

*On generics**

ANNA PAPAFRAGOU

Abstract

In this paper I argue against previous approaches to the semantics of generics which involved the notions of prototype, stereotype and relevant quantification. I assume that the logical form of generics includes a generic operator which, as Heim (1992) has suggested, can be construed as the modal operator of necessity. After demonstrating that the presence of the generic operator in a semantic representation, as well as its domain of quantification, are pragmatically supplied, I go on to show how the various interpretations generics may receive can be successfully accounted for within a relevance-theoretic framework.

1 Introduction

There have been two types of phenomena traditionally classified as 'generic' in philosophy and linguistics. The first involves reference to a kind, as exemplified in (1) and (2); the subject NPs do not refer to a particular orchid or group of orchids but rather to the kind Orchid itself:

- (1) Orchids are flowers.
- (2) The orchid is a flower.

The second involves propositions which do not describe specific episodes or isolated facts but instead report a regularity that summarises groups of particular episodes or facts. For instance, (3) reports some kind of general property attributed to the earth; it captures some sort of generalisation over particular events. Similarly (4) states something not

* I am grateful to Deirdre Wilson for her insightful comments and continuous support. I also thank Neil Smith, Paul Horwich, Giannis Veloudis and Villy Rouchota for offering me stimulating discussions on generics and comments on earlier versions of this paper. Finally, I would like to thank the State Scholarships Foundation in Greece for financially supporting my postgraduate studies at UCL.

about an individual elephant but rather about elephants in general — again a generalisation based on properties of specific members of the set:

(3) The earth turns around the sun.

(4) An elephant eats grass.

Following Krifka et al. (1995), I will call the subject NPs in (1) and (2) *kind-referring* (or *generic*) NPs, as opposed to *object-referring* NPs, and the predications involved in such sentences *kind predications*, in opposition to *object predications*. I will call sentences like (3) and (4) *characterising sentences* (or simply *generic sentences*); these will be opposed to *particular sentences*, which express statements about particular events, properties of particular objects, and the like. The respective predications will be classified as *characterising* (as opposed to *particular*) predications. Other common names for characterising sentences are 'habitual', 'dispositional' or '(g)nomical' (vs. 'episodic') sentences.¹

Obviously reference to kinds and characterising sentences have something in common: in both cases we abstract away from particulars, be they members of a kind or specific events. I will make no attempt here to reduce one to the other, since it turns out that there are significant linguistic differences among them and therefore it might be better to keep them separated (see Krifka et al. 1995). Instead, I would like to focus on a more general problem: namely, what is the nature of genericness itself? Is it better handled as a semantic phenomenon (as most of the linguistic terminology suggests), or as a pragmatic one? Should we analyse kind-referring NPs and characterising sentences as containing some sort of semantic operator in their logical form, or should we instead view genericness as a result of pragmatic processes operating during utterance interpretation?

My aim in this paper is to argue for a treatment of genericness which combines semantic and pragmatic aspects; more specifically, I want to claim that the interpretation of generics follows closely the interpretation of utterances containing modal expressions and should therefore make use of some appropriate version of the notion of possible worlds. My discussion proceeds as follows. In section 2 I critically review some previous attempts to deal with the semantics of genericness. In section 3 I present a framework for interpreting modal expressions and go on to apply it to the interpretation of utterances

¹I will use the terms 'generics' and 'generic propositions/statements' in relation to both characterising sentences and sentences containing kind-referring NPs.

exhibiting genericness; I insist throughout on a sound division of labour between the semantic and pragmatic component within a relevance-theoretic framework for semantics and pragmatics (Sperber & Wilson 1986/21995).

2 Previous approaches to generics

2.1 Preliminaries

When faced with a sentence conveying genericness, an initially plausible tendency is to treat it in terms of universal quantification. On this view, the propositions expressed by (1) and (3) are given by (1') and (3') respectively:

(1') (For all values of x) (if x is an orchid, then x is a flower)

(3') (For all values of t) (the earth turns around the sun at t)

However, as a number of authors have pointed out, universal quantification is actually unsuitable to capture the meaning of generic propositions, since it may prove occasionally too weak or too strong. It appears to be too weak in view of *essential* propositions, i.e. propositions of the sort in (1) which convey that a given property is a necessary and exceptionless attribute of all members of the class to which reference is made (Lyons 1977:195). More generally, generics as a rule express a principled (non-accidental) generalisation over the members of a class of entities or events. Use of the universal quantifier, however, does not distinguish between accidental and non-accidental properties (Dahl 1975, Lyons 1977:195, Carlson 1982:147). For instance, suppose that it is true for all students in London that they are underweight. This makes the universally quantified propositions in (5) and (6) true but does not make the generic propositions in (7) and (8) true — unless some specific, non-accidental relationship is taken to hold between students in London and being underweight:

(5) All students in London are underweight.

(6) Every/each student in London is underweight.

(7) Students in London are underweight.

(8) A student in London is underweight.²

When it comes to other sorts of generics, the universal quantifier is problematic for the opposite reason: it is too strong to capture the intended meaning of the generic proposition (Lawler 1972, Lyons 1977:196). Normally, what an utterance like (9) is taken to convey is that most mathematicians are exciting, not that they are all necessarily so:

(9) Mathematicians are exciting.

The same point applies to characterising sentences. The proposition expressed by the utterance in (10) is not falsified if Claudia walks to school four days out of five in a week:

(10) Claudia goes to school on foot.

So generics allow for exceptions, whereas universally quantified propositions do not. The latter should thus be rejected as a means of representing the former.

To preserve a quantificational semantic approach to generics, there have been proposals for a weakened (or 'modified' — see Clark 1973:43) universal quantifier which would translate as *almost all*, or *the majority of/most*. Closer attention, though, reveals that such a quantifier would still fail to mark the non-contingent character of generic statements, and so would only do half the job. Moreover, as Carlson (1977b) remarked, one can find generic sentences which are considered to be true but which cannot be satisfactorily described by any ordinary quantifier. For instance, the following are true characterising sentences, although less than half of all birds lay eggs (only the healthy and fertilised female ones), not more than five percent of the *Anopheles* mosquitos carry malaria, and the chance of a turtle having a long life is extremely small, as most turtles are eaten by predators early in life (examples from Krifka et al. (1995:44)):

(11) A bird lays eggs.

²It is possible on some occasions for an utterance such as (7) — but not (8) — to be interpreted as an accidental generalisation. On the pragmatic account I will put forth in section 3, this possibility comes out as a subtype of the more general case where the interpretation of generics involves quantification over different types of world: here, quantification would take in its scope just the actual world, so that the proposition expressed in (7) would come out as a description of a contingent fact about all actual students.

(12) An *Anopheles* mosquito carries malaria.

(13) A turtle lives a long life.

A more promising line seems to come from postulating the existence of a (mostly) phonologically unrealised generic operator. This operator functions as a quantificational adverb — an adverbial operator which relates one set of conditions to another set. Such structures have been proposed by a number of authors for representing sentences with conditionals or *when*-clauses (Lewis 1975, Heim 1982, Farkas 1981, Farkas & Sugioka 1983, Schubert & Pelletier 1989). Consider Lewis's (1975) example of a construction containing an explicit adverb of quantification:

(14) When m and n are positive integers, the power m^n can be computed by successive multiplication.

(14') Q (m and n are positive integers; m^n can be computed by successive multiplication)

Q is a quantifier expressed in English by *always*, which relates two propositions. Following Krifka et al. (1995), I will call them the *restrictor* and the *matrix* respectively. A similar, double-level analysis of generic sentences lacking an explicit adverb of quantification can be put forth, in which Q takes the form of GEN, the generic operator. Thus (9) will receive the propositional representation in (9'):³

(9') GEN (x is a mathematician; x is exciting).

In the above example, quantification operates on individuals who form members of a class: the operator GEN binds the variables in x . Characterising sentences are a different matter: here quantification operates over temporally bound occurrences of an event, that is, situations, occasions, or cases.⁴ To accommodate these, we introduce the variable s ,

³A certain caveat is in order here: in (9) Q functions as an unselective quantifier, binding any free variables in its scope. To capture a greater range of examples, we will have to assume that the quantifier indicates which variables it binds and which are to be bound existentially within the matrix. In the following discussion, I will disregard this complication.

⁴These notions are meant to be equivalent to the notion of *stage* introduced by Carlson (1977a).

which stands for situation and is bound by the generic operator.⁵ Thus (10) receives the logical form:

(10') GEN (Claudia goes to school in s; Claudia goes to school on foot in s).

The above solution accounts for the fact that genericness takes sentential scope in the case of characterising sentences. It also captures the intuition that the role of genericness is similar to that of adverbs of quantification such as *always*, *seldom* and *often*. The adverbs which come closest to the meaning of the generic operator are *generally/in general*, *typically*, *characteristically* and *habitually*. Thus (9) and (10) turn out to be semantically equivalent to the pairs (15)-(16) and (17)-(18) respectively:

(15) Generally, mathematicians are exciting.

(16) Generally, when one is a mathematician, one is exciting.

(17) Generally/habitually, Claudia goes to school on foot.

(18) Generally/habitually, when Claudia goes to school, she goes to school on foot.

Apart from its intuitive appeal, this approach manages to capture the quasi-universality of many generics: for (9) and (10) to be true, in the majority of cases to be checked (instances of mathematicians or Claudia's going to school) the state of affairs denoted by the generic proposition has to hold. If most mathematicians are exciting, we can truthfully say that mathematicians are generally exciting; if only a small subset of them is exciting, the assertion is not true. On the other hand, this 'neo-quantificational approach' (Kleiber 1985) avoids the shortcomings of universal quantification. The ability of adverbs of quantification to represent the nature of genericness is further demonstrated by their semantic incompatibility with object-referring NPs or particular sentences:

(19) A mathematician λ (from King's College I met yesterday) is generally exciting.

⁵One might wonder where the situation argument *s* comes from. Kratzer (1995) has suggested that episodic verbs have, in addition to their usual syntactic arguments, an argument for the location of the event described by the verb. This argument can be bound by quantificational adverbs (see also Krifka et al. 1995:31).

(20) ?Claudia generally went to school on foot.

Despite its initial plausibility, though, the neo-quantificational analysis of generics needs some further refinement before it can yield a satisfactory solution to the semantics of genericness. An immediate problem is to pin down the set of restricted occurrences over which quantification will apply. Consider the following examples — where (21) and (23) are modifications of (3) and (10) respectively:

(21) The earth turns around itself.

(22) In Haiti, the sun rises at 6 o'clock.

(23) Claudia walks to school.

The logical forms of these sentences are as follows:

(21') GEN (The earth moves in *s*; the earth turns around itself in *s*)

(22') GEN (In Haiti the sun rises in *s*; the sun rises at 6 o'clock in *s*)

(23') GEN (Claudia goes to school in *s*; Claudia walks in *s*)

The temporal intervals involved in the specification of the different situations in (21) to (23) may vary widely, and in any case cannot be part of the initial semantic representation of a characterising sentence. Instead, they are the product of pragmatic inferential work, i.e. they result from processing each utterance against a set of mutually manifest contextual assumptions.⁶ As a result, GEN in (21) comes out as quantifying over all situations in which the earth moves in any way; in (22) the domain of quantification is formed by situations in which the sun rises, i.e. only a subset of the situations in which the sun moves, separated by a one-day temporal interval; in (23) the generic operator ranges over the situations in which Claudia goes to school, which are again a subset of the

⁶The role of pragmatics in deriving non-overt restrictors has been recognised by several people in the literature on generics (see, for instance, Newton 1979, Kleiber 1985, Schubert & Pelletier 1989).

situations in which she is active, separated by one-day temporal intervals (without taking into account weekends, holidays, or situations when she is not a pupil).

More crucially, the neo-quantificational approach has to provide a more specific analysis for the semantics of GEN. In the next section I want to examine some precise proposals about the semantics of the generic operator, and look again at the interaction of semantic and pragmatic factors in the interpretation of generics.

2.2 The semantics of the generic operator

2.2.0 Of the various suggestions in the literature for specifying the semantics of the generic operator, I am going to discuss three which have tried to account for generics using the notions of prototype, stereotype and relevant quantification respectively.

2.2.1 Prototypes. The basic idea behind most prototype models is that, among the various members of a given category, one can choose some which are in a way more central or more 'representative' of the category as a whole (Rosch 1978). For instance, among various members of the category *bird*, one can distinguish some prototypical instances (such as *sparrow* or *robin*) which are closer to the prototype than others (say, *penguin*). A number of people have suggested that generics involve — at least in some cases — quantification over prototypes (Nunberg & Pan 1975, Heyer 1985, 1990). Thus, the utterance in (24) would convey that all prototypical instances of birds lay eggs:

(24) Birds lay eggs.

On this view the generic operator may be broken down into the normal universal quantifier plus a specification on the restrictor to the effect that it picks out only typical members of each concept (this can take the form of a typicality operator, as in Heyer's analyses).

It is obvious that this approach simply replaces the problem of identifying the content of the GEN operator with the problem of pinning down prototypes for each category. The latter difficulty is well-known and has been the subject of vast research in cognitive psychology (for an overview see Barsalou 1992); the current consensus seems to be that there is no unique notion capable of subsuming the various, often ad hoc criteria for forming a prototype. What for psychologists can be a fine proof of human cognitive flexibility, however, is inadequate for semanticists interested in determining the truth-

conditions of generic sentences. For them, to leave open the option of incorporating such a vague concept into the semantic representation would often result in semantic incongruity. Compare:

(25) A duck has colourful feathers.

(26) A duck lays whitish eggs.

As Krifka et al. (1995:47) point out, these examples are problematic since only male ducks have colourful feathers and only female ones lay whitish eggs. As the sets of male and female ducks are disjoint, the concept of typical duck is impossible to determine in a way suitable for both (25) and (26).

There are further difficulties. The prototype approach cannot by definition account for categories which lack a prototypical structure and yet appear unproblematically in generics; mathematical and legal terms (which can be defined in terms of necessary and sufficient conditions) are immediate candidates:

(27) Two and two equals four.

(28) Subletting extends to a maximum period of six months.

It is true that there have been experimental findings to the effect that even categories like *odd number* exhibit prototype effects (a fact that has received conflicting interpretations as evidence both for and against the prototype-based view of concepts — see Armstrong, Gleitman & Gleitman 1983 and Lakoff 1987a:148ff.). Defending the prototype analysis of generics, one could argue that (27) invokes a prototype for numbers which does not imply gradedness of membership in the category *number* but merely 'goodness-of-example'. Hence, the category *number* has both classical and prototype structure: the first is used in formal/mathematical definitions, the second in informal, everyday reasoning about numbers. No matter how we interpret Armstrong, Gleitman & Gleitman's findings, though, a prototype-based view of generics has to recognise that not all generics involve prototypes. In (27), even if there is a prototype for numbers, it does not participate in the interpretation of the utterance, which rather involves the definitional properties of the number concepts in the proposition expressed. A similar example is the following:

(29) A bachelor is unmarried.

Bachelor is a concept with a widely acknowledged prototypical structure (see Lakoff 1987a). It is clear, though, that whichever way we specify the prototype for *bachelor*, it cannot be of much help in determining the truth-conditions of (29): the proposition conveyed says something true of all unmarried male adults, not of all prototypical exemplars of them.

A directly related issue is how one is to decide for a given generic proposition whether it is to be interpreted with reference to a prototype or not. Heyer (1985) suggests that the kind of predicate will make clear whether it is assumed to be true of a kind as a whole or of the representative examples of a kind. He seems to believe that this involves some sort of *semantic* information, which will become available to the generic operator; the latter is thus expected to 'contribute to an understanding of the semantics of *prototype* propositions, if prototype propositions indeed contain an explicit reference to the *typical representatives* of a kind' (Heyer 1985:58, his emphasis). However, the appeal to a semantic typicality operator has to face a familiar set of objections to prototype semantics (Lakoff 1987b). It is preferable to think of a prototype interpretation of a generic proposition as being constructed on the basis of encyclopedic entries attached to concepts found in the predicate and the subject; this must involve some pragmatic mechanism constraining the accessibility and evaluation of encyclopedic information.

One might try to save a prototype approach in several ways. An obvious step would be to admit that the construction of prototypes depends on general cognitive and pragmatic considerations, and therefore falls outside the domain of semantics proper. It should be acknowledged that prototype formation takes into account multiple aspects of our encyclopedic knowledge, for which some sort of frame-based organisation is assumed. Within frames, prototypes represent the default values (i.e. the most frequently encountered values) of properties attributed to the members of a category. There is no reason to suppose — and in fact there is good reason to doubt — that each category has a single prototype (in other words, that the typicality operator picks out a unique member of a category every time). Frames allow for the representation of multiple prototypes (Barsalou & Billman 1989), especially when it comes to frequently encountered varieties of a certain kind. Applied to examples (25) and (26), this line of reasoning would predict that hearers use two prototypes for *duck*, which would assign to the two propositions their correct truth values.⁷ As for examples (23), (27) and (28), the present account would need

⁷Alternatively one could assume that both pairs in (25)-(26) and (27)-(28) are interpreted with respect to a prototype which is unspecified for sex. This, however, seems to lack plausibility, especially as far as the second pair is concerned and, in any case, will not work where there is considerable difference

some supplementary pragmatic machinery to explain why prototypes do not enter at all into the comprehension of these utterances.

Still, a prototype approach to generics fails to make the right generalisation, as is shown by two further types of example. First, certain generics involve ideal, rather than prototypical, exemplars of a category:

(30) Postgraduate students work hard.

On the present formulation of the prototype approach, such uses are unaccounted for. Second, characterising sentences pose a problem:

(31) Mary cleans her room without Hoovering under the bed.

It is hard to imagine what a prototype for a situation in which Mary cleans her room would look like. In any case, it does not seem to be a construct which is more or less ready-made and stored as such in the encyclopedic memory (as, for instance, prototypes for natural kinds may be). Given this sort of difficulty, it is probably no coincidence that prototype approaches to generics were based on kind-referring NPs rather than characterising predicates.

One could go on to look for some way of dealing with the objections I have raised. However, the model as it stands cannot give a satisfactory account of the semantics of genericness.

2.2.2 Stereotypes. The theory of stereotypes initially formed part of Rosch's theory of prototypes; in philosophy, it is associated especially with the work of Putnam (see Putnam 1975). A stereotype is a list of properties typically assigned to the things a given predicate applies to. It differs from a prototype in that a) it is a list of characteristics rather than a typical member of a category; b) speakers may know the stereotype for some predicate but not actually be acquainted with any prototypes of it (e.g. *ghost*, *witch-doctor* etc.). According to Putnam's (admittedly idiosyncratic) view, knowledge of stereotypes forms part of linguistic semantics and includes 'core facts' about the extension of a term, with which all speakers of the language should be acquainted. For instance, the stereotype of a tiger includes the predicate *striped*; this does not mean that all/most/normal tigers are

between the concepts for the male and female (e.g. there is no 'sex-neutral' prototype for *lion*).

striped (indeed stereotypes can be completely wrong), but that it is a widely acknowledged idea within a speech community that tigers are striped.

Let us apply this view to generics. Consider:

(32) Peacocks have richly ornamented tails with blue and green eyes.

(33) Peacocks are male.

Most speakers would accept (32) as true but not (33). Note that (32) is less likely to be true for any arbitrarily chosen peacock than (33), since only male peacocks have richly ornamented tails and only a subset of these has tails with blue and green eyes. Adopting proposals by Geurts (1985) and Declerck (1986), one could argue that this difference is caused by the fact that (32) but not (33) expresses the stereotype for peacocks in our culture. If we take generics to involve universal quantification over the set of exemplars which satisfy the stereotype for a concept, then (32) comes out as true: stereotypical peacocks in effect have colourful tails. Since, furthermore, the stereotype includes no reference to maleness, (33) is correctly predicted to be false.

As a variant of the prototype approach, the stereotype-based view of generics encounters the same problems which made the prototype solution untenable. For instance, it cannot cope with terms which apparently lack stereotypes, such as characterising predicates of the sort we saw in (31). Even if stereotypes for such predicates could be determined, it is implausible that these form part of their linguistic meaning (as was argued for natural-kind terms like *tiger* or *peacock*). A related problem arises for predicates which do have stereotypes but where these seem inadequate to account for the truth-conditions of the associated sentences. Consider *bachelor* again. According to Lakoff (1987a:85), a stereotypical bachelor is taken to be macho, pursue sexual conquest, date a lot of different women, hang out in singles bars etc. One would then expect (34) to be true and (35) false:

(34) Bachelors live alone and want to date a lot.

(35) Bachelors live with their mothers and are afraid of dating.

Both examples, though, seem equally likely to be judged true. It turns out that — *pace* Lakoff — (35) conforms to another social stereotype, that of the introvert and dependent unmarried man. The existence of multiple (and possibly conflicting) stereotypes cannot

be easily accommodated in a framework which attributes to stereotypes a *semantic* role; this would mean effectively that the predicate *bachelor* is semantically ambiguous, so that (34) and (35) can both be true at the same time.⁸ In addition, the postulation of widely diverging stereotypes, although empirically motivated, detracts from the initial plausibility and explanatory power of the notion of stereotype.

Furthermore, as Putnam recognises, stereotypes can actually be wrong. This would directly affect truth-conditions of sentences in an undesirable way: (36) would be assigned the value 'true' according to the conventional belief that foxes are sly, although the proposition conveys something strictly speaking false about actual, real-world foxes:

(36) Foxes are sly.

The stereotype-based analysis is unable to capture the fact that generics aim to make a claim of general validity about the actual world and not about culturally recognised norms. For all the above reasons, it seems that it does not provide a suitable basis for an analysis of the semantics of genericness.

2.2.3 Relevant quantification. The accounts discussed so far have assumed that genericness is essentially a semantic phenomenon, although the notions of prototype and stereotype have led us towards the semantics/pragmatics interface. By contrast, Declerck (1991) has put forth a purely pragmatic account of genericness. On his view, the generic operator might be spelled out as universal quantification over relevant entities. To determine the truth-conditions of (37), for example, we need to pin down the set of all relevant women for whom the generalisation is claimed to hold:

(37) Women are entitled to maternity leave.

The relevant set of entities is pragmatically specified on the basis of the hearer's world knowledge. Here the set of women claimed to be entitled to maternity leave is obviously the set of working women who are about to have a baby.

⁸Alternatively, we could decide to remove stereotypes from the semantic entry of predicates altogether and place them in the appropriate encyclopedic entries. I will not pursue this option further, since it would not fit the initial conception of a stereotype and would ultimately reduce to a very common phenomenon: the assignment of different contextual interpretations to what in the semantic component is a single core meaning.

A problem with this approach is that little content is given to the pretheoretical notion of 'relevant' quantification. As Krifka et al. (1995:46) are quick to point out, it is easy to find restrictions which would make *any* quantification come out true. In (38), if we take the relevant set of entities to be women who are bad conversationalists, we end up with a truism which is not what the utterance is intended to communicate:

(38) Women are bad conversationalists.

Although Declerck's analysis cannot be maintained as such, it contains one interesting idea, namely that the presence or absence of a generic operator in the logical form of an utterance is the result of pragmatic processing. This idea will be taken up again in section 3. For the moment I want to move on to a more promising approach to the semantics of genericness.

2.3 A modal account of generics

A number of people have proposed that generics can be best handled by a possible-worlds semantic framework, on a par with a variety of modal phenomena (Dahl 1975, Heim 1982). I am going to draw on such analyses to sketch how a modality-based approach to generics might go. However, my focus will be on the many pragmatic factors which come into play in the interpretation of both modals and generics, but which are normally ignored by the formal approaches. This subsection will outline and assess existing attempts to link genericness and modality; in section 3 I will modify this idea using a pragmatic theory of communication.

The fundamental tenet of the modal approach is that propositions may be descriptions — i.e. truth-conditional representations — of states of affairs in different types of worlds. In particular, they can be descriptions of states of affairs in the actual world or in a variety of alternative possible worlds. The idea behind the notion of possible worlds is that things might have been different from the way they are in reality. The real world is thus one of infinitely many possible worlds and differs from them only in that it happens to be actualised. Possible worlds are meant to exist independently of individuals who happen to entertain a given proposition, and of the way these individuals represent the proposition to themselves (for an extreme form of modal realism see Lewis 1973, 1986). Truth consequently depends on the relationship of a proposition to a world: saying that a

proposition is true or false in a given world is saying something about that world, namely that it does or does not contain the state of affairs described by the proposition.

To explain how modality functions in such a framework, we need to specify three factors which jointly underlie modal operators: following the model of Kratzer (1981), I will call them the *modal relation*, the *modal base* and the *ordering source*. The modal relation includes essentially the notions of possibility and necessity, or what the symbols \diamond and \square of modal logic are designed to capture. The modal base (or conversational background) involves a set of assumptions against which the modal relation can be understood; these sets are expressible with the phrase 'in view of -' (cf. also Kratzer 1977). Kratzer's theory includes, for instance, epistemic modal bases (where a given modality is understood in view of what is known), teleological modal bases (where modality is understood in view of one's aims), deontic modal bases (where modality applies in view of what is commanded) etc.; these are exemplified in the uses of *must* in (39a-c) respectively:

- (39) a. Judy must be John's sister: they have the same mannerisms.
 b. I must buy a bicycle to get to college quickly every day.
 c. Soldiers must do their duty.

In (39a), according to the available evidence, Judy is necessarily John's sister. In (39b), the speaker's aim to get to college quickly every day makes it necessary for her to buy a bicycle. In (39c), given certain orders, it is necessary for soldiers to perform their duty. Kratzer recognises a long list of modal bases, which further include categories such as stereotypical or buletic (related to wishes); all of these are formally treated as sets of possible worlds over which quantification by the modal operator takes place. Not all conversational backgrounds are the same for every possible world — obviously the epistemic conversational background in a world where there is no space travel is different from the one in the actual world. Kratzer points out that modal expressions are context-dependent and vague, since the sort of modal base which will be selected for their interpretation is determined pragmatically (often by some rule of accommodation of the sort proposed by Lewis 1983).

Not all worlds contained in a modal base are taken into account in the interpretation of an utterance containing a modal. For instance, in (39a) the speaker does not consider the possibility that Judy has the same mannerisms as John because she is madly in love with him and tries to imitate his every movement, or the possibility that their mannerisms are mere coincidence. Why should that be? A first approximation is that worlds containing

these states of affairs are further removed from the actual world and not easily accessible⁹ from it; they represent 'wilder' possibilities, which need not be taken into account. More precisely, they are removed from what we take to be the normal course of events, or the causal relations that hold among states of affairs. Therefore we need an *ordering source* which will rate worlds in terms of similarity to the normal course of events and specify the minimal degree of 'distance' from normalcy a world should satisfy. The closer a world is to normal conditions (usually, what holds in the actual world), the more probable it is that a modal relation will hold in that world. Modal bases come with different ordering sources: in the case of epistemic modal bases, a departure from what is known is not such a serious offence; in the case of deontic bases, however, even a slight breach of what is commanded is more problematic.

In the system I have described, a proposition is a human necessity (cf. the relation expressed by *must*) in a world *w* in view of a modal base and an ordering source iff it is true in all those accessible worlds which come closest to the normal (incidentally, since the order is not total, there can be more than one 'most normal' world). A proposition is a human possibility in a world *w* in view of a modal base and an ordering source iff its negation is not a human necessity in that world. The idea of an ordering of possible worlds along the lines of similarity belongs to David Lewis, whose theory of counterfactuals has been widely influential in this area (Lewis 1973; cf. Lewis 1986).

Let us now go back to the analysis of generics and see how the modal approach can be put to work there. Heim (1982) has suggested that the generic operator can be construed as the modal operator of necessity. Consider the utterance:

(40) Dogs have four legs.

(40') GEN (x is a dog; x has four legs)

What (40) conveys, on the modal view, is the following: everything which is a dog in the worlds of the modal base is such that, in every world which is closest to normal according to the ordering source, it will have four legs. On its preferred reading, (40) requires a realistic modal base, that is, it is interpreted according to the set of facts assumed to hold in the actual world; thus the possibility is left open for dogs in other sorts of worlds (e.g. in science fiction) to have a different number of legs. (40) can also be ascribed analytic status, in which case quantification will range over all possible worlds — including the

⁹Accessibility was first formally defined by Kripke 1971.

actual one: everything that is a dog in those worlds will have four legs under conditions closest to normal. Note that in neither case do exceptions (e.g. three-legged dogs) falsify the statement made by (40): what is being conveyed is merely that a world where dogs have four legs is more normal than a world in which dogs don't have four legs.

By allowing for variations of the modal base and the ordering source, the modal approach can explain the multiplicity of generic interpretations. Consider (41):

(41) Paul smokes.

The utterance is interpreted according to a stereotypical modal base and ordering source; it explicitly communicates that in all worlds most similar to the world in which Paul shows his typical behaviour, whenever it is appropriate, Paul smokes. In order to make a generalisation about typical behaviour, one needs some evidence from actual (and probably recurrent) behaviour; this accounts for the fact that (41) is felicitously uttered in case Paul has been observed to smoke. On the modal analysis, a situation where Paul is given an appropriate chance and does not smoke is not inconceivable, but is more remote from normalcy.

A stereotypical modal base can also account for the interpretation of certain kind-referring NPs. Consider (32) repeated as (42):

(42) Peacocks have richly ornamented tails with blue and green eyes.

Note that the modality-based approach avoids the disadvantages of the stereotype-based approach to generics, since it correctly recognises that some but not all genericness can be explained through a stereotypical conversational background and ordering source.

(43) A Christian is forgiving.

This example involves a deontic modal base (and ordering source), consisting in particular of those possible worlds in which religious and moral obligations hold: every Christian in the most normal circumstances is forgiving in those worlds. (43) does not make a statement about *actual* Christians but only about Christians in ideal (deontic) worlds; in fact, (43) does not even presuppose/imply the existence of Christians in the actual world, since the actual world does not belong to the modal base. It follows that actual nonforgiving Christians are not counterexamples to this generalisation (cf. the remarks on (40) above).

The picture I have outlined is certainly appealing as an abstract description of how the semantics of modality and genericness operate. Kratzer's original proposal is particularly successful on two points: on the one hand, in the assignment to modal expressions of a weak semantics, which together with additional contextual considerations yields the epistemic, deontic etc. interpretations; on the other hand, in the reduction of the core meanings of modal expressions to the two logical notions of necessity and possibility. The application of the framework to generics manages to bring together two areas which clearly have much in common. However, there are a number of problems, both traditional and new, which the theory has to face.

First, the possible-worlds model is not meant as a psychologically plausible model of how speakers represent and handle alternative possibilities. In fact the model as it stands rather runs counter to experimental findings, according to which human subjects have difficulty with the systematic mental manipulation of even a few alternatives (Johnson-Laird 1982). The psychological reality of possible worlds is the subject of a long and heated debate which falls outside the scope of the present paper.¹⁰ One way out would be to suggest that worlds do not correspond to full-fledged representations of states of affairs but to partial specifications of them, or *situations*.¹¹ Another would be to assume that sets of possible worlds are not individually represented but subsumed under a single description.¹² In the best of cases, the theory as outlined above should be taken to capture some facts about the speaker's semantic competence while remaining neutral as to how this competence is to be cognitively represented (see similar suggestions in Stalnaker 1986:120-1).

Now some more immediate points about Kratzer's analysis. We saw that it is not designed to fit with a cognitively informed pragmatic theory. Still, it seems that two of the three components of modality, namely the conversational background and the ordering source, involve non-linguistic knowledge and consequently belong to pragmatics (Kratzer

¹⁰For some criticisms see Smith 1983, Smith & Smith 1988.

¹¹This notion does not necessarily correspond to the use of Barwise & Perry (1983); for its application to generics see Krifka et al. (1995:57-8).

¹²Kratzer (1980), arguing against the view that possible worlds are incompatible with finite brains, shows that to know a set does not necessarily require a discrete mental representation of all its members; therefore, a given set of possible worlds need not require an individuated mental representation but may be stored under a single description (see also Partee 1989:117ff.). A similar proposal for the mental representation of time was put forth by Hans Kamp and is discussed in Johnson-Laird 1982.

herself makes vague but repeated references to the role of contextual factors in the overall comprehension of modals). The way pragmatics is dealt with in her account, however, is at least dubious.

For one thing, the idea of a fixed inventory of conversational backgrounds favouring one modal interpretation over others presupposes a rigid and inflexible conception of the role of context in comprehension. In this sense it reflects earlier views in the literature on communication, according to which context was a determined/'given' chunk of information fixed independently of the utterance (see, e.g., Brown & Yule 1983). As Sperber & Wilson (1986) have shown, however, constructing the context for understanding an utterance is part of the interpretation process, constrained by general pragmatic principles.

Moreover, the idea that possible worlds have an absolute ordering based on their similarity to an ideal needs amendment if it is to be compatible with any adequate account of utterance interpretation. Speakers certainly do not possess such an ordering: it would put too great a burden on human cognitive capacities with little foreseeable gain, especially since the notion of similarity is too vague to be used on its own (see the well-known puzzles of Goodman 1970 and Lewis 1973). So even if we knew how to give a psychologically tractable picture of possible worlds, the similarity metric would still remain a mystery.

Finally, there are more specific objections to the modal theory of generics. For instance, the present analysis does not explain why there are no generics that require an epistemic conversational background. Utterances of the sort in (44) cannot be interpreted in view of the speaker's knowledge/beliefs, that is, on a par with, say, deontic conversational backgrounds which require a set of assumptions about what is ordered/dictated by moral laws etc.:

(44) Lions are dangerous.

Last but not least, the modality-based approach is unable to resolve potential clashes in the selection of modal bases and ordering sources. Consider the examples from Krifka et al. (1995:56):

(45) A turtle is long-lived.

(46) A pheasant lays speckled eggs.

Following the modal definition of GEN, (45) is true if and only if every turtle in the modal base is long-lived in all the most normal worlds with respect to the ordering source. The sentence should be evaluated against a 'realistic' background, in which the laws of biology hold. However, the worlds in which no turtle ever dies a premature death are themselves biologically highly abnormal. This is a case where normalcy conditions contradict each other: the laws for a single organism are not compatible with the laws for a whole ecosystem. In (46) we again construct an interpretation according to biological normalcy. The proposition expressed is true if and only if every pheasant in the worlds of the modal base lays speckled eggs in every most normal world. To capture the intended universal quantification, we need to consider only worlds containing pheasants able to lay eggs (i.e. female fertilised pheasants). However, according to the laws of biology, such worlds should also include male pheasants for fertilisation to take place; and of course such pheasants would not lay eggs. These are not only paradoxes for ornithology but also for a semantic theory which does not allow enough room for pragmatics in the construction of the proposition expressed by the utterance.

So far I have argued that, although the modal approach to generics offers an interesting insight into their semantics, the way it is set up acts as a barrier to its psychologically plausible application. Obviously, the question to ask is: how can the insights of the formal approaches be integrated into a psychologically sound account of utterance interpretation? This is the issue I address in the next section.

3 Towards a relevance-theoretic account

3.1 Genericness in semantic representation: a pragmatic choice

My purpose in this section is to show how the interpretation of generics is accounted for within a relevance-theoretic framework (see Sperber & Wilson 1986²1995; for a brief summary of the theory I refer to Carston 1988). As far as the semantics of the generic operator is concerned, I assume that genericness involves universal quantification over a set of contextually determined sets of alternative states of affairs. For the sake of simplicity I will continue to refer to them as possible worlds, although it will be clear in what ways I intend them to differ from the formal notion defined in the previous section. Possible worlds are meant to be objective, external-to-the-individual states of affairs; nevertheless, their internal representation may be partial and incomplete, leaving unspecified various points which can be filled in by inference or are not of immediate

importance for a given mental process. In what follows I want to take up Declerck's observation that the decision to place GEN in the semantic representation of a given sentence is itself the product of pragmatic reasoning. What I hope to demonstrate is that the generic operator is not necessarily 'plugged in' the initial logical form of a characterising sentence, but its presence or absence is determined by considerations of relevance.

An immediate consequence of postulating a semantically encoded generic operator GEN would be that both NPs and predicates would come out as multiply ambiguous between generic and various non-generic senses. One kind of evidence against this assumption comes from a test devised by Kripke (1977) to check alleged ambiguities: would we be surprised to find two languages with two different words for, say, the generic and the non-generic meaning of a verbal predicate? Most speakers would say 'yes', which suggests that our expectations do not favour an individual encoding of genericness in referential expressions or predicates.

Further, Kripke proposes, we could check whether there are languages that separately encode the alleged senses. Genericness is a fairly pervasive and cross-linguistically robust type of meaning: if it is encoded by a linguistic operator such as GEN, we could expect languages to develop some specialised way of grammatically marking it (as they mark other semantic categories such as tense, aspect or number). Dahl (1995), after researching a large amount of cross-linguistic material on characterising sentences, concludes that natural languages follow a Minimal Marking Tendency with respect to generics: that is, generics either lack overt tense-aspect marking or they use the least marked form that the morphological system of the language offers. Furthermore, whatever form is used for generics is also used in other, non-generic environments. Even in the fairly small number of languages where there seem to be explicit grammatical markers of genericness, their semantics are not always clear.¹³ In addition, in those groups generics sometimes share their expression with other types of sentences. Dahl (1995) draws a parallel between the difficulty of finding clear grammatical distinctions that coincide with the episodic/generic distinction and the absence of a specifically 'generic' article in the languages of the world.

¹³Krifka et al. (1995:8) give the example of Swahili which includes the verbal prefix *hu-*, standardly taken to encode habituality. As the authors add, such markers are only a sufficient, and not a necessary, condition, for a generic reading: the same meaning in Swahili could be expressed using the present tense. Comrie (1985:40-1) is also cautious about such examples and rejects the idea of a habitual or a universal tense, underlining the importance of pragmatic factors in the derivation of 'generic' meanings in tense-aspect-modality systems.

He concludes that the generic/non-generic distinction 'more often than not is only indirectly reflected in speakers' choices between grammatical markers'.

It seems thus that there is no straightforward association of genericness and overt/specialised grammatical marking. In other words, genericness is not encoded by any specific type of construction. Instead, a variety of lexical and grammatical means may serve to indicate (not always decisively) that a generic interpretation is intended. In English, such mechanisms include bare plural count nouns and bare mass nouns for kind-referring NPs, and the construction *used to*, deverbal adjectives in *-able* and verbal predicates in the middle voice for characterising sentences:

- (47) a. Beavers build dams.
 b. Water is rare.
 c. John used to smoke cigars.
 d. This introduction to particle physics is highly readable.
 e. He weighs 75 kilos.

Whenever there are no compelling semantic reasons for choosing a generic over a non-generic semantic representation, a sentence will remain underdetermined in this respect. What exactly does it mean for a sentence to be underdetermined between generic and non-generic readings? Obviously it means that it is not clear in which sort of world the given proposition is expected to hold, that is, in the actual or some other kind of possible world. Resolving the underdeterminacy is part of the hearer's attempt to understand the utterance; once the hearer has decided on the sort of world in which the proposition communicated is assumed to hold, he can inferentially complete the decoded logical form into a truth-evaluable representation. The absence (or scarcity) of particular indicators of genericness in the languages of the world shows that this is a task hearers can carry out in a wide array of cases.

Usually, if there is no explicit indication to the contrary (e.g. a mood indicator, intonational cues, word order variation), the world which the proposition expressed will be taken to describe is the actual world. Indeed, it seems that this is a sort of 'default' option which is characteristic of many deictic phenomena which pertain to the 'here and now' of the act of utterance. For instance, the preferred interpretation of (48) is that it is raining here and now:

- (48) It is raining.

Let me digress slightly to see how such examples can be handled. There have been various formulations to the effect that utterances like (48) have a preferred interpretation. Fodor (1987:78) notes that the default values in (48) are almost mandatory: 'ceteris paribus, you can't say 'it's raining' and expect your hearer to default to rain in Calcutta'. Similarly, Comrie (1985:49) observes that 'one would expect a non-past tense, other things being equal, to be given the interpretation of present time reference' and, using Grice's maxim of relation, explains that 'this is the interpretation most relevant to the situation at hand'. A relevance-theoretic account of the interpretation of utterances like (48) would proceed as follows. There is an immediately accessible context in which the utterance in (48) can be processed: this consists of assumptions about the moment and location at which the utterance is produced, including the hearer's current thoughts, intentions, possible courses of action etc. Since these are likely to be influenced by the novel piece of information in (48), use of the 'here-and-now' context is also bound to yield adequate effects for the hearer. As a result, it is the context which will be selected for comprehension of the utterance in (48) and which will furnish the spatio-temporal values for enriching the logical form of the utterance into a truth-evaluable representation.¹⁴

This sort of argumentation is considered pretty much standard for propositional enrichment involving spatio-temporal correlates (see Carston 1988). I would argue that similar machinery can deal with enrichment involving modal correlates, i.e. specifying the world which a certain proposition is taken to describe. Other things being equal, the world which a proposition is taken to describe will be the actual one. Considerations of relevance indeed suggest that more immediate cognitive gain will come from a proposition taken to describe some aspect of the actual world rather than a proposition taken to describe a merely possible state of affairs. This has a number of consequences; for instance, it predicts that, other things being equal, whenever a proposition can be interpreted as non-generic, it will be. This is borne out in cases like (49), uttered during a conversation on the contemporary role of the Church:

(49) The clergy prays for the future of humanity.

¹⁴Of course these default values may change in highly accessible contexts. If (48) is uttered while watching a film, it might well be interpreted as 'It is raining in the film'. Actually it has been suggested that temporal correlates might be more basic than spatial ones, but this is an issue which has no immediate relevance here (for some discussion see Partee 1985).

Arguably, on its preferred interpretation, (49) is construed as making a claim about members of the clergy in the actual world — or at least in realistic worlds resembling the actual one in crucial respects — not in alternative worlds. The idea is that the actual world typically forms a sort of zero-point or default value for the world-location of a proposition in utterance comprehension, just as the 'here and now' provides default values for time and space location.¹⁵

If the hearer decides that the utterance was not meant as a description of a state of affairs in the actual world but as a generalisation over a range of possible worlds, he constructs a logical form with an underspecified generic operator, identical to the modal operator of necessity. Recasting the modal proposal, I take the underspecified semantics of genericness to be fleshed out by two sorts of pragmatic considerations: on the one hand, a specification of the *kind* of states of affairs that affect the interpretation of the generic proposition; on the other hand, a specification of the *range* of such states of affairs, i.e. how close these need to be with respect to some normalcy standard. I want to suggest that both sorts of considerations follow straightforwardly from the search for relevance and are therefore naturally captured by a general theory of utterance interpretation.

Consider the multiplicity of modal bases introduced in the previous section. It is clear that what they represent is the various sets of assumptions that can be brought to bear on the interpretation of generics. These assumptions involve the sorts of world that a generic may be taken to describe (e.g. realistic, stereotypical, deontic etc.). There is no need to assume that there is a finite list of such worlds, or that each one enjoys a full-fledged specification in mental representation. Instead, it seems that, in setting up representations of alternative realities, we restrict our attention to worlds which are quite similar to the world as it really is. I will look briefly at the role of similarity among worlds in the next section, before going on with the interpretation of generics.

3.2 Relevance and similarity among worlds

The problem of similarity among worlds has seriously beset previous modal approaches. We saw that not all worlds of a specific kind need be taken into account when it comes

¹⁵These default values are bound to receive some feedback during human cognitive development: precisely because they yield greater effects, they will increase in accessibility over time, according to a relevance-theoretic account of cognition and memory organisation (Deirdre Wilson, p.c.).

to the evaluation of generics, but only a limited range of them, determined by an ordering source.

I have already expressed doubts as to the usefulness and feasibility of an abstract ordering of possible worlds in terms of similarity. The point here is that, whatever the ontological status of possible worlds, their similarity ratings are a cognitive and thus subjective matter. Even as ardent an adherent of modal realism as Lewis rejects the possibility of objectively measuring similarity among worlds (e.g. by mathematical methods). He admits: 'Overall similarity consists of innumerable similarities and differences in innumerable respects of comparison, balanced against each other, according to the relative importances we attach to these respects of comparison. Insofar as these relative importances differ from one person to another, or differ from one occasion to another, or are indeterminate even for a single person on a single occasion, so far is comparative similarity indeterminate' (Lewis 1973:91). In a more revealing excerpt Lewis points out that our perception of similarity makes little use of the vast inventory permitted by logic: interlocutors expect each other to remain within a relatively limited range of inter-world similarity, and it is natural to have vocabulary conventionally reserved for use within this range (1973:94). In what follows I will take seriously the idea that humans obey cognitive and communicative constraints in their assessments of similarity (and, in particular, similarity among worlds); indeed, I will show that the 'limited vagueness' of similarity which Lewis noted can be attributed to relevance-based limitations in construing plausible alternatives to the actual world.¹⁶

The question I will first try to answer is this: why is it that, while interpreting a generic proposition, hearers take it to represent a generalisation over a range of worlds which resemble the actual one as much as is allowed by the specific type of possibility involved? A reason seems to emerge from the discussion above about the increased relevance of propositions describing the actual world. Indeed, why should a rational being entertain alternative possibilities if they had absolutely nothing to do with the way the world actually is? The issue relates directly to the way representations of alternative realities achieve relevance: by introducing a set of separate assumptions (deontic, stereotype-based, or other) into a body of assumptions about the real world which is held maximally constant, humans are typically able to derive cognitive gains which they could not have produced by manipulating solely representations of actuality. So in most cases, in order

¹⁶Interestingly Kratzer (1989) in her analysis of counterfactuals points out — without explanation — that in determining the set of propositions relevant for their truth only 'humanly graspable' propositions should be considered; cognitive factors are thus indirectly introduced in the formal analysis.

to achieve relevance, representations of a given world should be descriptions of a world not simply possible but also potential, that is, similar enough to the actual world to be able to be, become or have been actual itself.

How can any world be similar to the actual one? There are obviously quantitative and qualitative aspects to similarity. On the one hand, the two worlds have to share a number of propositions (that is, a number of propositions has to be true in both worlds); on the other, the sort of propositions they share has to be important, or non-trivial. Surely thoughts about alternative states of affairs may depart quite radically from our mental representations of reality. However, every departure from the current construal of the world demands considerable cognitive effort and therefore has to be offset by adequate cognitive effects: what these effects consist of is a different picture of what the *actual* world could be (become/have been), if a set of conditions (e.g. moral rules, obligations etc.) had been met. Moreover, apart from the cognitive side, there is a communicative side to modality. When inviting the hearer to entertain a representation of an alternative state of affairs, and see what follows from it, the speaker should be careful to convey a representation which the hearer can reconstruct from the evidence he possesses, that is, from his own representation of actuality. Since the actual world is (by definition) the only world which is realised, it is going to be the one which furnishes a vast range of assumptions mutually manifest to the interlocutors; it would be communicatively infelicitous (Grice would say 'un-cooperative') to introduce possible worlds which depart from these mutually manifest assumptions for no good reason and in unpredictable ways. If the hearer cannot see how to construct a representation of a possible world while preserving some fundamental assumptions about the actual world, he will probably give up trying and refuse to process the utterance altogether (as is the case with some extreme counterfactuals).

The problem of deciding how much of the structure of the actual world to preserve when constructing alternative possibilities has been much discussed in philosophy and linguistics. As Stalnaker puts it, 'some opinions acquire a healthy immunity to contrary evidence and become the core of our conceptual system, while others remain near the surface, vulnerable to slight shifts in the phenomena' (Stalnaker 1970:126). The former correspond to assumptions we are unwilling to part with when constructing a possible world — for instance, laws of nature, or particularly compelling empirical facts; the latter correspond to contingent details, which, if they were otherwise, would not seriously influence the picture we have about the world as it is. I will not pursue the issue any further but merely take for granted that some of our ideas about the world are cognitively

and metaphysically more central, and therefore unlikely (because of processing cost) to be revised or abandoned when we set up plausible alternatives to actuality.¹⁷

A final point on the construction of alternative possibilities. Very often the introduction of a new proposition into a world will bring chain modifications in other parts of that world. Take a world in which I have two sisters instead of one. This world will not differ from the actual one only in terms of a single state of affairs described by the proposition *p* 'I have two sisters': it will also differ in terms of a range of propositions which would also be true once *p* was true (such as 'My family has five members', or 'There are four females in my family'). This is because whatever aspect of the world makes *p* true also makes the rest of these propositions true. This relation between propositions (called 'lumping' by Kratzer 1989) is obviously very important in specifying the exact content of a possible world. However, when it comes to the psychological representation of possibilities, not all propositions lumped by a novel proposition will be computed and added to the representation of a possible world. Considerations of processing effort will ensure that only aspects which may yield cognitive gains will be explicitly represented and processed. These cognitive gains do not include, for instance, the trivial implications of *p* I mentioned; rather they include implications which are capable of entering into further inferential processes and yield contextual effects (e.g. 'I get less money from my parents', 'I share my room with two people', and so on). To the extent that dropping, adding or modifying assumptions about the actual world does not have any immediate relevance, we can safely suppose that it is left undone. As I have already hinted, when constructing a representation of a possible world, we manipulate that part of it which is cognitively rich and assume the (unspecified) rest to resemble the actual world, even though we recognise upon reflection that this is not strictly speaking correct.

3.3 The interpretation of generics

I now come to a more specific explanation of how generics are interpreted, illustrated by a variety of examples. Consider (50):

(50) Two dots are joined by a single line.

¹⁷For more detailed discussion of the philosophical issues involved here, see Stalnaker 1968, 1970, 1980, Goodman 1955, Lewis 1973.

In order for the hearer to flesh out the underspecified semantics of generics, he must find a set of worlds such that, in all of these worlds which are similar enough to reality, for every pair of dots, there exists a single line which joins this pair. The set of worlds we are after includes the worlds in which mathematical laws hold, and arguably it is the set of all possible worlds (including the actual one). Assumptions concerning mathematical knowledge are easily activated by the mathematical concepts included in the logical form of (50); therefore, it is natural to interpret (50) with respect to 'mathematically possible' worlds. The proposition is true in all such worlds, that is it expresses a mathematical necessity. Such propositions do not allow any exceptions: if there were a single instance of a pair of lines that was joined by, say, two lines, (50) would not express a true generalisation.

Example (51) is similar:

(51) Lions are animals.

The interpretation of (51) relies on assumptions about the nature (i.e. the biological make-up) of lions, which can thus be taken to hold in every world in which the laws of nature hold. Such worlds by definition include the actual world and all worlds which share with it some basic structure. In every such world, whatever is a lion is also an animal, since whatever fact makes something a lion guarantees that it will also be an animal (by virtue of a 'lumping' relation). If we try to imagine a world in which a lion is not an animal, one of two things will happen: i) we will come up with an inconsistent, hence impossible, world, in case we try to remain close to the actual world; ii) we will radically depart from our assumptions about the actual world, so that the novel world will contain phenomena of very little relevance to us. We conclude that in no world which it is rational for a relevance-oriented organism to consider are there any lions which are not animals.

The above utterance can also be interpreted as expressing a linguistic necessity, that is, a proposition true in all worlds similar enough to the actual one where language is interpreted in the same way. Here linguistic definitions (e.g. of the term 'lion') would play the same role as laws of mathematics or nature did in previous examples.

Let me move on to another sort of utterance:

(52) Peacocks have richly ornamented tails.

On the modal analysis, (52) is analysed as involving a stereotypical modal base. Indeed, the hearer is not entitled to assume that the utterance expresses a generalisation holding

for all peacocks in all worlds, but rather for peacocks in those possible worlds satisfying a specific natural-kind stereotype; the reason is that the encyclopedic entry of the concept *peacock* contains a particularly accessible piece of information, namely the fact that male members of the species have a richly ornamented tail. Given the activation of the property *ornamented tail* towards the end of the utterance, the hearer can rationally conclude that the context in which the utterance should be processed includes only assumptions about male peacocks. In interpreting (52), he assumes that the generic operator ranges over worlds which contain only the relevant subset of these birds (i.e. the male ones) and indeed, on this interpretation, (52) is true. Of course, another hearer lacking the corresponding item of encyclopedic knowledge might take (52) to convey a generalisation over all members of the species in all worlds which reasonably resemble the actual one — and consequently conclude that (52) is false. A pragmatic account that allows for different contextual considerations to be brought to bear in determining the proposition expressed by an utterance can naturally explain such differences.¹⁸

A number of cases which presented difficulties for the modal approach as outlined in section 2.3 can now be more readily accounted for. Recall:

(53) A turtle is long-lived.

(54) A pheasant lays speckled eggs.

(53) was problematic because natural laws for turtles clash with natural laws for the ecosystem to which turtles belong, and consequently it is difficult to determine the sort of possible world which (53) describes. It turns out, though, that what we are concerned with are in effect two different interpretations of the utterance. On one interpretation, considerations of relevance narrow down the choice of possible worlds to ones in which laws of biology hold *with respect to turtles*: although these laws might be contradicted by more general aspects of biological evolution, hearers do not represent those aspects to themselves, since they are not relevant for the interpretation of (53). What the utterance conveys is that ideal turtles, or turtles in isolation, have the potential to live long — which

¹⁸Another relevance-theoretic proposal about the interpretation of (52) would involve enrichment of the referential expression *peacocks* into *male peacocks*; enrichment can also be invoked to explain analytic and deontic examples. This proposal has been made with respect to the generic interpretation of indefinite description in Rouchota (1994). However, such an approach could not be extended to characterising sentences and would fail to capture the parallel between generics and explicit modal expressions.

is a true proposition. On the second interpretation, the hearer takes into consideration worlds in which laws of biology hold for the whole ecosystem, and thus constructs a different array of possible worlds for interpreting (53). In the most natural of these worlds (including the actual one), the proposition conveyed by (53) is false. Since the second interpretation requires a sort of world that comes closer to (and indeed includes) the actual one, my account predicts that it will be preferred; this prediction seems to be borne out. The fact that (53) can be construed in these two very different ways is verified by the observation that it can be followed by different kinds of response:

- (55) a. Yes, a turtle lives even longer than a frog.
 b. No, a turtle rarely survives the dangers in its environment.

The generic in (54) is similar to (53). The problem here was that, although quantification is supposed to range over worlds which include female fertilised pheasants only, normalcy considerations would suggest that such worlds have to include male birds as well. The obstacle is lifted if we realise that only worlds in which the specific subset of pheasants exist matter for the interpretation of the utterance; these worlds include the actual one and a range of worlds that resemble it *as far as egg-laying of pheasants is concerned*. A world inhabited solely by female fertilised pheasants may be highly abnormal if the whole body of laws of biology is brought to bear on it; this is not what hearers of (54) do, since maintaining this aspect of similarity to the actual world will bring no cognitive gains — in other words, it is not relevant to the interpretation of the utterance.

So-called 'deontic' generics like (56) can be explained along the same lines:

- (56) A gentleman escorts a lady home.

The proposition in (56) holds for all gentlemen inhabiting a range of ideal worlds, where rules of good behaviour are obeyed. In all other respects these worlds look quite like the actual world; worlds in which ladies are always accompanied by their parents when they go out with gentlemen do not matter for evaluating the relevance of the proposition.

What about the characterising sentence in (57)?

- (57) This machine squeezes oranges.

The comprehension of (57) requires a type of world in which a certain machine performs the function for which it was made. Laca (1990) — cited in Krifka et al. (1995:54) —

points out that this is normal for characterising sentences about artifacts, but fails to give an explanation for this phenomenon. I would suggest that the reason lies in the existence of functional (rather than, say, perceptual) stereotypes for concepts of artifacts, which are widely used in their categorisation and processing. Such stereotypes are formed through our everyday interaction with artifacts, which involves their manipulation and use for a number of purposes. Thus, an immediate context for the interpretation of (57) will include assumptions about worlds where the machine conforms to the stereotype. Since such worlds may not include the actual one, it is possible for (57) to be acceptable even if no actual occurrence of the machine squeezing oranges has been observed.

Let me just offer some final remarks on the modal approach to generics. A question that might arise is whether, apart from the unrealised modal operator of necessity, there exists in generics an operator of possibility; the consequence would be that the interpretation of some generics would involve a subset of possible worlds of a certain sort. An empirical investigation shows that generalisations of this type do not appear. For independent reasons, Heim (1982:184) excludes the possibility of a phonologically null possibility operator on syntactic grounds. This solution seems satisfactory since it also serves to bring together generics and other modal environments such as bare conditionals.

A second interesting question is one I have mentioned already: why haven't we taken into account in interpreting generics those possible worlds corresponding to the speaker's knowledge (i.e. epistemic possible worlds)? In our modal analysis we have made use of 'realistic' (actual-like), stereotypical and deontic possible worlds, exactly on a par with the analysis of overt modal expressions in natural language. However, we have not made use of epistemic possible worlds which are needed for the analysis of epistemic necessity expressed by an overt modal operator such as *must*. Compare the epistemically modalised (58) to (59):

(58) This must be the postman.

(59) Dancers at the Royal Ballet are well-paid.

On our analysis, (59) conveys that in all (realistic) worlds which remain close to the actual one, whenever someone is a dancer at the Royal Ballet, they are well-paid. Why don't we say instead, following the analysis of (58), that the state of affairs described in (59) holds in all worlds similar enough to the actual one where the speaker's knowledge holds? The question seems to be a matter of principle rather than empirical fact, since intuitively the new formulation hasn't brought forth anything tremendously different. If we adopt an

'externalist' view, on which knowledge is defined as belief true in the actual world, then epistemic possible worlds are a subset of realistic possible worlds; consequently, generics construed with respect to realistic conditions are also construed with respect to the speaker's knowledge. If again we adopt an 'internalist' view, on which knowledge is defined as belief the individual holds as true, we return to the trivial assumption that the speaker's communicative contribution conforms to her stock of beliefs: what the speaker puts forth as a generalisation is obviously compatible with what she herself takes to be true. So in a sense epistemic modality is trivially present whenever we analyse a generic as holding in a realistic or stereotypical kind of world.¹⁹

4 Conclusion

In the vast literature on generics, it has often been suggested that the phenomena studied under genericness do not actually constitute a natural class (see Lawler 1972, Kleiber 1985). In this paper I have taken the opposite view. My main concern has been to expand on the widely acknowledged but rarely developed idea that generics are context-dependent constructions. In particular, I have tried to show that their context-dependence results from an underspecified modal semantics, which is pragmatically developed into a truth-evaluable representation; in the course of this development, considerations of relevance determine the set of possible worlds over which the generalisation communicated by a generic is supposed to hold. In my discussion I have unavoidably glossed over a number of important differences between kind-referring and characterising sentences, and finer distinctions holding within each of these categories; the result is a clearer, albeit still sketchy, picture. Although a large number of issues remain problematic in the area of generics, I hope to have demonstrated that a modal-based approach is a promising starting point for solving some of the puzzles.

¹⁹Deontic generics are a different lot, since they are concerned with ideals, moral imperatives, commands etc.

References

- Armstrong S., L. Gleitman & H. Gleitman (1983), What some concepts might not be. *Cognition* 13, 263-308.
- Barsalou L. (1992), *Cognitive Psychology*. Hillsdale, N.J.: Erlbaum.
- Barsalou L. & D. Billman (1989), Systematicity and semantic ambiguity. In D. Gorfain (ed.) *Resolving semantic ambiguity*, 146-203. New York: Springer-Verlag.
- Barwise J. & J. Perry (1983), *Situations and attitudes*. Cambridge, MA: MIT Press.
- Brown G. & G. Yule (1983), *Discourse analysis*. CUP.
- Burton-Roberts N. (1977), Generic Sentences and Analyticity. *Studies in Language* 1, 155-96.
- Carlson G. (1977a), A unified analysis of the English bare plural. *Linguistics and Philosophy* 1, 413-56.
- Carlson G. (1977b), *Reference to kinds in English*. Ph.D. dissertation, University of Massachusetts, Amherst. Publ. 1980 by Garland Press, New York.
- Carlson G. (1982), Generic terms and generic sentences. *Journal of Philosophical Logic* 11, 145-81.
- Carlson G. & F. Pelletier, eds. (1995), *The generic book*. Chicago: The University of Chicago Press.
- Carston R. (1988), Implicature, explicature, and truth-theoretic semantics. In R. Kempson (ed.), *Mental representations: The Interface between Language and Reality*, 155-81. Cambridge: CUP.
- Clark R. (1973), Prima facie generalisations. In G. Pearce & P. Maynard (eds.), *Conceptual change*, 42-54. Dordrecht: Reidel.
- Comrie B. (1985), *Tense*. CUP.
- Dahl Ö. (1975), On Generics. In E. Keenan (ed.), *Formal Semantics of Natural Language*, 99-112. London & New York: CUP.
- Dahl Ö. (1995), The marking of the episodic/generic distinction in tense-aspect systems. In G. Carlson & F. Pelletier (eds.), 412-25.
- Declerck R. (1986), The manifold interpretations of generic sentences. *Lingua* 68, 149-88.
- Declerck R. (1991), The origins of genericity. *Linguistics* 29, 79-102.
- Farkas D. (1981), *Intensionality and Romance Subjunctive Relatives*. Univ. of Chicago Dissertation. Bloomington: IULC.
- Farkas D. & Y. Sugioka (1983), Restrictive if/when clauses. *Linguistics and Philosophy* 6, 225-258.
- Fodor J. (1987), A situated grandmother? Some remarks on proposals by Barwise and Perry. *Mind and language* 2, 64-81.
- Geurts B. (1985), Generics. *Journal of Semantics* 4, 247-55.
- Goodman N. (1955), *Fact, fiction, and forecast*. Cambridge, Mass.: Harvard University Press.
- Goodman N. (1970), Seven strictures on similarity. In L. Foster & J. W. Swanson (eds.), *Experience and theory*. University of Massachusetts Press.
- Harper W., R. Stalnaker & G. Pearce, eds. (1981), *IFS: Conditionals, belief, decision, chance, and time*. Dordrecht: Reidel.
- Heim I. (1982), *The semantics of definite and indefinite noun phrases*. Ph.D. dissertation, University of Massachusetts, Amherst.
- Heyer G. (1985), Generic descriptions, default reasoning, and typicality. *Theoretical Linguistics* 11, 33-72.
- Heyer G. (1990), Semantics and knowledge representation in the analysis of generic descriptions. *Journal of Semantics* 7, 93-110.

- Johnson-Laird P.N. (1982), Formal semantics and the psychology of meaning. In S. Peters & E. Saarinen (eds.), *Processes, beliefs, and questions*, 1-68. Dordrecht: Reidel.
- Kleiber G. (1985), Du côté de la généricité verbale: Les approches quantificationnelles. *Langages* 79, 61-88.
- Kratzer A. (1977), What 'must' and 'can' must and can mean. *Linguistics and Philosophy* 1, 337-55.
- Kratzer A. (1980), Possible-worlds semantics and psychological reality. *Linguistische Berichte* 66, 1-14.
- Kratzer A. (1981), The notional category of modality. In H.-J. Eikmeyer & H. Rieser (eds.), *Worlds, words, and contexts*, 38-74. Berlin: de Gruyter.
- Kratzer A. (1989), An investigation of the lumps of thought. *Linguistics and Philosophy* 12, 607-53.
- Kratzer A. (1995), Stage-level and individual-level predicates. In G. Carlson & F. Pelletier (eds.), 125-175.
- Krifka M., F. Pelletier, G. Carlson, A. ter Meulen, G. Chierchia, G. Link (1995), Genericity: An Introduction. In G. Carlson & F. Pelletier (eds.), 1-124.
- Kripke S. (1971), Semantic considerations on modal logic. In L. Linsky (ed.), *Reference and Modality*, 63-72. Oxford: Clarendon Press.
- Kripke S. (1977), Speaker's reference and semantic reference. In P. A. French, T. E. Uehling & H. K. Wettstein (eds.), *Contemporary perspectives in the philosophy of language*, 6-27. Minneapolis: University of Minnesota Press.
- Laca B. (1990), Generic objects: some more pieces of the puzzle. *Lingua* 81, 25-46.
- Lakoff G. (1987a), *Women, fire, and dangerous things*. Chicago: University of Chicago Press.
- Lakoff G. (1987b), Cognitive models and prototype theory. In U. Neisser (ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorisation*, 63-100. New York: Cambridge University Press.
- Lakoff R. (1972), The Pragmatics of Modality. *Papers from the VIII Regional Meeting of CLS*, 229-46.
- Lawler J. (1972), Generic to a fault. *Papers from the VIII Regional Meeting of CLS*, 247-58.
- Lewis D. (1973), *Counterfactuals*. Oxford: Blackwell.
- Lewis D. (1975), Adverbs of quantification. In E. Keenan (ed.), *Formal Semantics of Natural Language*, 3-15. London & New York: CUP.
- Lewis D. (1983), Scorekeeping in a language game. In R. Baeuerle, U. Egli & A. von Stechow (eds.), *Semantics from different points of view*, 172-87. Berlin: Springer-Verlag.
- Lewis D. (1986), *On the plurality of worlds*. Oxford: Blackwell.
- Lyons J. (1977), *Semantics*. 2 vols. CUP.
- Newton B. (1979), Scenarios, modality, and verbal aspect in Modern Greek. *Language* 55, 139-67.
- Nunberg G. & C. Pan (1975), Inferring Quantification in generic sentences. *Papers from the XIth Regional Meeting of CLS*, 412-28.
- Partee B. H. (1985), Situations, worlds and contexts. *Linguistics and Philosophy* 8, 53-8.
- Partee B. H. (1989), Possible worlds in model-theoretic semantics: A linguistic perspective. In S. Allén (ed.), *Possible worlds in Humanities, Arts and Sciences*, 93-123. Berlin - New York: de Gruyter.
- Putnam H. (1975), The meaning of meaning. In K. Gunderson (ed.), *Language, mind and knowledge*, 131-93. Minneapolis: University of Minnesota Press.
- Rouchota V. (1994), On indefinite descriptions. *Journal of Linguistics* 30/2, 441-475.
- Rosch E. (1978), Principles of categorisation. In E. Rosch & B. Lloyd (eds.), *Cognition and categorisation*, 27-48. Hillsdale, N.J.: Erlbaum.
- Schubert L. & F. Pelletier (1989), Generically speaking. In G. Chierchia, B. Partee & R. Turner (eds.), *Properties, Types and meaning*, vol.2, 193-268. Dordrecht: Kluwer.

- Smith N. V. (1983), On interpreting conditionals. *Australasian Journal of Linguistics* 3, 1-24.
- Smith N. V. & A. Smith (1988), A relevance-theoretic account of conditionals. In L. Hyman & Ch. Li (eds.), *Language, speech and mind: Essays in honour of Victoria Fromkin*. London: Routledge.
- Sperber D. & D. Wilson (1986, ²1995), *Relevance: Communication and cognition*. Oxford: Basil Blackwell.
- Stalnaker R. (1968), A theory of conditionals. In N. Rescher (ed.), *Studies in Logical Theory*. Repr. in W. Harper et al., ed. (1981), 41-555.
- Stalnaker R. (1970), Probability and conditionals. Repr. in W. Harper et al., eds. (1981), 107-28.
- Stalnaker R. (1980), A defense of conditional excluded middle. Repr. in W. Harper et al., eds. (1981), 87-104.
- Stalnaker R. (1986), Possible worlds and situations. *Journal of Philosophical Logic* 15, 109-23.