

Andrea Bonomi - Paolo Casalegno

CONTEXTS, QUANTIFICATION AND GENERIC SENTENCES

I. THE CONTEXTUAL NATURE OF QUANTIFICATION

We shall argue for the following theses.

I. A satisfactory semantic analysis of natural language requires each occurrence of a quantified noun phrase to be endowed with an index for contexts. The common noun of the noun phrase must be interpreted as relativized to that context. In other words: a common noun in a noun phrase must be understood as standing for a *contextualized* property.

II. To each noun phrase there corresponds a property of contexts, and one can find cases, in natural language, in which a quantification over contexts is restricted by the property corresponding to a noun phrase. We shall illustrate this phenomenon by discussing the so-called generic reading of the indefinite article.

We shall not give here a precise definition of the term 'context'. Roughly speaking, a context is a portion (as little as you want) of a more comprehensive state of affairs. If we regard a classical model *M* as a complete picture of the world, contexts can be represented by means of 'partial submodels' of *M*, i.e. structures which embody only partial information as to which individuals of the domain of *M* have which properties (or are related by which relations). Here, anyway, we shall confine ourselves to the discussion of a few linguistic facts: on the one hand, we shall avoid too broad methodological issues; on the other, we shall omit all technical details. Our aim is to sketch some problems and to give you an inkling of why we find them interesting.

So let us begin by illustrating our first thesis.

I.1. The Traditional View

When you say that *all the dogs* are sleeping, you hardly mean that all the dogs in the world are sleeping. Rather, a context is likely to make a particular class of dogs relevant, and what you mean is that all the dogs of *this* class are sleeping.

And if you say that *the dog* is sleeping, usually this statement does not entail that there is just one dog in the world, and this dog is sleeping. Once more, a context will serve to pick out the intended dog, so that your statement ascribes something to *this* dog.

Such truisms remind us that the reference to the global *universe* of discourse is not sufficient to characterize the use of quantified noun phrases in natural language. *Contexts* must also be taken into account.

Occasionally this truth has emerged in contemporary logic and philosophy of language. But it seems to us that its relevance has not yet been duly appreciated. According to a deep-rooted way of reasoning, that we shall call here for brevity *the traditional view*, contexts must eventually be accounted for in pragmatics, not in formal semantics. The reference to a global universe is all we need, and the truisms

we mentioned at the outset can be *exorcized* in different ways. For instance, by saying that formal semantics is concerned only with idealized situations. Or that, in each case, a context can be seen as a suitably restricted *universe*, so that the very notion of context turns out to be dispensable. Or that a proper paraphrase of the quantified noun phrase can account for the intended restriction. And so on.

In order to show the inadequacy of this view, we shall start from some reflections on a little parable. The story is the following.

I.2. Some Simple Facts

John's life in his country house is gladdened by the presence of five nice dogs. The only problem, with these lively pets, is that sometimes they run off and pick a fight with the cats they meet. The neighbours have complained, and John is a little worried. So, when he hears some dog bark in the distance, he gives his wife a look of anxiety. Perhaps some of the dogs ran into Bill's cat. But she reassures him. She says:

(1) All the dogs are sleeping.

If John feels in better spirits, the only reason is that he accepts (1) as *true*. But how is that possible, since John is aware that some dog is barking and, as a consequence, is not sleeping? The natural answer is that both John and his wife give the noun phrase 'all the dogs' a restricted interpretation, by referring to a proper context. All the dogs selected by *that* context are sleeping: this is why (1) is true. What happens to dogs elsewhere is not relevant.

All this is so obvious that the traditional view cannot ignore it. But, according to that view, a universal quantification involves *all* the individuals of a given universe: so, in our example, this policy would entail the falsity of (1) - against our intuition. To get round this difficulty, two (alternative) arguments are offered. A quick look at them is sufficient to see their inadequacy.

a. Ellipsis. What is suggested is that in such a sentence as (1) the quantified noun phrase must be assumed as an ellipsis of a more extended expression: for example, in (1), 'all the dogs' must be read as 'all the dogs we have' or, more simply, 'all our dogs'.

Against this kind of justification, we do not intend to deny that, in some particular cases, a quantified noun phrase has an elliptical form. What we deny is that this approach can offer a *systematic* treatment of quantification. For we must wonder: what is, each time, the "right" paraphrase? Is there any systematic procedure that enables us to associate the given noun phrase with its extended version, where all the required specifications are made explicit? Notice that, even in our simple example, the suggested expansion can be inappropriate: for example if John and his wife are not the masters of all the dogs at issue. (Imagine that two of them live in the villa next door, but John feels responsible for their misdeeds because their master, who is spending his holidays by the sea, has entrusted these dogs to his care). Now, if each time an *ad hoc* paraphrase must be found, and no systematic procedure is available, we are no longer able to assign *recursive* truth-conditions to this class of sentences. Each time an appeal to unstable intuitions is essential to fix the logical form of the sentence (whatever you mean by this term). And this is an unwelcome outcome for our semantics.

On the other hand, the task of formal semantics is not to say what is in each case, the relevant context, but more interestingly, to account for the *systematic* relations between contexts (however they are selected) and truth-conditions. And this can be done in a recursive way.

b. The pragmatic determination of the domain. The adopted strategy is, this time, simpler. One admits that, in such a sentence as (1), the domain of quantification does not involve all the dogs in the world. And the consequent suggestion is: in your model $M = \langle D, F \rangle$ just take D of the size you need. After all, the traditional view is not committed to any particular assumption on the size of the domain. So in the truth-definition the clauses concerning our quantifiers can be left unchanged, whilst the determination, in each case, of the relevant universe is a pragmatic matter which can be ignored if the idealized perspective of (model-theoretic) semantics is adopted.

We believe this is the most common argument the advocates of the traditional view refer to (implicitly, in many cases) when they have to justify the apparent inadequacy of the usual approaches to quantification in natural languages. So it deserves a more careful consideration.

(i) The suggested argument does not fit our intuition. For, if we consider the example at issue we cannot help admitting that, although John accepts (1) as true, the portion of world (with the respective universe of things) John is acquainted with *is large enough to falsify that sentence*. (In fact, he knows that there is a dog which is not sleeping: he hears it bark!). On the other hand, John's behaviour is perfectly natural, and the only way to account for this is to say that by means of the quantified noun phrase John is referring to a *restricted* domain which *cannot* coincide with the given (total) universe (where not all the dogs are sleeping) and is contextually determined. So the moral is that the notion of a general universe and the notion of context are both needed: to conflate them is tantamount to making the idea of *restricted* quantification a nonsense.

(ii) Certainly the *absolute* invariance of the domain is an untenable dogma. But what about a notion of universe which involves, in many cases, a change of universe from sentence to sentence? Imagine that, after uttering (1), John's wife says

(2) The barking dog is a pest.

Well, if you accept the argument suggested by the advocates of the traditional view, in order to recognize (2) as true (as it is assumed to be in our example) you must presuppose, *as necessary*, a switch of universe. In fact, for (1) to be true, a universe without barking dogs is needed (all the dogs are sleeping), but a different universe is relevant in the case of (2), where a reference is made to a barking dog. Is this fickleness of universes plausible? After all, to give up the absolute invariance of the domain (what would it be like to be the universe of *all* the possible discourses?) does not entail the rejection of the (reasonable) idea that a universe is at least something *relatively* stable.

Both points above have to do with intuitive considerations. But a more compelling counterexample can be presented. In our opinion it represents a conclusive proof of the inadequacy of the argument suggested in (b) and, more in general, of the *empirical* inadequacy of the traditional view. Because of this theoretical significance, the counterexample we have in mind needs a separate discussion.

1.3. The Ubiquity of Contexts

Here is a new version of the above homely scene. This time John's wife says:

(3) All the dogs are sleeping peacefully for they do not hear the barking dog.

The reason why this utterance cannot be accounted for along the lines of the strategy we have just mentioned is quite simple. A *single* sentence is at issue here. So just one universe must be selected (on "pragmatic considerations") as the relevant one. As we have seen in the previous discussion, there are two plausible candidates. But if the more restricted universe is picked out (the universe which includes only sleeping dogs), the sentence turns out to be false (or lacking a truth-value), since no dog is barking. On the other hand, if the large universe is selected (where there is a barking dog), the sentence is, once more, false - for not all the dogs are sleeping. In short, according to the traditional view such a sentence as (3) is not satisfiable; but this clashes with our assumption that, in the given situation, (3) is true.

The traditional view is not the only victim of this kind of counterexamples. Think, for instance, of Bar-Hillel's solution. The idea is that, in order to account for the contextual nature of reference and quantification, what we have to evaluate is not a sentence, but an assertion (or a "judgement", as he says): i.e. a pair $\langle S, c \rangle$, where S is a sentence and c the intended context. Now, in the face of such utterances as (3), this policy is not effectual either. And the reason is the same as before: no *single* context can give both noun phrases the intended interpretation.

A radical solution seems to emerge from these reflections: *each* quantified noun phrase in a sentence is to be connected with a specific context.

I.4. Zooming

Our last informal reflections hinted at an essential connection between the referential power of a noun phrase and the suitable delimitation of a local setting that specifies, within the given model, which properties and relations characterize a relevant set of individuals. The usual assimilation of a model to a complete picture of the world might suggest a metaphor: referring presupposes a sort of zooming, which each time isolates the intended portion of world.

This metaphor is tempting, but it requires a qualification. In zooming we are bound by the spatial coherence of the scene. This is why the term 'portion' (or 'segment') of world is appropriate. We can widen the scene, or we can make it smaller. But what occurs in the scene depends on the *physical* point of view we have adopted. On the other hand, the notion of context we have in mind does not have this realist flavour. The way properties are delimited is subject to no particular restriction: it only depends on the pragmatic needs of the discourse. Perhaps the "spatial" characterization of the partial situations associated with referring noun phrases may be attractive. But to see that it is misleading, just think of a situation in which John's five dogs are, respectively, in Gressoney (Aosta), Sidney, Toronto, Nairobi, Shanghai. (Suppose, for instance, they are taking part in different beauty contests). After some phone calls, John's wife says:

(4) All the dogs are fine.

Now, what keeps these dogs "together" (as a partial domain of quantification) is neither the spatial homogeneity of the situation, nor any other intrinsic property of the world. The cohesive force of the context is determined only by pragmatic reasons, where intentionality plays a central role.

The following conclusion can be drawn from our last remarks. Formal semantics must exhibit structures whose articulation is "fine-grained" enough to make a distinction between the given universe and the set of countless possible contexts. This is the main difference between the theoretical framework we are going to describe and the traditional view. And it is not a negligible difference, if, as we

believe, there is a *logic of contexts* to be captured. But the way such a context is selected in each case cannot be explained in semantic terms. It is a matter of pragmatics. Perhaps we are here close to the intrinsic limits of semantic expressibility.

I.5. Qualifications

What we have been saying so far is of course only a starting point. We have not discussed (and we will not discuss here) several problems which arise as soon as one tries to transform the rough sketch we have provided into a proper formal analysis of noun phrases. For instance, we have been very vague as to where, exactly, the context-indices should be inserted. Should they be appended to noun phrases themselves? Or rather to common nouns? And how to deal with noun phrases containing a complex common noun which in turn contains other noun phrases, such as 'the boy who danced with all the girls at the party' or 'the woman all women envy'?

Another problem is this. There obviously are cases in which the interpretation of a noun phrase does not seem to be subject to any contextual restriction. If we say 'Every natural number is either odd or even', we are presumably referring to the whole set of natural numbers, not to a contextually determined subset of the natural numbers. Now, this forces us to choose between two possibilities. We can insist that *each* occurrence of a noun phrase must be thought of as endowed with a context-index, and claim that the context relevant to the interpretation of such noun phrases as 'every natural number' in the last example is the *maximal* context, coinciding with the whole universe of discourse. Alternatively, we can admit that there are cases in which the indexing must be dropped. We shall not examine here the pro and con of these two possible approaches.

Further problems arise from the fact that the contexts to which the interpretation of noun phrases is relativized often seem to be quantified over. This observation opens up a fascinating line of research, which could lead us to rewrite in terms of contexts the analysis of the linguistic phenomena concerning tense and aspect. We will not pursue this matter here. However, quantification over contexts will play a central role in the discussion of our second thesis, to which we now turn.

Our second thesis is: there are cases in which a noun phrase must be interpreted as expressing a property of contexts. Take the noun phrase 'a dog'. Suppose the value of the contextual parameter relevant to the interpretation of this noun phrase has been determined somehow. For example, suppose it is clear that we are talking about John's dogs. Now, which property of contexts can be said to correspond to the noun phrase 'a dog' understood in this way? An obvious answer is: the property of being a context in which a dog belonging to John is present. Similarly, the noun phrase 'two dogs' can be thought of as expressing the property of being a context in which two dogs belonging to John are present, the noun phrase 'all the dogs' can be thought of as expressing the property of being a context in which all the dogs belonging to John are present, and so on. All this is trivial. What is by no means trivial is the fact that this way of looking at noun phrases is exactly what we need to explain certain linguistic facts. As we said at the beginning, we shall illustrate this point by discussing the so-called generic reading of the indefinite article in Italian.

II. PROPERTIES OF CONTEXTS: GENERICS

When we learnt elementary logic, we were told that, in normal circumstances, the indefinite article corresponds to an existential quantification. We were warned, however, that there also exists a different - generic - reading of the indefinite article

corresponding, more or less, to a universal quantification. Let us consider this issue from the point of view of a general theory of contexts.

II.1. The Problem

The generic reading is exemplified by sentences such as:

- (5) Uno scout possiede un coltello a serramanico (*A boy-scout owns a spring-knife*)
- (6) Un ragno ha otto zampe (*A spider has eight legs*)
- (7) Un pirata beveva rum (*A pirate used to drink rum*)
- (8) Un' aquila costruisce il nido sulla roccia (*An eagle builds the nest on the rock*)

These sentences can be interpreted as (roughly) equivalent to the following ones:

- (5') Tutti gli scout possiedono un coltello a serramanico (*Every boy-scout etc.*)
- (6') Tutti i ragni hanno otto zampe (*Every spider etc.*)
- (7') Tutti i pirati bevevano rum (*Every pirate etc.*)
- (8') Tutte le aquile costruiscono il nido sulla roccia (*Every eagle etc.*)

Now, how can the generic reading of the indefinite NPs be accounted for?

Our starting point was an interest in the role the notion of context - or, to use a more common term, the notion of 'circumstance'¹ - can play in the semantics of natural language. Now, if we are ready to admit that sentences like (5)-(8) involve a universal quantification over circumstances, the existence of the so-called generic reading of the indefinite NPs can be accounted for without attributing to the indefinite article any ambiguity. In other words, the indefinite NP can preserve its usual meaning. Take, for instance, sentence (5) above. Our idea is that the content of (5) is more or less the following:

- (9) Every circumstance in which there is a boy-scout is a circumstance in which there is a boy-scout who owns a spring-knife.

This is, of course, a very rough formulation, but it should suffice to explain what we have in mind. The point to be emphasized is that (9) does not contain any explicit universal quantification over boy-scouts. What in (9) corresponds to the indefinite NP 'uno scout' are two occurrences of 'there is a boy-scout', i.e. two existential quantifications. The required interpretation results from the interplay between these existential quantifications and a universal quantification over circumstances. So, if something like (5) is indeed an adequate paraphrase of (5), it is not true that indefinite NPs are ambiguous between an existential and a generic reading. In a sense, the generic reading of indefinite NPs does not exist: indefinites are interpreted uniformly in all contexts; other elements of the context are responsible for the reading of sentences like (5)-(9) we are interested in. On this point, we find ourselves in agreement with the most popular theory of indefinites: the so-called Lewis-Heim-Kamp theory. The difference is the following: according to that theory, indefinite NPs must be uniformly interpreted as expressions devoid of quantificational force; as for us, we think that indefinites always correspond to an existential quantification.

¹Latin *circumstantia*, a standing around, condition < *circum-*, around + *stare*, to stand.' (Webster Dict.)

Our discussion will be divided into two parts. In the first - and longer - part we discuss the relation between generic reading and focus. In the second part we touch upon some questions concerning the actual formalisation of our proposal.

II.2. Generic Reading and Focus

Generic reading and focus. It has been observed that the generic reading of an indefinite NP seems to be allowed only when the NP occurs in subject position, not when it occurs in object position. Consider the following examples:

- (10) Truman Capote adulava una donna ricca (*T. C. flattered a rich woman*)
 (11) Una zitella protegge un gatto randagio (*A spinster protects a stray cat*).

(10) can only mean that Truman Capote used to flatter a certain rich woman, it cannot mean that Truman Capote used to flatter every rich woman (or, at least, every rich woman he came across). In (11) the first indefinite NP can have generic reading, but the second cannot: in other words, (11) can mean that for every spinster x there is a stray cat y such that x protects y , but it can mean neither that every spinster protects every stray cat, nor that every stray cat is protected by some spinster.

The fact exemplified by (10) and (11) can be explained by saying that an indefinite NP cannot receive a generic interpretation if it is contained in the VP of the sentence (this hypothesis has actually been put forward by several people: by Greg Carlson, for instance, and in a more developed form by Molly Diesing and Angelika Kratzer). Here, however, we will consider a different hypothesis (also proposed by Manfred Krifka): according to this hypothesis, the generic reading of 'una donna ricca' in (10) and 'un gatto randagio' in (11) is impossible because the NPs in question do not belong to the presupposition of the sentence containing them.

We are using here the term 'presupposition' in the sense of the focus-presupposition distinction. Usually, this distinction is drawn in terms of new vs old information. The focus of a sentence is described as the part of the sentence which conveys new information, whereas the presupposition is described as the part of the sentence which expresses information already available to both speaker and hearer. The focus-presupposition structure of a sentence is related to its stress pattern. The focus can be any constituent containing the main stress of the sentence (more precisely: any constituent whose main stress under normal intonation coincides with the main stress of the sentence). The presupposition is what we obtain when we remove the focus from the sentence and replace it by a variable. Here is an example:

- (12) A. Che cosa ha mangiato Giovanni? (*What has Giovanni eaten?*)
 B. Giovanni ha mangiato CARMELLE (*Giovanni has eaten TOFFEES*).

Suppose that the main stress of (12B) is on 'caramelle' (we express this by writing 'caramelle' in capital letters). So the possible foci of (12B) are the constituents containing 'caramelle', i.e.: 'caramelle', 'ha mangiato caramelle' and the whole sentence. In the present context the right choice is 'caramelle', for this is the part of the sentence conveying new information. The corresponding presupposition is obtained by removing 'caramelle' from the sentence and by replacing it by a variable:

- (13) Giovanni ha mangiato X (*Giovanni has eaten X*).

One might ask, at this point, what exactly is the (already available) information expressed by this presupposition. (13) contains a free variable, it is an open

formula, and an open formula does not seem to be able to express information. The usual answer to this question is that the presupposed information is the information expressed by the existential closure of the presupposition: in the present case by

(14) $\exists X$. Giovanni ha mangiato X,

i.e. by 'Giovanni ha mangiato qualcosa' ('*Giovanni has eaten something*'). Such an answer is not completely unproblematic, but we will ignore here the difficulties it gives rise to.

Now consider the following question:

(15) A. Chi ha mangiato caramelle? (*Who has eaten toffees?*)

It is obvious that (12B) would not be an appropriate answer to (15A). We get an appropriate answer if we move the main stress of the sentence:

(15) B. GIOVANNI ha mangiato caramelle.

The reason why (15B) is, whereas (12B) is not, an appropriate answer to (15A) should be clear. In this case the new information is conveyed by 'Giovanni', so 'Giovanni' must be the focus: but this is possible only if 'Giovanni' is the word containing the main stress of the sentence.

Let us go back to (10) and (11). Our hypothesis is that an indefinite NP cannot have generic reading unless it is contained in the presupposition of the sentence. Now, if (10) is pronounced with normal intonation, the main stress is on 'ricca':

(10') Truman Capote adulava una donna RICCA.

As a consequence, 'ricca' is certainly part of the focus of the sentence, and this implies in its turn that 'una donna ricca' cannot be entirely contained in the presupposition. If our hypothesis is correct, this is the reason why 'una donna ricca' cannot receive generic interpretation. Similarly for (11). If (11) is pronounced with normal intonation, the word containing the main stress is 'randagio':

(11') Una zitella protegge un gatto RANDAGIO.

Therefore, 'randagio' is part of the focus, and 'un gatto randagio' cannot belong entirely to the presupposition: so it cannot be understood generically.

We can conclude that examples like (10) and (11) (or, more precisely, (10') and (11')) are compatible both with the hypothesis that an indefinite cannot be generic if it is in the VP, and with the hypothesis that an indefinite can be generic only if it is contained in the presupposition. To distinguish between the two hypotheses, however, all we have to do is to reconsider (10) and (11) pronounced with a different intonation, for instance with the main stress on the verb:

(10'') Truman Capote ADULAVA una donna ricca

(11'') Una zitella PROTEGGE un gatto randagio.

(10'') is the proper intonation if (10) is to be used as an answer to the question: 'Come si comportava Truman Capote con una donna ricca?' ('*How was T.C.'s behaviour towards a rich woman?*'). Similarly, (11'') is the intonation appropriate if (11) is to be used as an answer to 'Come si comporta una zitella con un gatto randagio?' ('*How does a spinster behave towards a stray cat?*'). Now, it is clear

that, with this intonation, a generic reading of 'una donna ricca' and 'un gatto randagio' becomes possible. (10") *can* mean: Truman Capote used to flatter every rich woman (or at least: every rich woman he met); and (11") *can* mean: every spinster protects every stray cat (or at least: every stray cat she comes across). And this is exactly what we would have predicted on the basis of our hypothesis in terms of presupposition and focus. Since in (10") the main stress is on 'adulava', 'adulava' is the only possible focus: so 'una donna ricca' is included in the presupposition; this is why the generic reading is allowed. Likewise, in (11") the only possible focus is 'protegeva': so 'una zitella' and 'un gatto randagio' belong to the presupposition, and this is why both NPs can be assigned a generic reading. On the other hand, it is by no means clear how the availability of such a generic reading could be reconciled with the idea that NPs in the VP cannot be generic (unless, of course, one is ready to claim that the intonation of (10") and (11") corresponds to a syntactic structure in which, in spite of all appearances, 'una donna ricca' and 'un gatto randagio' are *not* in the VP).

So let us pretend that our hypothesis in terms of presupposition and focus is the correct one, and let us explore it a little further. The crucial question now is: *why* there should be such a correlation between the presupposition-focus structure of sentences and the conditions under which an indefinite can receive generic interpretation? Well, why not suppose that this question has an answer which is, in a sense, trivial? Perhaps, there is nothing to be explained. What we mean is that one might claim that the presupposition-focus distinction *as such* corresponds to a certain semantic mechanism which is *directly responsible* for the fact that only indefinites occurring in the presupposition can be understood as generics. We do not know whether such an idea is actually tenable. But it seems to us that it is much more plausible than it looks at first sight. So we will try to develop it a little.

II.3. Quantifying over circumstances

We will argue that the semantic interpretation of a sentence with a given presupposition-focus structure obeys the following principle:

Principle A. Let S be a sentence with a certain presupposition-focus structure and let P the (existential closure of the) corresponding presupposition. Then the meaning of S is: every circumstance in which P is true is a circumstance in which S is true.

This is of course a very rough formulation. But for the moment it can suffice. Let us see, by way of example, how (9), the logical form we have attributed to (5), can be derived from Principle A. If we pronounce (5) with normal intonation, the main accent is on 'serramanico'. So the focus of (5) can be the whole VP, and the corresponding presupposition will be 'uno scout X'. If we bind the variable X by means of an existential quantifier, we obtain ' $\exists X$. uno scout X', which means: 'there is a boy-scout (who is or is doing something)'. Now, Principle A tells us that (5), with the presupposition-focus structure we are considering, has the following meaning:

(16) every circumstance in which ' $\exists X$. uno scout X' is true is a circumstance in which (1) is true

or equivalently

(17) every circumstance in which there is a boy-scout who is or is doing something, is a circumstance in which there is a boy-scout who owns a spring-knife.

But obviously *every* circumstance in which a boy-scout is present is a circumstance in which that boy-scout is or is doing something (for instance, every circumstance containing a boy-scout is a circumstance in which that boy-scout is a boy-scout). So we can conclude that (17) is equivalent to (9).

So far so good. Now the problem is to show that Principle A is not purely ad hoc, that it can be used not only to derive logical forms like (9), but also to account for other facts. Consider the following examples:

- (18) In ufficio Maria giocava a CARTE (*In the office Maria used to play cards*)
 (19) Per darsi coraggio Garibaldi beveva GRAPPA (*To give himself courage Garibaldi used to drink grappa*)
 (20) Dopo pranzo Alberto fuma la PIPA (*After lunch Alberto smokes the pipe*)

Now let us modify these examples as follows:

- (18') Maria giocava a carte in UFFICIO (*Maria used to play cards in the office*)
 (19') Garibaldi beveva grappa per darsi CORAGGIO (*Garibaldi used to drink grappa to give himself courage*)
 (20') Alberto fuma la pipa dopo PRANZO (*Alberto smokes the pipe after lunch*).

Each sentence of the first set has an interpretation which the corresponding sentence of the second set does not have; and vice versa. Let us take (18) and (18'). Among the possible interpretations of (18) we find the following:

(21) whenever Maria was in the office, she played cards.

Now, this meaning *cannot* be expressed by means of (18'). On the other hand, (18') can mean:

(22) whenever Maria played cards, she was in the office,

or, to put it otherwise: Maria played cards only in the office. And *this* meaning cannot be expressed by means of (18). The other two pairs of examples exhibit a similar contrast. How can we explain such a contrast? Well, if we assume that the interpretation of a sentence having a certain presupposition-focus structure obeys Principle A, a contrast of the kind described is exactly what we would expect. In (18) the main stress is on 'carte'. As a consequence, the focus of (18) can be 'giocava a carte'. The corresponding presupposition is 'in ufficio Maria X'. So Principle A tells us that we can interpret (18) as follows: every circumstance in which Maria was in the office is a circumstance in which Maria was in the office and played cards, and this is exactly the interpretation of (18) paraphrased in (21). Now let us turn to (18'). Here the main stress is on 'ufficio': therefore, 'ufficio' must be part of the focus and 'in ufficio' cannot belong to the presupposition; so in this case the application of Principle A cannot produce the interpretation given in (21). On the other hand, if we assume that the focus of (18') is 'in ufficio' and that the presupposition is 'Maria gioca a carte X', Principle A gives us the interpretation according to which Maria played cards only in the office, i.e. (22). Finally, since in (18) the main stress is on 'carte', 'carte' must be part of the focus, and 'Maria giocava a carte' cannot be the presupposition: this is why in this case the interpretation expressed in (22) is unavailable.

III. METATHEORY/ RULES, EXCEPTIONS AND VAGUENESS

Examples (18)-(20) and (18')-(20') still are of a rather special kind. But it is not hard to see that Principle A produces acceptable interpretations no matter what sentence we are considering and no matter how the boundary between focus and presupposition is drawn. A point which deserves emphasis is that Principle A can be applied when the presupposition of a sentence contains an indefinite NP but does not coincide with it. This observation can help us to unravel some of the tangles which surround the topic of generic indefinites.

It is often claimed that an analysis of generic indefinites in terms of universal quantification is wrong. Several different arguments are used to support this claim. The weakest is the argument based on the fact that the truth of a sentence containing a generic indefinite seems to be compatible with the existence of counterexamples (for instance, the truth of 'A pirate used to drink rum' seems to be compatible with the existence of some teetotal pirate). We find this argument weak because a sentence which is strictly speaking false can nevertheless be acceptable for pragmatic reasons easy to understand. But the argument based on the existence of counterexamples has a more embarrassing variant.

III.1. Eagles and Ravens

Let us go back to our example (8): 'Un' aquila costruisce il nido sulla roccia'. One might reason as follows. Sentence (8) is true. But most eagles die too young to build nests. Moreover, only male eagles build nests (we do not know whether this is actually the case; but let us suppose it is). This shows that an analysis of (8) in terms of a universal quantification is wrong. (8) does *not* mean that every eagle builds its nest on the rock. Not only are there counterexamples. The point is that *most* eagles are counterexamples. Several different solutions are suggested: some authors claim that a sentence like (8) must be understood not as about eagles, but as about the natural kind 'eagle'; others tell us that we must resort to default logics, etc. This line of reasoning is quite common in the literature on generics. Common but not unavoidable. If what we have been saying so far is by and large correct, the compatibility of (8) with the fact that most eagles do not build nests can be accounted for in a much simpler way. Let us reconsider our sentence. The main accent is on 'roccia', so 'roccia' must be contained in the focus. But this does not yet suffice to determine the focus. We can choose. As the focus we can take 'costruisce il nido sulla roccia': what we then obtain, by means of Principle A, is the interpretation paraphrased in (8'), equivalent to a universal quantification over eagles: an interpretation which, in the light of the reasoning sketched above, can perhaps be regarded as problematic. But as the focus of the sentence we can also choose 'sulla roccia'. And it is easily seen that, in this case, the interpretation produced by Principle A is the following:

(23) every circumstance in which there is an eagle that builds a nest is a circumstance in which there is an eagle that builds a nest on the rock.

It should be obvious that (23) can be true - literally true! - even if most eagles never build nests. (23) is not about all eagles: it is about all eagles that build nests. Admittedly, an analysis of this sort cannot be extended to every problematic case. But the fact that at least *some* problematic cases can be dealt with in this way - without any need to enrich our ontology by the introduction of natural kinds, without resorting to default logics, and so on - seems to us to be important.

Our perplexities about default theories of generics deserves some qualifications. Actually, our approach is quite compatible with the view that many generic statements are used "by default". We are indeed aware that, to account for the existence of obvious counterexamples, the universal quantification over circumstances - on which our treatment of generics is founded - must be suitably restricted to a set of *relevant* circumstances. What we add is that this *pragmatic* feature is not a peculiarity of generics, but a general phenomenon which, as such, is not an essential component of a *semantic* theory of this class of expressions. Let us deepen this point a little more.

Take a generic statement like:

(24) Un corvo è un uccello di grandi dimensioni (*A raven is a large bird*).

Now, according to the default theories of generics, a sentence such as (24) licenses inferences that are *valid* and, at the same time, *defeasible*. The point is that from (24) and, for example,

(25) Paw è un corvo (*Paw is a raven*)

the conclusion

(26) Paw è un uccello di grandi dimensioni (*Paw is a large bird*)

is considered as sound, even if there are well-known counterexamples: if Paw is, for example, a new-born raven, it is not, of course, a large bird; this is why the inference at issue is said to be defeasible, even if valid.

Our idea is that we are here in front of an alternative.

(i) We can say that (24) - (26) is a shortened version of a more explicit (and complex) inference, which is really valid, but not defeasible. For example, this extended version of the inference would make it explicit that a raven is a large bird if it is an adult specimen, so that the case of the new-born raven is ruled out. As to the usual *semantic* notions (truth, entailment, validity, etc.), their classical properties (monotonicity, in particular) are preserved.

(ii) Otherwise, we can take (24) - (26) *as such*; but in this case what we can attribute to this kind of inferences is just a sort of *pseudo-validity*: which really is, this time, a non-monotonic notion. Now, this notion of pseudo-validity is a very interesting one, and it is worth investigating if we want to reconstruct the way people make inferences in their everyday life. But it belongs in pragmatics, not semantics. To realize it, consider the following facts.

III.2. Charity Principles.

(a) The existence of counterexamples is licensed by a "liberal" use of words which is not confined, of course, to generics. Some vagueness characterizes the way we use words in general. Consider this example:

(27) Raven = Large bird [...] with lustrous black feathers [...]. (Webster Dict.)

Such a statement is not, formally speaking, a generic, but rather a metalinguistic definition. What is interesting, from our point of view, is that it presents the same kind of problems as our original generic sentence (24). The point is that (27) is pragmatically acceptable even if patently incomplete (for the restriction to the case of adult specimen is omitted). This kind of liberality - which is manifest even in the way we *define* a word - is a very general pragmatic phenomenon, quite independent

of generics. A certain latitude in the use of words is witnessed by other kinds of examples:

(28) Jim *always* smokes

(29) *Every* day, exactly at five o'clock, Kant used to go out for a walk.

Countless exceptions must be admitted in both cases, since, for example, Jim does not smoke when he is sleeping, neither did Kant go out when he was in bed because of an illness. Should we say that 'always' and 'every' do not have their intended meaning (which amounts to a universal quantification), and that the "default" interpretation is part of their *semantics*? We do not think so. And we believe that the same holds of generics. What an expressions *means* should be kept distinct from the ways it can be *used*.

(b) The pragmatic nature of the relationship between generics and default is shown by the *contextual* character of the admissibility of counterexamples. In some cognitive situations, as we have just seen, (many) exceptions are allowed. In others, no counterexample is admitted: for instance in arithmetics, where two sentences like:

(30) Un numero primo è divisibile per due (*A prime number is not divisible by two*)

(31) Ogni numero primo è divisibile per due (*Every prime number etc.*)

are *definitely* false for the *same* reason, i.e. the existence of a *single* exception. This point is also illustrated by examples of a quite different kind. The following statements are true for in both cases no counterexample to the nested generic is admitted:

(32) It is false that a cat is unable to swim. For yesterday I saw Sylvester cross the river.

(33) It is false that every cat is unable to swim. For yesterday I saw Sylvester cross the river.

That means that the admissibility of counterexamples does not characterizes the generics *as such* (otherwise (32) would be true), but it depends on the cognitive context at issue.

These examples are interesting because they show that when the default is, as it were, empty (i.e. when counterexamples are ruled out by the context, like in arithmetics), generic and universal statements are roughly equivalent. But since universal statements can *themselves* have, in natural language, a default interpretation (as witnessed by examples like (28) and (29)), it is plausible to think that - when the default is not empty - a generic statement can be reduced to a universal statement (of a particular kind, as we shall see in a moment) *plus* default.

(c) What we have just said seems to be confirmed by the formal treatments of generics in non-monotonic logics, where such a statement as (24) is often translated into logical forms like:

For *every* "typical" raven...

For *every* "normal" world, if x is a raven in that world, then ...

and so on;

that is, some kind of universal quantification (plus default) is commonly used to reconstruct the intended meaning.

III.3. Special Operators: A Problem

But suppose now that the default interpretation is considered as an essential feature of the *semantics* of generics and that a primitive operator is introduced to this aim. Let GEN be such an operator (as suggested, for instance, by M. Krifka), so that a statement of paradigm:

(34) An A is B

is translated into something like:

(35) GEN(A(x) : B(x)).

Now, what happens in that, when we have to interpret formal structures like (35), the natural solution is, once again, to resort to some kind of universal quantification (plus default). For example:

(35') For every x, if x is A and it is not known that x is not B, then x is B;

(35'') Every situation in which x is A is, unless otherwise demonstrated, a situation in which x is B

and so on.

But the point is that universal statements, in natural languages, can *themselves* have a default interpretation, so that they can be associated with logical forms of the type we have just seen. From this point of view, there is no distinction with respect to generics; in both cases we have two components: (i) some kind of universal quantification (over individuals, or circumstances, etc.); (ii) the specification of the "default" character of this quantification. Our proposal is to consider (i), i.e. a universal quantification over circumstances, as an essential *semantic* feature of generics (which is to be accounted for in the relevant logical forms), whilst (ii) is a more general, pragmatic feature - not confined to generics, which must each time be fixed contextually. So no special operator is required: quantification over circumstances is all we need, apart from the relevant pragmatic specifications. (We will see later some differences between plain quantification over individuals and generics as cases of quantification over circumstances.)

Furthermore, if the default interpretation were part of the intrinsic meaning of generics, a statement like (30) would turn out to be necessarily true (against our intuition), for if its sense is given by something like (35') or (35''), the existing counterexample (number two) is neutralized by the fact it is known to be divisible by two. What is worse, in view of the same argument (applied to each odd number), the statement

(30*) An odd number is divisible by two

would turn out to be necessarily true!

III.4. Some metatheoretical speculations

What shall we conclude from these remarks? First of all, that a semantic treatment of generics in terms of universal quantification over circumstances makes sense, in many cases, only if the domain of quantification is suitably restricted to a set of relevant circumstances. But this notion of relevance is not part of semantics, for it varies from context to context. As often happens in pragmatics, we must face a phenomenon which is not easily definable. This is why, in our opinion, the term

'default' is unintentionally used to cover different kinds of data. Here are some examples.

(a) *Gestalt*. We have learnt from psychology that there is a natural tendency, in human mind, to disregard possible irregularities in favour of simple patterns of thought. For instance, we perceive as a circle what actually is an incomplete shape. Why not think that a similar tendency to pick regularities makes a statement like:

(36) Un corvo è nero (*A raven is black*)

true even if the existence of albino ravens is well-known?

(b) *Ellipsis*. To make communication more efficient, some specifications that are taken for granted can be omitted. This is the case of our example (24), where it is clear from the context that we are speaking of adult ravens.

(c) *Presupposed information*. Common knowledge does not need to be made explicit. For instance, the well-known fact that only a female creature procreates and suckles its young is presupposed in a statement like:

(37) Un mammifero allatta i suoi piccoli (*A mammal suckles its young*)

which is true (or acceptable) in spite of lots of counterexamples, for male mammals do not suckle, of course.

(38) *Contextual dependence*. If in my department I say:

(39) Uno studente del Corso di Logica apprezza il manuale di Smullyan (*A student from the Logic Course appreciates Smullyan's textbook*)

this is not a statement about any possible student who attends a Logic Course (with obvious counterexamples). What makes exceptions unlikely is the reference to a *given* restricted domain.

(e) *Tenses*. To stress the relationship between generics and default, it is often said that such a sentence as:

(39) Una tartaruga vive a lungo (*A turtle lives a long time*)

is accepted as true in spite of countless exceptions. Once more, this is a more general phenomenon which is not peculiar to generics, but it is inherited by them. The point is that a tense like the present often expresses something *virtual* (determined by cultural, biological, etc., reasons), rather than an actual state. To consider a quite different example, suppose I have in front of me a brand-new computer which has never been used, and that you are interested in a particular sequence of operations written on the blackboard. Well, if I have all the relevant information, by referring to *this* computer I can truthfully say:

(40) This machine computes the sequence on the blackboard in less than three seconds

although it will *never* compute (and has never computed) that sequence. What I mean, by uttering (40), is just a virtuality. And the same happens with (39).

Caveat. What we are arguing for is that in some cases a generic has a *modal* dimension (expressed by a particular tense), not that it always has it. For instance, no idea of necessity (or intrinsic virtuality) is expressed by sentences like (5) or (38) on its generic interpretation.

Closely related to this thesis concerning the alleged modal dimension of the generic reading is the claim that the generic reading of an indefinite noun phrase is possible only when the predicate expresses an essential property. To support this claim, Krifka, for example, observes that

(41) A madrigal is polyphonic

is perfectly acceptable, whilst

? (42) A madrigal is popular

is not. The reason - Krifka says - is that being polyphonic is an essential property, while being popular is an accidental one. (This example is actually due to Lawler). Now, we are not sure that this explanation is correct. As a matter of fact, there are generic statements in which the property of being popular is involved. In Italian, the following sentence has a generic reading (actually, it has *only* this reading, since the relative clause is in the subjunctive mood):

(43) Un filosofo che sia apparso al Maurizio Costanzo Show è popolare (A philosopher who has appeared (*subj.*) at the Maurizio Costanzo Show is popular).

And what about

(44) Un quadro di Van Gogh è costoso (A painting by Van Gogh is expensive) ?

The notion of essential property is vague, but it would be very difficult to argue that our last two examples, on their generic reading, refer to essential properties.

III.5. Kinds of Universal Quantification. Differences in Use.

The aim of our reflections on the existence of counterexamples was to remove an obstacle to the reconstruction of generics in terms of universal quantification over circumstances. On the other hand, we spoke of a *rough* equivalence between a generic like 'An A is B' and a plain universal quantification (over individuals) like 'Every A is B'. But the adjective which qualified our statement was not casual, for some subtle differences in the pragmatics of these two kinds of expressions are to be taken into account.

If what we said at the outset is correct, the use of a universally quantified NP is characterized by the reference to a *given* context. (Recall our example with the NP 'All the dogs'.) This contextual restriction is one of the reasons why a plain universal quantifier is inappropriate to express a "generic" statement. As we have just said, a given context is required, and there is a reference to all the individuals of *this* context. On the contrary the universal import of generics is determined by a universal quantification over circumstances or contexts (the circumstances described by the indefinite NP). So we have here two symmetric situations: in the case of plain universal quantifiers, the *intended* context is referred to in order to introduce a restriction to the domain the quantifier ranges over, whilst in a generic statement the (*indefinite*) NP is used to restrict the class of circumstances we are quantifying over.

The reference to a *given* context makes a plain universal quantification odd if used for a definition, i.e. when a (ideally) new class of entities is introduced. Just to give a simple example, consider the oddity of:

? (45) Every even number is a number divisible by two (if used for a definition)

whilst

(46) An even number is a number divisible by two

is perfectly in order if the statement at issue is a definition. (Notice that (45) is quite appropriate to express a theorem concerning a class of numbers (i.e. the even numbers) which has *already* been introduced into the domain of discourse.)

We have the same contrast in a different kind of situations, as witnessed by the contrast between:

? (47) Every raven is a large bird with lustrous etc.

and:

(48) A raven is a large bird with lustrous etc.

Sentences like (47) are hardly used to give a definition (think of a dictionary, for instance), whereas a generic like (48) is quite appropriate.

What makes a generic suitable to introduce a definition or to express a regularity, a modal notion, etc., is its referring to a whole set of (conceivable) circumstances or contexts, not only to a single context as in the case of a definite NP.

No accident if this kind of contrast disappears when the generic too is to be interpreted as a universal quantification over circumstances with respect to a *given* context. Notice the affinity, this time, between:

(49) L'anno scorso, al campeggio, ogni boy-scout diceva le preghiere (*Last year, at the camp, every boy-scout used to say grace*)

and:

(50) L'anno scorso, al campeggio, un boy-scout diceva le preghiere (*Last year, at the camp, a boy-scout used to say grace*) (on its generic reading)

where the context is fixed by the phrases 'last year', 'at the camp'. In this case there is a strong similarity in meaning between the statement with the plain universal quantifier and the generic statement, for the latter is to be interpreted in terms of a universal quantification over the subcontexts of the given context (recall the suggested translation: For every circumstance, last year, at the camp, in which there is a boy-scout, this boy-scout says grace), which involves exactly the boy-scouts of that context (last year, at the camp). But when no particular context is referred to (as it happens, typically, in the case of definitions) the contrast is, once gain, perceivable. Consider:

? (51) Ogni boy-scout è un membro dell'organizzazione fondata da Baden-Powell (*Every boy-scout is a member of the organization founded by Baden-Powell*)

vs

(52) Un boy-scout è un membro dell'organizzazione fondata da Baden-Powell (A boy-scout is a member of the organization founded by Baden-Powell).

The conclusion is that, in spite of the "rough" equivalence between a generic and a plain universal quantification there are cases in which only one of them can be properly used. We have qualified this phenomenon as pragmatic for the relevant conditions of appropriateness are largely dependent on the situation of discourse. For instance, contra the idea that a generic must refer to an *open* class of individuals, we have already given a couple of examples in which the objects involved by the generic statement are not an open class, but a set of objects *delimited* by a given context (see examples (38) and (50).) On the other hand, it should be clear that, if this set of entities is very small (if for instance only two students attend our Logic Course, or if there were just three boy-scouts at the camp, last year) the possibility of a generic reading is quite remote. But it is hard to imagine a clear-cut criterion that, independently of the particular situations of discourse, can filter out all the cases in which the generic reading of the indefinite NP sounds improper.

Our (metatheoretical) explanation of these phenomena is that an *indefinite* NP can be used in a generic statement to restrict the intended quantification over contexts or circumstances, whilst a *definite* NP - which on the contrary needs the reference to a given context - cannot have a generic interpretation. In fact, the following examples (with indefinite NPs) have a generic reading:

- (53) (a) Un cinghiale causa danni (*A wild boar causes damage*)
 (b) Due cinghiali causano danni (*Two wild boars cause damage*)
 (c) Molti cinghiali causano danni (*Many wild boars cause damage*).

This generic interpretation can be paraphrased as: In every circumstance in which a wild boar is present (two wild boars are present, many wild boars are present), this wild boar (these wild boars) cause(s) damage. But this kind of interpretation is not allowed with a *definite* NP. Consider:

- ? (54) (a) Tutti i cinghiali causano danni (*All the wild boars cause damage*)
 (b) La maggior parte dei cinghiali creano danni (*Most wild boars cause damage*).

These statements cannot have the interpretation whose paraphrase is: in every situation in which all the boars (most wild boars) are present, these wild boars cause damage.

IV. EXHAUSTIVENESS CONDITIONS

To provide further evidence for the correctness of our Principle A, we turn now to a phenomenon which has at first sight nothing to do with the problem of generic indefinites. We might call this phenomenon the Exhaustiveness Condition of answers. Take example (15B). We observed above that the focus - i.e. 'Giovanni' - is the part of the sentence which conveys new information. But we must now observe also that the information conveyed by the focus must exhaust, so to speak, all the relevant information. This means, in practice, that (15B) is not an appropriate answer to (15A) unless Giovanni is *the only one* who has eaten toffees. This is the Condition of Exhaustiveness implicit in (15B). A qualification is in order here. Obviously, (15B) does not imply that nobody in the world has ever eaten toffees except Giovanni. But it does imply that nobody except Giovanni has eaten toffees within the small portion of reality the discourse is about. Imagine the following

circumstance. Maria's mother finds out that a lot of toffees have disappeared from the cupboard. She asks her daughter: 'Chi ha mangiato caramelle?' ('*Who has eaten toffees?*'). Maria says: 'GIOVANNI (ha mangiato caramelle)'. Now suppose that some of the toffees have been eaten by Giovanni (Maria's brother), and some by Maria herself. Then it is true that Giovanni has eaten toffees. Nevertheless, Maria is not telling the truth: she is actually lying. Why? Well, because of the Exhaustiveness Condition of answers. What her answer implies is not only that Giovanni has eaten toffees, but also that Giovanni is the only toffee-eater around. Now, it seems to us that the Condition of Exhaustiveness can be accounted for by means of Principle A. Let us apply the principle to the analysis of (15B). What we get is roughly the following:

(55) Every (relevant) circumstance in which somebody has eaten toffees is a circumstance in which Giovanni has eaten toffees.

IV.1. A First Approximation

But (56) is just a roundabout way to say that nobody except Giovanni has eaten toffees (if somebody else had eaten toffees, there would be at least one circumstance in which somebody has eaten toffees and Giovanni has not eaten toffees).

A slightly more complex example. Consider

- (56) A. Chi aiutava l' insegnante a mettersi il cappotto? (*Who helped the teacher to put his coat on?*)
 B. Uno SCOLARO aiutava l' insegnante a mettersi il cappotto (*A schoolboy helped the teacher etc.*)

In this case the answer can be understood in two different ways. Both interpretations, however, involve a Condition of Exhaustiveness. First interpretation: When the teacher put his coat on, he was always helped by a schoolboy, and *only by a schoolboy* (although not necessarily by the same schoolboy every time). Second interpretation: there was a schoolboy x such that x , and *only* x , always helped the teacher to put his coat on. The derivation of the first interpretation from Principle A is completely unproblematic. Since the focus of (56B) is 'uno scolaro', a straightforward application of the principle gives us

(57) Every circumstance in which somebody helped the teacher to put his coat on was a circumstance in which a schoolboy helped the teacher to put his coat on.

This is exactly what we wanted. What about the other interpretation of (56B)? Well, in order to derive it, all we have to do is to realize that the existential quantification corresponding to the indefinite NP 'uno scolaro' can also have wide scope with respect to the universal quantification over circumstances introduced by Principle A. What we get in this way is:

(58) For some schoolboy x : every circumstance in which somebody helped the teacher to put his coat on, was a circumstance in which x helped the teacher to put his coat on.

A moment's reflection shows that once again this is just what we wanted. Examples could and should be multiplied. But the two examples discussed above should have been sufficient to explain in which way we think that Principle A can be used to justify the Condition of Exhaustivity. More generally, we hope that, in

the light of what we have said, the idea that Principle A captures the semantic content of the presupposition-focus distinction is not completely implausible. A word of caution is perhaps in order here. The terms 'presupposition' and 'focus' have been used in the analysis of a wide range of linguistic facts. We do not intend to claim that all these facts can be accounted for on the basis of something like Principle A. On the contrary, we would be very much surprised if all those facts turned out to be reducible to a unitary explanation. What we claim is that our Principle A reflects an important semantic mechanism which is systematically related to the intonation of sentences, and can explain at least some of the phenomena traditionally reduced under the labels of 'presupposition' and 'focus'. It would be interesting to rediscuss, from this point of view, the traditional characterization in terms of old vs new information, the so-called 'contrastive' value of focus, and so on. But all this must be left for another occasion. To complete the picture, we just add a remark, and point to some of the - alas! - many flies in the ointment.

IV.2. The Role of Tense

Unfortunately, it is false that the semantic content of the presupposition-focus structure is always determined by Principle A. Here is an unmistakable counterexample:

- (59) A. Chi (nel nostro dipartimento) ha scritto un libro su Kant? (*Who (in our department) has written a book on Kant?*)
 B. Quattro ASSOCIATI hanno scritto un libro su Kant (*Four associated professors have written a book on Kant*).

In (59B) the focus is 'quattro associati' and the presupposition is 'X hanno scritto un libro su Kant'. If we apply Principle A, we obtain the following interpretation: every circumstance containing a member of our department who has written a book on Kant is a circumstance in which four associated professors have written a book on Kant. Which means: every book on Kant written by some member of our department has been written by four associated professors. Now, such an interpretation of (59B) is extremely implausible. What (59B) really means is: there are four associated professors such that each of them has written a book on Kant and they are the only members of our department who have written books on Kant. Unfortunately, there is no way to derive *this* interpretation of (59B) from Principle A.

We need something else. We are forced to admit that the semantic content of a presupposition-focus structure can also be determined by a principle of the following kind:

Principle B. Let S be a sentence with a given presupposition-focus structure, and let P be the (existential closure of the) corresponding presupposition. Then the meaning of S is: the "union" of the circumstances in which P is true is a circumstance in which S is true.

Here by the 'union' of a set of circumstances A we mean the circumstance we obtain by amalgamating, so to speak, the circumstances which are elements of A. For instance, the union of the circumstances satisfying the presupposition of (59B), i.e. the union of the circumstances in which 'X.X has written a book on Kant' is true, is the circumstance containing every member of our department who has written a book on Kant, every book on Kant written by a member of our

department, and nothing else. It is easy to see, then, that Principle B accounts for the exhaustiveness condition expressed by (59B).

But now the problem is: is there any rule to establish when a certain presupposition-focus structure is to be interpreted according to Principle A and when it is to be interpreted according to Principle B? This is a difficult question. To a certain extent, the choice between the two principles seems to depend on the tense. In particular, with 'simple past' and 'present perfect' the application of Principle A seems to be, if not impossible, at least unnatural. In deriving the exhaustiveness condition for example (15B) above from Principle A, we were actually cheating. We could do it because in the case of (15B) Principle A and Principle B give almost indiscernible results. But since the tense of (15B) is present perfect, Principle B would be more appropriate. On the other hand, with imperfect, Principle A turns out to be all right, as is shown by the examples discussed above. But we must not jump to the conclusion that there is a simple and direct correlation between the tense and the semantic principle to be applied. Imperfect is also compatible with Principle B, at least under certain conditions. Take

- (60) A. Chi aveva una moglie sudamericana? (*Who had a South-American wife?*)
 B. Due MARINAI avevano una moglie sudamericana (*Two sailors had a South-American wife*).

In (60) 'una moglie sudamericana' is in the scope of 'due marinai', and this rules out the possibility of applying here Principle A. Principle B, on the other hand, gives the right interpretation. It is certainly tempting to conjecture that Principle A and Principle B result from the interaction of two elements: a common element which is the contribution of the presupposition-focus structure, and an element derived in some way from the tense. Unfortunately, we do not see at the moment how this idea could be developed in detail.

Notice, however, that if we were able to explain why the present perfect requires Principle B, we would have an explanation of why the present perfect is incompatible - or almost incompatible - with the generic reading of indefinites. For a moment's reflection shows that the logical form we have attributed to sentences with generic indefinites *cannot* be derived on the basis of Principle B. Let us go back once again to our example (5). What do we get if we try to interpret (1) by applying Principle B? We get

- (61) The union of the circumstances in which there is a boy-scout is a circumstance in which there is a boy-scout who owns a spring-knife.

It goes without saying that (61) is not the interpretation of (5) we are interested in.

IV.3. Notes on Formalization

The key role in our analyses is played by the notion of context or circumstance. So the first problem is how to provide a formal representation of circumstances (of a given state of the world). What we propose is to identify them with "submodels" of a given model. A submodel is, roughly speaking, a fragment of a classical model. A classical model has a fixed domain: given any individual in the domain and any property expressible in the language, the model specifies whether that individual has that property or not. Similarly for relations. A submodel only contains partial information. Its domain is (a subset of) the domain of the classical model to which it belongs. For some elements of its domain the submodel can specify that they have or do not have some property, or that they are or are not related by some relation. But in a partial model the extension and the

counterextension of properties and relation can remain partially, or completely, undetermined, even with respect to the domain of the submodel itself. Suppose x is an element of the domain of a partial model M : M can contain the information that x is a boy-scout, or it can contain the information that x is not a boy-scout, but it can also leave the question completely undecided.

Let us turn now to the problem of how a logical form like (9) can be translated into formal terms. Since we are identifying circumstances with partial models, we might be tempted to translate (9) as follows:

(62) For every partial model M : if there is an x such that in M x is a boy-scout, then there are x and y such that in M x is a boy-scout, y is a spring-knife, and x owns y .

(62) seems to reproduce faithfully the content of (9). But unfortunately (62) does not work. The reason is simple. (62) turns out to be false unless the set of boy-scouts is empty. Let us assume that the big model representing reality (or at least, the part of reality we are interested in) contains a boy-scout, say Walter. Then there exists a partial model L with the following properties: 1) the only element of the domain of L is Walter; 2) the only information contained in L is that Walter is a boy-scout. L is among the partial models quantified over in (62). So, all we have to do to prove that (62) is false is to show that there are no x and no y such that in M x is a boy-scout, y is a spring-knife, and x owns y . But this is obvious: L has been defined in such a way that no spring-knife is present in it.

Is there any way out? Well, we might resort to the notion of 'expansion' of a submodel. Let M and N be two submodels. We say that N is an expansion of M if the domain of M is a subset of the domain of N and if N contains at least all the information already contained in M . A reformulation of (62) could now be the following:

(63) For every partial model M : if there is an x such that in M x is a boy-scout, then there is an expansion N of M and there are x and y such that in N x is a boy-scout, y is a spring-knife, and x owns y .

At first sight, (63) avoids the difficulty pointed out above. We do not have to worry any longer about the fact that partial models containing a boy-scout may be too small to contain also the boy-scout's spring-knife, since we are now allowed to look beyond the partial models we start with and to take also their expansions into account. But a closer look at (63) reveals that it too involves a difficulty. Let us assume that the big model representing reality contains a boy-scout - Edward - who owns a spring-knife, and many other boy-scouts who do *not* own a spring-knife. Then sentence (5), with 'uno scout' interpreted as a generic, is obviously false. But (63) turns out to be true, and so it cannot be taken as an adequate logical form of (5). The reason why, under the circumstances described, (63) is true is the following. Let M be any partial model whatsoever. Let us construct a new partial model M' which is exactly like M except that: 1) Edward and his spring-knife (call it Donald) are added to the domain; 2) M' contains all the information contained in M plus the following: Edward is a boy-scout, Donald is a spring-knife, and Edward owns Donald. It is clear that M' is an expansion of M and that in M' there are x and y such that x is a boy-scout, y is a spring-knife, and x owns y . Since the construction of M' can always be performed, no matter what the initial model M is, we have proved that every partial model - and a fortiori every partial model containing a boy-scout - can be expanded to another partial model in which there is a boy-scout who owns a spring-knife. So (63) is true.

This difficulty arises because in (63) we have no guarantee that the boy-scout who is assumed to exist in a partial model M is identical with the boy-scout who is said

to exist in the extension N of M. So what we have to do is to modify (63) in such a way as to make this identity explicitly. The result of such a modification can be the following:

(64) For every partial model M: if there is an x such that in M x is a boy-scout, then there is an expansion N of M and there are x and y such that in M x is a boy-scout, and in N y is a spring-knife, and x owns y.

As far as we can see, (64) is acceptable, at least as a first approximation. Needless to say, many other problems arise as soon as we attempt to extend the formalization in a systematic way. But we will not go into further details. Our aim here was simply to give you an inkling of the sort of difficulty that face us, and of the strategies we can adopt to solve them.

Another interesting problem is the following, which has to do with the Exhaustiveness Condition. As a matter of fact, to account for this condition we have resorted to Principles A and B, according to the relevant tense. Now, in both cases our formalism has to face the same predicament. Let us illustrate it with respect to Principle (B).

Consider our example (15) once again. To deal with the exhaustiveness assertion that it entails, we suggested to associate with (15) a logical form which is roughly equivalent to:

(15') The "union" of the (relevant) circumstances in which it is true that somebody has eaten toffees has an expansion in which it is true that Giovanni has eaten toffees.

[Actually, for the sake of simplicity, we did not use the notion of expansion. But now know that it is an essential part of our formalism.]

The idea was that (15') is true iff nobody except Giovanni has eaten toffees - in consonance with the Exhaustiveness Condition. Unfortunately (15') turns out to be inadequate from this point of view. To realize it, consider a state of affair in which Giovanni *and* Paolo have eaten toffees. Well, it is easy to see that the union of the circumstances in which it is true that somebody has eaten toffees (i.e. a circumstance in which Giovanni and Paolo have the property of eating toffees) has an expansion in which it is true that Giovanni has eaten toffees. In other terms, (15') is true even if two people have the property at issue: which is a violation of the Exhaustiveness Condition.

To avoid this difficulty, we must replace the notion of truth in a circumstance with a more fine-grained one. Let us say that a circumstance c is *characterized* by a statement A iff A is true in c and, for every c' such that $c \geq c'$ (i.e. for every submodel c' which is "smaller" than c), if A is true in c', then $c = c'$. Principle B can now be expressed as follows:

Principle B. The "union" of the circumstances that are characterized by P has an expansion which is characterized by S.

As a result, (15') is replaced by:

(15*) The union of the (relevant) circumstances which are characterized by the statement that somebody has eaten toffees has an expansion which is characterized by the statement that Giovanni has eaten toffees.

It can be easily proved that this new version meets the Exhaustiveness Condition. [A similar modification of Principle A is needed.]

APPENDIX

Contexts and Demonstratives

1. Bare Contexts

Reflecting on the first sections of our paper, the reader might argue as follows. As far as good. But what you need is, after all, a sort of parametrization of the noun phrases. And a well-known device to get this result is already available for a specific case: Kaplan's contexts provide the parameters needed to account for indexicals. If you introduce other kinds of contexts the resulting picture might prove messy: how many theoretical entities are required? So a couple of points are to be considered: (i) in what sense the notion of context you have in mind is not reducible to the notion already available in the analysis of indexicals? (ii) conversely, is it possible to account also for indexicals in the framework you are going to build up?

In such a logic of demonstratives (or indexicals, more in general) as Kaplan's, contexts have a very *austere* nature. They are elegant and simple devices which provide the parameters needed by indexical expressions. For example, with respect to a given utterance, a context specifies the value of the parameter: AGENT (of the utterance), in order to assign a referent to the indexical expression 'I'; it specifies the value of parameter: TIME, in order to assign a referent to the indexical expression 'now'; and so on. From this point of view, "context is a package of whatever parameters one needs to determine the referent [...] of the directly referential expressions of the language" (Kaplan, 1989: 591).

Formally speaking, in this logic of demonstratives, contexts are *new* theoretical entities which are *added* to the usual machinery of model-theoretical semantics in order to account for indexicals. This *ad hoc* character of contexts is justified by the nature of the expressions the logic of demonstratives is about: "pure" indexicals (as 'I', 'here', 'now', etc.) and "true" demonstratives ('that'), that is expressions whose reference is determined by the associated *demonstration*.

But it is not clear how this machinery works if other, less "pure", expressions are taken into account. For example, we are not sure that the most typical (and common) use of demonstratives like 'that' is characterized by the demonstration, which is an essential ingredient of "true" demonstratives in Kaplan's sense. In most cases, the demonstration is *not* essential (and, in fact, is completely absent), for the interaction between the built in sortal ('book', for instance) and the context is sufficient to determine the intended referent. If you say 'This book is boring' you are likely to use no demonstration, since (usually) the situation in which the utterance takes place is definite enough to specify what is the relevant book. In our opinion, it is easier to conceive of uses of demonstrative phrases without an associated demonstration than to imagine cases in which what is lacking is the built-in sortal. And the reason is obvious: any demonstration is potentially ambiguous without an associated sortal term. Suppose you say 'That is expensive' pointing to a table. If, for instance, a book is on the table, a possible way of reacting to your utterance might be: 'That *what?* The book? The table?'. And the specification of the suitable sortal term (an (expanded) common noun, in most cases) will settle the matter.

2. Furnished Contexts

But this is not the point. In other terms, we do not have to discuss here whether the most typical use of demonstratives entails demonstration rather than the essential presence of a sortal. Neither can we tackle a more general question, i.e. the question: is the notion of direct reference really a workable one? The question, rather, is that "impure" demonstratives (i.e. phrases with a built-in sortal as 'this book' or 'that ashtray on the table') are very common in ordinary discourse. Now, to account for the way they are used, the notion of context as a "package" of parameters, is no longer sufficient. What we need, now, is something with a more articulated structure: local states of affairs in which individuals have properties (recall 'this *book*') and relations (recall 'that ashtray *on* the table'), and are identified thanks to those properties and relations.

It is instructive to see how a class of expressions related to impure demonstratives is treated in Kaplan's framework. What we have in mind is, roughly speaking, the class of definite descriptions used in a demonstrative manner. As is well-known, Kaplan has a special operator to generate these expressions: given a definite description α , '*dthat* (α)' is a term whose referent is determined both by the demonstration and the property the description expresses. More exactly, ignoring here for the sake of simplicity such parameters as time and assignation, the denotation of a *dthat*-term is given by the clause:

$$|dthat(\alpha)|_{c,W} = |\alpha|_{c,c(W)}$$

which means: the denotation of '*dthat* (α)' in the world W with respect to the context c is the denotation of the description α in the world of the context c . But notice that $c(W)$ - the world of the context c , which may not coincide with W - is, in model-theoretic terms, a complete or global state of affairs: something related to the usual notion of model, as regards its "size". As a result we run, once more, into the problems of the traditional view. In order to pick out the intended denotation, we are told to look not at a reasonable segment of universe, but at the whole universe itself, although such a description as 'the barking dog' has no denotation with respect to a world where there are countless barking dogs. Certainly nothing prevents here the advocates of this theory of contexts from introducing "partial" worlds (or, more exactly, "pieces" of worlds) to get round this situation. But, in this case, a new theoretical entity is needed besides contexts (and worlds): and this proves that the notion of context elaborated by that theory is of no use when the noun phrase at issue is not an indexical or a "pure" demonstrative.

Is there any reason to keep these problems separate, to look for a notion of context that accounts for indexicals and related expressions, and another one that accounts for quantified expressions, definite and indefinite descriptions, and so on? After all, it is a truism to say that in both cases we have to do with the *contextual* nature of language.

Somewhere, in the previous discussion, we remarked that contexts, as they are conceived of in the logic of demonstratives, are *ad hoc* devices which provide the parameters needed by indexicals. They are added *from outside*, so to speak, to usual first-order models (ignoring here the possible-world machinery imposed by other intensional notions). But if we scan a model, conceived of as a complete state of affairs, we find out that all the necessary information is already present in it. The parameters we need in order to determine the referents of such indexical expressions as 'here', 'I', 'you', are given by the fact that, *somewhere* in the model, or, more exactly, in a model belonging to a given collection of models,

someone has the property of speaking to *someone else*. And this fact, because of the completeness of the model, must be *part* of the model itself; we have just to isolate it from the rest. If we conceive of contexts as model-theoretic constructs that we get by looking at the *internal* structure of the model, the ad hoc character we have spoken about disappears. As a result, a single notion of context is required to account for indexicals *and* quantified expressions, for in both cases we must refer to a local specification of the properties and relations characterized by the model. It is this specification which determines who is the speaker denoted by the expression 'I', or what are the dogs we must take into account when we utter or interpret sentences with such noun phrases as 'all the dogs' or 'the barking dog'.

*Dipartimento di Filosofia
Università degli Studi di Milano
Via Festa del Perdono 7
20122 Milano (Italy)*