite singulars may not refer to kinds, we can predict that collective readings are impossible. This is, indeed, borne out:

(63) *A lion gathers near acacia trees when it is tired.

The distribution of the indefinite singular is restricted compared with that of the bare plural, but in ways that are different from those of the definite singular. Consider the following pair (Lawler 1973):

(64) a. A madrigal is polyphonic.
b. *A madrigal is popular.

While (64.a) receives a generic interpretation, (64.b) cannot. In contrast, both (65.a) and (65.b) are fine.

(65) a. Madrigals are polyphonic.
b. Madrigals are popular.

Burton-Roberts (1977) provides a number of additional examples, among which are the following:

(66) a. Kings are generous.
b. *A king is generous.

(67) a. Rooms are square.
b. *A room is square.

(68) a. Uncles are garrulous.
b. *An uncle is garrulous.

Lawler (1973) claims that this difference between bare plural and indefinite singular generics is due to the fact that the latter are restricted to properties that are, in some sense, "necessary," "essential," "inherent," or "analytic." Thus, whereas polyphonicity is an essential property of madrigals, popularity is not, hence the unacceptability of (64.b).

The problems with this approach is that it falls short of a complete explanation: why is it indefinite singulars, rather than bare plurals or definite singulars, that are singulars have this property. Moreover, it fails to account for sentences such as the following:

(69) A madrigal is a popular song.

Although (69) it seems to be saying exactly the same as (64.b), it is perfectly acceptable.

Krifka *et al.* (1995) propose an account of this phenomenon, based on the fact that indefinite singulars may not refer to kinds. They suggest that all cases where the indefinite singular generic is disallowed are cases of direct kind predication. That is to say, just like (42) expresses a property directly of the kind *dinosaur*, and not of individual dinosaurs, (65.b) expresses a property directly of the kind *madrigal*. Specifically, unlike (65.a), the logical form of (65.b) does not involve the generic quantifier. Since indefinite singulars cannot occur in cases of direct kind predication, (64.b) is ruled out.

This approach amounts to disposing with the quantificational account of genericity except for a small number of cases such as (65.a). It follows that characterizing generics are, in fact, the exception, rather than the rule.

However, it is not clear that the claim that (65.b) is a case of direct kind predication can be maintained. If we apply the relevant tests, it appears that these are cases of characterizing generics rather than direct kind predication: the sentences in (70) are grammatical, and (71) exhibits a scope ambiguity.

(70) a. Madrigals are always popular.

b. Kings are usually generous.

c. Rooms are sometimes square.

d. Uncles are never garrulous.

(71) Madrigals are popular with exactly one music fan.

Burton-Roberts (1977) proposes that indefinite singulars carry a normative force. He considers the following minimal pair:

(72) a. Gentlemen open doors for ladies.

b. A gentleman opens doors for ladies.

Burton-Roberts notes that (72.b), but not (72.a), expresses what he calls "moral necessity." Burton-Roberts observes that

> if Emile does not as a rule open doors for ladies, his mother could utter [(72.b)] and thereby successfully imply that Emile was not, or was not being, a gentleman. Notice that, if she were to utter… [(72.a)] she might achieve the same effect (that of getting Emile to open doors for ladies) but would do so by different means…For [(72.a)] merely makes a generalisation about gentlemen (p. 188).

Sentence (72.b), then, unlike (72.a), does not have a reading where it makes a generalization about gentlemen; it is, rather, a statement about some social norm. It is true just in case this norm is in effect, i.e. it is a member of a set of socially accepted rules and regulations.

We have seen above (section 2.3) that Greenberg (1998) and Cohen (forthcoming) propose that generic bare plurals are ambiguous: they may express a characterizing generic, amenable to some sort of inductivist treatment, or they may express a rule, amenable to a treatment within the framework of the rules and regulations view. In contrast, indefinite singulars are not ambiguous: they only express rules. Thus, given the scenario with the supermarket described in section 2.3, only (73.b) is true:

(73) a. A banana sells for \$.49/lb.

b. A banana sells for \$1.00/lb.

The rule may be a linguistic rule, i.e. a definition. Since polyphonicity forms a part of the definition of a madrigal, (64.a) is fine. The acceptability of (69) stems from the fact that it has the classical *form* of a definition, even though it is not, in fact, the approved definition of a madrigal.

6 The use of generics

If generics are, indeed, so prevalent, a natural question arises: what are they good for? Why do we use them as often as we do? One possible answer is that generics are used to state default rules.

Our beliefs about the world are almost never certain. In most cases, the conclusions we draw are plausible, but not guaranteed to be true. For example, when I turn the ignition key in my car, I expect it to start. I do not know for certain that it will; sometimes there is some malfunction, and the car fails to start. But it is a reasonable assumption that the car will start, an assumption I am ready to retract if I find out that this is not the case. It is not irrational to assume that the car will start although I do not have complete confidence in it; quite the reverse. The alternative would be to subject the car to a comprehensive inspection by a mechanic every time I am about to start it—clearly an impractical solution, and, in fact, an unnecessary one.

Rules of inference that allow us, for instance, to conclude that the car will start without actually establishing it conclusively, are usually referred to as *default rules*. The most important property that distinguishes them from classical logical rules of inference is that they are *nonmonotonic*: the conclusions may be retracted given additional information. For example, if I see that the car's battery has been stolen, I will no longer expect the car to start. Not so for classical logical rules: if we know that all men are mortal and that Socrates is a man, we conclude that Socrates is mortal, and no amount of additional information will invalidate this conclusion.

There is a sizeable body of research on nonmonotonic reasoning. See, for isntance, Ginsberg (1987) for a classic collection of papers. Of particular relevance to our concern here is the fact that, when researchers discuss a default rule, they often characterize it, informally, in natural language; and they usually use a generic to do this. It is, therefore, an appealing idea that the use of generics is often to express default rules. We can say that one utters (1.a) in order to express the following default rule: if we know that an individual is a raven, we should conclude, by default, that it is black.

We have seen above that Krifka (1987) proposes that the meaning of a generic is a default rule. But one need not be committed to the claim that the *meaning* is a default rule, to propose that the *use* is that of stating a default rule. What one does need to be committed to is that the meaning of generics supports the conclusions that follow from an appropriate system of default rules.

The problem is that there is little consensus on which inferences are sound and which ones are not. For example, ordinary (monotonic) entailment is transitive. If we believe that A entails B, and B entails C, we can conclude that A entails C. But what about nonmonotonic inference? Sometimes such a conclusion appears valid. Suppose we believe the following:

(74) a. Tweety is a robin.
b. Robins are birds.
c. Birds fly.

We seem justified in concluding, on the basis of this, that Tweety flies.

But now suppose we believe the following:

(75) a. Tweety is a penguin.

b. Penguins are birds.

c. Birds fly.

Are we justified in concluding that Tweety flies? Intuitively, the answer is no. This is because we also believe another rule:

(76) Penguins don't fly.

Somehow, (76) ought to take precedence over (75.c), and thus we should conclude that Tweety does not fly. In a sense, (76) is more specific than (75.c), and it appears that this is the reason why it overrides it. There have been a number of attempts to formalize this notion of specificity, and to give a semantics for generics that can support it (see, among others, Etherington and Reiter 1983; Brewka 1991; Morreau 1992; Cohen 1997; Pelletier and Asher 1997).

Suppose we believe (77.a) and (77.b). We are surely justified in concluding (78) (Pearl 1988).

(77) a. Red birds fly.

b. Non-red birds fly.

(78) Birds fly.

Any theory of generics of the meaning of generics ought to be able to account for this trivial inference.

But what about inference in the opposite direction? Can we conclude either (77.a) or (77.b) from (78)? And, in general, can we conclude (79.b) from (79.a)?

(79) a. ψs are ϕ.

b. ψs that are χ are ϕ.

The general answer appears to be no. For example, we certainly do not conclude (80) from (78).

(80) Dead birds fly.

In order to account for this fact, Pearl (1988) proposes that we can conclude (79.b) from (79.a) only if we have a rule stating

(81) ψs are χ.

For example, we can conclude (82.b) from (78), because we also have the rule (82.a).

(82) a. Birds lay eggs.

b. Birds that lay eggs fly.

Thus (80) does not follow, because we do not have a rule saying

(83) Birds are dead.

A problem with this requirement is that it appears to be too strong, blocking desirable inferences. For example, the inference from (78) to (77.a) is blocked, whereas it appears that this is a conclusion worth having.

Pelletier and Asher (1997) go to the other extreme: they propose that concluding (79.b) from (79.a) is always licensed, unless we already have a rule stating

(84) ψs that are χ are not ϕ.

Thus, (80) does not follow, because we presumably have a rule stating

(85) Dead birds don't fly.

In contrast with Pearl's approach, Pelletier and Asher's view appears too liberal, allowing inferences that appear not to be warranted. For example, we should not, intuitively, conclude either (86.a) or (86.b) from (78)—whether a sick bird flies or not depends on the type and severity of the disease.

(86) a. Sick birds fly.

b. Sick birds don't fly.

Yet, according to Pelletier and Asher, if we cannot conclude (86.b), we should be able to conclude (86.a).

An alternative approach is to allow the derivation of (79.b) from (79.a) only if the property ϕ is independent of the property χ. Thus, (77.a) follows from (78), because the ability to fly is independent of one's color; but (86.a) does not follow, because being able to fly is not independent of one's health. The problem with this approach is that it is not easy to specify an appropriate notion of independence; but see Shastri (1989); Bacchus (1990); Bacchus *et al.* (1993); Cohen (1997) for proposals.

Another question is how to treat exceptions to rules. In most work on default reasoning, being an exception to one rule does not affect the applicability of other rules. For example, suppose we have the following rules:

(87) a. Mammals bear live young.

b. Mammals have hair.

Suppose Pat is a platypus, so she violates (87.a), by not bearing live young. Are we still justified in applying (87.b) to conclude that she has hair? The answer appears to be yes.

There are cases, however, where this strategy leads to wrong conclusions. Suppose we have the following additional rule:

(88) Mammals have a uterus.

Now it appears that we are not allowed to apply (88) and to conclude that Pat has a uterus.

Theories of generics that interpret them as expressions of quantification over normal individuals would block the conclusion that Pat has a uterus (see, in particular, Pelletier and Asher 1997). This is so because by failing to satisfy (87.a), Pat has shown herself to be an abnormal mammal, and hence other rules should not apply to her either. The problem is that (87.b) will not follow either. Pelletier and Asher propose to add such a conclusion as a special case, but a more general approach is probably desirable.

Theories that make use of the notion of independence, on the other hand, have better prospects of accounting for such cases. Rule (87.b) is applicable, because the property of bearing live young is independent of the property of having hair. In contrast, bearing live young is not independent of the property of having a uterus, hence we are not justified in concluding that Pat has a uterus.

We have seen that the rules governing our default reasoning can be seen 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" (p. 203), we may safely conclude that when formalizing our common-sense intuitions, it is beneficial to look closely at the language we use to express them.

References

Asher, N., and Morreau, M. 1995. What some generic sentences mean. In Carlson and Pelletier (1995) 300–338.

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

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.

Bacon, J. 1974. Do generic descriptions denote? *Mind* 82.331–347.

Bartsch, R. 1972. The semantics and syntax of number and numbers. In *Syntax and Semantics*, ed. by J. Kimball volume 2 51–94. New York: Seminar Press.

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

Barwise, J., and Perry, J. 1983. *Situations and Attitudes*. Cambridge, Massachusetts: MIT Press.

Barwise, J., and Seligman, J. 1992. The rights and wrongs of natural regularity. In *Philosophical Perspectives*, ed. by J. Tomberlin.

Bennett, M. 1974. *Some Extensions of a Montague Fragment of English*. University of California at Los Angeles dissertation.

Bolinger, D. 1980. *Syntactic Diffusion and the Indefinite Article*. Bloomington: Indiana University Linguistics Club.

Brewka, G. 1991. Cumulative Default Logic: In defense of nonmonotonic inference rules. *Artificial Intelligence* 50.183–205.

Burton-Roberts, N. 1977. Generic sentences and analyticity. *Studies in Language* 1.155–196.

Carlier, A. 1989. Généricité du syntagme nominal sujet et modalités. *Travaux de Linguistique* 19.33–56.

Carlson, G. 1977. *Reference to Kinds in English*. University of Massachusetts at Amherst dissertation. Also published 1980, New York: Garland.

—— 1982. Generic terms and generic sentences. *Journal of Philosophical Logic* 11.145–181.

—— 1989. On the semantic composition of English generic sentences. In Chierchia *et al.* (1989) 167–192.

—— 1995. Truth-conditions of generic sentences: Two contrasting views. In Carlson and Pelletier (1995) 224–237.

Carlson, G., and Pelletier, F. J. (eds.) 1995. *The Generic Book*. Chicago: University of Chicago Press.

Cavedon, L., and Glasbey, S. 1994. Outline of an information-flow model of generics. *Acta Linguistica Hungarica* 42.

Chierchia, G. 1992. Anaphora and dynamic binding. *Linguistics and Philosophy* 15.111–183.

—— 1995. Individual-level predicates as inherent generics. In Carlson and Pelletier (1995) 176–223.

Chierchia, G., Partee, B. H., and Turner, R. (eds.) 1989. *Properties, Types and Meaning*. Dordrecht: Kluwer.

Cohen, A. 1996. *Think Generic: The Meaning and Use of Generic Sentences*. Carnegie Mellon University dissertation. Published 1999, Stanford: CSLI.

—— 1997. Generics and default reasoning. *Computational Intelligence* 13.506–533.

—— 1999. Generics, frequency adverbs and probability. *Linguistics and Philosophy* 22.221–253.

—— 2000. Is there an unambiguous level of representation? In *Interface Strategies*, ed. by H. Bennis, M. Everaert, and E. Reuland. Amsterdam: Royal Netherlands Academy of Arts and Sciences.

—— forthcoming. On the generic use of indefinite singulars. In *Uses of Indefinite Expressions*, ed. by C. Piñon and P. Dekker. Stanford: CSLI.

Cohen, A., and Erteschik-Shir, N. 1997. Topic, focus and the interpretation of bare plurals. In *Proceedings of the 11th Amsterdam Colloquium*, ed. by P. Dekker, M. Stokhof, and Y. Venema 31–36.

Dahl, Ö. 1975. On generics. In *Formal Semantics of Natural Language*, ed. by E. L. Keenan 99–111. Cambridge: Cambridge University Press.

—— 1995. The marking of the episodic/generic distinction in tense-aspect systems. In Carlson and Pelletier (1995) 412–425.

Dayal, V. S. 1992. The singular-plural distinction in Hindi generics. In *Proceedings of the Second Conference on Semantics and Linguistic Theory*, ed. by C. Barker and D. Dowty number 40 in Working Papers in Linguistics 39–57. Ohio State University.

Declerk, R. 1986. The manifold interpretations of generic sentences. *Lingua* 68.149–188.

—— 1991. The origins of genericity. *Linguistics* 29.79–101.

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

DIESING, M. 1992. *Indefinites*. Cambridge, Massachusetts: MIT Press.

DOBROVIE-SORIN, C., and LACA, B. 1996. Generic bare NPs. Unpublished manuscript.

ETHERINGTON, D. W., and REITER, R. 1983. On inheritance hierarchies with exceptions. In *AAAI 83*.

FARKAS, D., and SUGIOKA, Y. 1983. Restrictive if/when clauses. *Linguistics and Philosophy* 6.225–258.

GERSTNER, C. 1979. Über generizität. Master's thesis University of Munich.

GEURTS, B. 1985. Generics. *Journal of Semantics* 4.

GINSBERG, M. L. (ed.) 1987. *Readings in Nonmonotonic Reasoning*. San Mateo, CA: Morgan Kaufmann.

GREENBERG, Y. 1998. Temporally restricted generics. In *Proceedings of SALT VIII*.

GRICE, H. P. 1975. Logic and conversation. In *Syntax and Semantics 3: Speech Acts*, ed. by P. Cole and J. L. Morgan. Academic Press.

HEYER, G. 1985. Generic descriptions, default reasoning, and typicality. *Thoretical Linguistics* 12.33–72.

—— 1990. Semantics and knowledge representation in the analysis of generic descriptions. *Journal of Semantics* 7.93–110.

JÄGER, G. 1999. Stage levels, states, and the semantics of the copula. In *ZAS Papers in Linguistics* volume 191 65–94. Zentrum für Allgemeinee Sprachwissenschaft.

KISS, K. É. 1998. On generic and existential bare plurals and the classification of predicates. In Rothstein (1998) 145–162.

KONINGSVELD, H. 1973. *Empirical Laws, Regularity and Necessity*. Wageningen, The Netherlands: H. Veenman & Zonen B.V.

KRATZER, A. 1981. The notional category of modality. In *Words, Worlds, and Contexts: New Approaches to Word Semantics*, ed. by H. J. Eikmeyer and H. Rieser 38–74. Berlin: de Gruyter.

—— 1995. Stage-level and individual-level predicates. In Carlson and Pelletier (1995) 125–175.

KRIFKA, M. 1987. An outline of genericity. Technical Report SNS-Bericht 87–25 Seminar für natürlich-sprachliche Systeme, Tübingen University Germany.

—— 1995. Focus and the interpretation of generic sentences. In Carlson and Pelletier (1995) 238–264.

Krifka, M., Pelletier, F. J., Carlson, G., ter Meulen, A., Link, G., and Chierchia, G. 1995. Genericity: an introduction. In Carlson and Pelletier (1995) 1–124.

Landman, F. 1989. Groups, II. *Linguistics and Philosophy* 12.723–744.

Lawler, J. 1973. Studies in English generics. University of Michigan Papers in Linguistics 1:1.

Link, G. 1983. The logical analysis of plural and mass terms: A lattice-theoretical approach. In *Meaning, Use and Interpretation of Language*, ed. by R. Bäuerle, C. Schwartze, and A. von Stechow 302–323. Berlin: de Gruyter.

—— 1995. Generic information and dependent generics. In Carlson and Pelletier (1995) 358–382.

McCarthy, J. 1986. Applications of circumscription to formalizing commonsense knowledge. *Artificial Intelligence* 28.89–116.

McNally, L. 1998. Stativity and theticity. In Rothstein (1998) 293–307.

ter Meulen, A. 1986. Generic information, conditional contexts and constraints. In *On Conditionals*, ed. by E. C. Traugott, A. ter Meulen, J. S. Reilly, and C. A. Ferguson. Cambridge: Cambridge University Press.

—— 1995. Semantic constraints on type-shifting anaphora. In Carlson and Pelletier (1995) 339–357.

Morreau, M. P. 1992. *Conditionals in Philosophy and Artificial Intelligence*. University of Amsterdam dissertation.

Nunberg, G., and Pan, C. 1975. Inferring quantification in generic sentences. In *Papers from the Eleventh Regional Meeting of the Chicago Linguistic Society*.

Ojeda, A. 1993. *Linguistic Individuals*. Stanford: Center for the Study of Language and Information.

Parsons, T. 1970. An analysis of mass terms and amount terms. *Foundations of Language* 6.363–388.

Partee, B. H. 1991. Topic, focus and quantification. In *Proceedings of the First Conference on Semantics and Linguistic Theory*, ed. by S. Moore and A. Z. Wyner 159–187. Cornell University.

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

Pelletier, F. J., and Asher, N. 1997. Generics and defaults. In *Handbook of Logic and Language*, ed. by J. van Benthem and A. ter Meulen 1125–1177. Amsterdam: Elsevier.

Postal, P. 1969. Anaphoric islands. In *Papers from the Fifth Regional Meeting of the Chicago Linguistic Society* 205–239.

QUINE, W. V. O. 1960. *Word and Object*. Cambridge, Massachusetts: MIT Press.

REITER, R. 1987. Nonmonotonic reasoning. In *Annual Review of Computer Science*, ed. by J. F. Traub, N. J. Nilsson, and B. J. Grosz 147–186. Palo ALto: Annual Reviews Inc.

ROOTH, M. E. 1985. *Association with Focus*. University of Massachusetts at Amherst dissertation.

ROTHSTEIN, S. (ed.) 1998. *Events and Grammar*. Dordrecht: Kluwer.

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.

—— 1989. Generically speaking, or using discourse representation theory to interpret generics. In Chierchia *et al.* (1989) 193–268.

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

DE SMET, L. B. 1997. *On Mass and Plural Quantification*. University of Groningen dissertation.

STRZALKOWSKI, T. 1988. A meaning representation for generic sentences. Technical Report 423 Courant Institute of Mathematical Sciences, New York University.

DE SWART, H. 1991. *Adverbs of Quantification: A Generalized Quantifier Approach*. Groningen University dissertation. Also published 1993, New York: Garland.

—— 1996. (In)definites and genericity. In *Quantifiers, Deduction and Context*, ed. by M. Kanazawa, C. Piñon, and H. de Swart. Stanford: Center for the Study of Language and Information.

VENDLER, Z. 1971. Singular terms. In *Semantics: An Interdisciplinary Reader in Philosophy, Linguistics and Psychology*, ed. by D. D. Steinberg and L. A. Jakobovits. Cambridge: Cambridge University Press.

WILKINSON, K. 1991. *Studies in the Semantics of Generic Noun Phrases*. University of Massachusetts at Amherst dissertation.