

## Generics and Concepts

Greg Carlson  
University of Rochester

Draft version 8-06

### 0. Introduction.

In the experimentally oriented literature on concepts, one often runs across discussions which include language to the following effect:

The concept 'dog' is characterized (in part) by 'has four legs'  
'Eats meat' is a feature of LION  
'Flies' is a part of the prototype of BIRD  
....

The language employed in such descriptions fairly regularly consists of one part (the concept) that is expressed as a noun or a short phrase based on a noun, and the other part is a property, expressed as a predicate. (There is a relation expressed between them, the nature of which may vary and which need not concern us immediately.) The point is, when these components are combined, and dressed up a little to make them sound like grammatical English, the resulting sentences sound something like these:

1.     a. Dogs have four legs  
       b. Lions eat meat.  
       c. A bird flies.

Sentences like these express generalizations: they are not about particular events or characteristics, which might be exemplified by sentences such as the following:

2.     a. Paul's dog is out in the yard  
       b. That lion is eating some meat I just gave it.  
       c. A bird flew by.

These are not generic sentences. What is striking is that the conceptual literature tends not to use sentences like these in the discussions (at whatever the level of formality intended in the discussion). Rather, the language employed to discuss concepts and their structures bears a striking resemblance to the language examined in the literature on generics and habituals (e.g. Krifka et al, 1995). That is, taking a very broad and hazy perspective, there appears to be a striking confluence of interests between the study of concepts in the psychological and cognitive science literature, and the study of certain types of sentences in the formal semantics literature. Despite these initial confluences, the relationship between the meanings of generic sentences, and the study of conceptual structures remains unclear. While I cannot presently express any systematic ideas about what that relationship might be, this paper explores this relationship a little further, gives

some reason to think that further exploration of the relationship is warranted, and makes some tenuous suggestions about ways this exploration might be pursued.

The term “meaning” even as applied to natural language is notoriously difficult to define. However, I assume the normal machinery of what is currently known as “formal semantics” or “model theoretic semantics” in all its currently practiced varieties (which include dynamic approaches (Kamp and Reyle, 1993, Groenendijk and Stokhof (1991), Chierchia (1995)), situation semantics (Kratzer, 1989, Barwise and Perry, 1983), and the various offshoots of Montague grammar (Partee, 1974; Dowty, Wall and Peters, 1981).) This body of literature focusing on natural language semantics, for all its variety, is theoretically fairly coherent and well-understood, sharing fundamental assumptions and a kind of commitment to the agenda of truth-conditional semantics. It embodies a relatively clear view of what “meaning” (or at least a major component of it) might be. The main problem has to do with understanding concepts, at its best an only slightly clearer notion than “meaning”.

## 1. “Concepts”

The term ‘concept’ itself is also notoriously difficult to pin down; there is no single generally accepted intended sense to guide everyone in its application. But we can pin them down just a little better. At the most general level, concepts are spoken of in two major ways, which I will call ‘objective’ and ‘subjective’. The former are usually intended in philosophically-oriented discussions of concepts (though not always), where it is truly, without abstraction, possible to talk about two different individuals apprehending, grasping, coming to understand (etc.) the same concept. It is ‘out there’, independent of individual psyches; it is what is represented when we speak of a mental representation *of* a concept. In the formal semantics literature, any identification of Fregean senses with concepts has this character.

On the other hand, the discussions of concepts in the psychological and experimental literature tend to assume that what is being studied directly in observing behavior are the concepts themselves. In this sense, two individuals cannot share the same concepts any more than they can share the same sense impressions or experiences. While we *can* talk of individuals sharing concepts (sense impressions, experiences), this is talk in abstraction; all we can really do is determine whether individuals A and B have similar concepts, and that’s pretty much it—they can never be truly identical. (And then, of course, one can easily take the representational view, where we call the objective things ‘out there’ ‘concepts’, and then apply that selfsame term to the representations of those concepts themselves. This is where things tend to become a little confusing, for it tends to mix up what the exact object of study really is.)

For present purposes, the term ‘concept’ will be intended to refer to subjective properties of individual psyches, as that seems the most dominant assumption within the experimental literature I wish to make some connection with here. This choice in no way endorses any view that ‘concept’ should be confined to this sense alone, nor that there is

anything suspect about discussion of objective concepts. It is a matter of a perspective I wish to focus on here.

But in making this practical terminological choice, one gains only a little in terms of the clarity with which the term is defined. One way to illustrate the range of choice remaining is to operationalize the notion: that is, to ask “What do you want concepts to do for you?” “What do you want these to give some account of?” Depending upon one’s answer, one gets very different ideas of what a psychological concept might be. At the most basic level, it appears that everyone wants concepts to give an account of categorization. But, does one also wish for the structure of concepts to account for vagueness of category boundaries as well? Here, for instance, are some other kinds of things that Ray Jackendoff (1989) explicitly asks psychological concepts to give some account of:

- our ability to understand language
- the calculation of consistency with other sources of knowledge
- the performance of inferences
- the formulation of responses
- translating structures from one language (or domain) to another

Prasada and Dillingham (2005) ask of concepts that they enter into explanations of the character of things in certain ways. Wittgenstein (1952, para. 570) suggests that: “Concepts lead us to make investigations; are expressions of our interests, and direct our interests”, that is concepts direct our activities in a very global sense. Depending on how one answers these questions about what to expect from a theory of in terms of human (and animal) behavior, one gets very different views of how ‘rich’ a theory of concepts needs to be. I am going to leave this issue fairly open here, subject to one reservation: that concepts should give a full account of our ability to understand language.

## **2. A brief semantics for generic sentences**

Generic sentences take a wide variety of forms. Some of the most common include:

3. Cats have fur.
4. Mice squeak
5. Policemen carry nightsticks
6. The digital computer is found in many homes
7. A crayon melts all too easily in the sunlight.

In each instance, the subject of the sentence is an expression which does not make direct reference to any particular individual of that class. Thus, ‘cats’, ‘the digital computer’, and ‘a crayon’ in these examples are understood as applying to or picking out no one particular individual (though they may easily be so understood, in other contexts).

Broadly speaking, there are two views on what the semantics of such sentences might be. The first, is the quantificational approach. This assumes that as quantifier that has

variable-binding properties sets conditions on the number or proportion of individuals that may satisfy the formula, in order to give an account of truth-conditions. Thus, (3) might be rendered as (3'):

3'.  $\mathbf{Q}x [ \text{cat}(x) \rightarrow \text{have-fur}(x) ]$

This is imprecise in any number of ways, as it does not introduce intentionality and/or defeasibility, and I evade the precise nature of the connective (here rendered as ' $\rightarrow$ '); see Krifka et al (1995) for more details. But the essential point is that this analysis treats the subject as a predicate and not as a referring expression. Here, Q might be thought of as being an adverbial like 'normally' or 'usually' (that is, 'cats normally or usually have fur' seems an approximate paraphrase). The truth-conditions of such a formula depend crucially on substituting individuals as values of x (and, intuitively, only those individuals that are cats would have any bearing on the truth or falsity of the whole formula).

Another view is the one I've been espousing for some time. On this view, the subject noun phrase is not treated as a predicate, but rather as a fully referential expression which denotes a kind of thing. Thus, the noun phrase "cats" will refer to the kind cats (represented here as c), and predication will be attributed without mediating quantification. Thus, (3) might be represented by (3''):

3''. Have-fur (c)

The cost of approaching the semantics of generics making use of this strategy is that one needs to give some account as to why any predicate that can be meaningfully applied to a regular individual can also be applied meaningfully to a kind. This requires that there be some type-raising or type-shifting operators, whereas this correspondence follows naturally from the quantificational approach..

8. *Fluffy* has fur. *Cats* have fur.
9. *Mickey* squeaks. *Mice* squeak.

However, the quantificational approach does not give an adequate account of a number of other factors, which include the aspectual category of the resulting sentence, interactions with negation, and sources of intensionality. See Wilkinson (1991) for one extended discussion.

Rather, I'd like to turn to the question of how closely we can identify the meanings of these generic noun phrases (whether denotational, or predicational) with the corresponding concepts. Certainly if this were a reasonable hypothesis to pursue, then it would give us an excellent start on understanding why there seems such an affinity for phrasing property relations in the conceptual literature making use of generic sentences. One preliminary observation about the type of language used to describe concepts and affiliated properties is that they all tend to be rather brief, whereas brevity is no more a feature of generic sentences in any language than any other types of sentences. So not

only do we talk about flying geese and meowing cats and leaf-bearing trees, but we also talk about:

10. *Unpainted kitchen appliances that are just beginning to rust* need to be replaced within a period of two to three months by a qualified kitchen professional in order to prevent any possibility of bacterial contamination.

11. *Friendly but slightly confused medical professionals without appropriate training who nonetheless have medical degrees from top-ranked teaching hospitals* have much to contribute to society beyond their incomplete medical expertise.

The principled point is that linguistic expressions can be nearly unbounded in complexity—limited by the syntax of the language—with the consequence that, if these are mapped directly onto certain types of mental states that we’re calling concepts, then we stand in need of a device (let us call it ‘conceptual combination’) which can produce the arbitrarily large number of corresponding brain states in order to provide these phrases with appropriate denotations. While many people find this a perfectly reasonable and even desirable consequence (e.g. Jackendoff 1990), the philosophical underpinnings of the truth-conditional semantic enterprise makes such an outcome highly undesirable (and, in fact, wrong). The very notion of truth is itself an indexical one, and it is not wholly dependent on aspects of human mental representations. For instance, to take an example from generic sentences, imagine that the vast majority of people have a belief that snakes are slimy. One can even imagine (contrafactually) that everyone believes this of snakes. But what determines whether they are in fact slimy is not how people categorize them or what people might believe about them, but rather the nature of snakes themselves (which are in fact not slimy). Moving from truth to reference, the extensive literature on proper names stemming from Kripke (1972) and the literature on natural kind terms stemming from Putnam (1975) expressly show the inadequacy of any theory of meaning for such things that depends upon people’s mental states or abilities. The fundamentals of the theory of quantification and of modality, likewise develop shortcomings if thought of as mentally-interpreted constructs. Partee (1980) explicitly discusses whether linguistic meaning in general can be represented adequately by purely psychological notions. She shows, to put it briefly, that the space of meanings is simply too “big” to be represented by any brain or mind states (assuming that the mind/brain is finite in resources). While she couches her demonstration in terms of individuals, times, and possible worlds, it bears some emphasis that her conclusions do not depend on which particular parameters of evaluation one chooses, but will generalize to all sets of parameters that individuate meanings at least as finely as these. Further, it is well-known that, if anything, these parameters divide up meaning too coarsely, and not too finely. Thus, her demonstration possibly understates rather than overstates the volume of the corresponding mathematical space necessary to define linguistic meanings.

Just because we cannot identify concepts with linguistic meanings does not mean that concepts must play no role whatsoever. In fact, I believe that the hazy confluences noted above between talk of concepts and generic sentences is not coincidental, and that the

nature of concepts can play a serious role in determining how some generic sentences are understood. Maintaining an objectivist (denotative) semantics accounting for truth-conditions (and maybe more), the question becomes one of what corresponds in the mind to the meanings of phrases denoting kinds and phrases denoting their features. It is most plausible to think that kinds correspond to something in the mind, and these seem to be typically labeled concepts. In the next section I outline some phenomena associated with the semantics of generic sentences that might well be understood as connected to or influenced by conceptual structures.

### 3. Generic sentences and concepts

3.1. k- and t- properties. Prasada and Dillingham (2005) explore in some depth the connection between kinds and their properties. Their research strongly suggests that people represent a principled connection between the type of thing something is and some of its properties. They call these “k-properties”. However, other properties the kind may have are not judged to call into account a principled sort of connection, but depend upon simple “statistical” cooccurrences alone (or primarily). They label these latter properties “t-properties”. For example, our judgment that dogs have four legs enters into certain types of explanation (e.g. this has four legs *because* it is a dog), and into normative judgments (if this is a dog, it *should* have four legs). On the other hand, our judgment that dogs are brown is an instance of a t-property, because this judgment seems to be based on sheer statistical regularity alone, and does not enter into the explanatory or normative accounts. Here are a few other examples drawn from their article:

<u>k-properties</u> (“principled”)	<u>t-properties</u> (“statistical”)
airplanes have wings	barns are red
trains travel on tracks	cars have radios
cheetahs run fast	dogs wear collars
doctors diagnose ailments	golfers wear plaid pants

Now, it is not that properties ascribed to kinds are k-properties or t-properties once and for good. Rather, it depends upon what kind of thing they are being predicated of. For instance, redness said of barns appears to be a t-property, whereas redness said of blood would appear to be a k-property. Thus, the distinction between these types of properties is the consequence of predicating one thing of another, i.e. of the meaning of a wholes sentence. The point of the Prasada and Dillingham article is to substantiate this distinction experimentally.

3.2 Greenberg (2003). Yael Greenberg has done some work on the formal semantics of generic sentences which appears to lead us to something that looks like it could well be the same distinction. She is interested in giving an account of an interesting and systematic difference in interpretation between English indefinite singulars as subjects of generic sentences, and bare plurals. For the most part, the matter of singular vs. plural to

the side, the two are pretty often substitutable for one another in a generic sentence, and they will seem to mean almost exactly the same thing, as the following illustrate:

12. A dog has four legs / A grizzly bear hibernates during the winter.
13. Dogs have four legs / Grizzly bears hibernate during the winter.

However, as has been known at least since the observations of Lawler (1973), there are many instances where the two differ, and the indefinite singular version simply seems strange or possibly ungrammatical:

14. a. Madrigals are popular / Rooms are square / ?Men are blond/ Uncles like marshmallows.  
b. # A madrigal is popular / # A room is square / #A man is blond / #An uncle likes marshmallows

Lawler originally characterized the relationship required by the indefinite singular as one where the property is “essential”. But Greenberg revises this understanding since essence is too strong, to one of requiring a “principled connection”, as the properties. The contribution of her work is to characterize the notion of “principled connection” in a model-theoretic framework.

Part of Greenberg’s case makes use of observations about “nonce” categories. These are categories often expressed by fairly complex noun phrases carve out little-used or unusual categories. What, for instance, follows from a person being a French writer born in 1954, apart of whatever follows from being a French writer, and being born in 1954? Yet, for most such ad hoc categories, one could find fairly consistent properties associated with that group that would be judged to be simply accidental. So, suppose that, by sheer happenstance, French writers born that year by and large, and in contrast to those born in 1953, 1955, etc., write quite technical papers. If we use the bare plural construction, even as unmotivated a connection as this sounds quite fine, whereas the indefinite singular examples (when interpreted generically) are seriously degraded:

15. a. French linguists born in 1954 write very technical papers  
b. #A French linguist born in 1954 writes very technical papers

Or, to exhibit one example from Greenberg, consider the following contrast:

16. a. Famous carpenters from Amherst give all their sons names ending with ‘t’.  
b. # A famous carpenter from Amherst gives all his sons names ending with ‘t’

Now, this does not mean all such ad hoc categories lead to strangeness in the indefinite singular. If one perceives a principled connection between the kind and its predicated property, it becomes fine. Consider the phrase “bananas that have been sat on by a rhinoceros”, a category with very low occurrence frequency (a google search reveals no matches whatsoever). If we say:

17. Bananas that have been sat on by a rhinoceros are flat.

one can easily understand the principled connection: the sitting causes the flatness of the otherwise-shaped banana. Thus, we predict that the indefinite singular expression of the corresponding sentence will be acceptable. And, it is:

18. A banana that has been sat on by a rhinoceros is flat.

So it is not a matter of whether the category is ad hoc, but rather whether a principled connection exists.

If we return to the Prasada and Dillingham materials and substitute indefinite singulars for the subjects of their proposed k- vs. t- properties, we intuitively seem to find that the t-property examples are judged to be less acceptable. Consider the following sampling drawn from their categories of natural, artifactual, and social kinds:

k-properties

A car has four wheels  
A train travels on tracks  
A cheetah runs fast  
A lemon is sour  
An artist is creative  
A gymnast is flexible

t-properties

A barn is red.  
A shower cap is transparent  
A pigeon sits on statues  
A rock is jagged  
A biker has tattoos  
A Hindu lives in India

It would appear that Prasada and Dillingham, and Greenberg are talking about much the same thing.

3.3. Property stability. Lyn Frazier has pointed out to me that t-properties appear to be less “stable” than judged k-properties over time. That is, while most generic properties are amenable to change, the k-properties are, on the whole, more persistent than the t-properties, as a whole. If for instance, bears ceased hibernating, they would still be bears, but one would have the distinct impression that there must have been a big change (whether in bears, or the environment) to account for this. On the other hand, if one noted that bears no longer performed in circuses, no corresponding big change would be assumed because one was not assuming this to be a highly stable property (being a t-



property).

Now, it is also a fact that generic sentences, while commonly expressed in the simple present tense in English, can be also expressed in the past (or the future). However, if one uses a generic in the past that could have been easily used in the present as well, consider the intuitive effect:

19. a. Dogs ate meat.
- b. Teachers assigned homework
- c. Airplanes had wings.

The examples in (19) would also be true in the present as well.

20. a. Dinosaurs laid eggs.
- b. Romans wore togas.
- c. Pony Express riders were expert horsemen.

One senses some difference between the examples of (19) and (20). In examining (19) in contrast to (20), we see two possible effects. One, we can find the examples of (19) a little odd or unacceptable. We could also accord them a non-generic interpretation as a means of trying to avoid this oddness. Understanding them this way would yield interpretations similar to saying “SOME dogs ate meat (on a certain occasion)”, or “SOME airplanes had wings.”

Frazier and Clifton have run a preliminary study which applies this observation to the study of k- and t- properties. A comparison was made between the k-properties list of Prasada and Dillingham, and their t-properties with sentences that have been put into the past tense. Since t-properties are more amenable to change, the reasoning was that it would be simpler to imagine a circumstance where they are no longer true but once were. However, k-properties, because of their relative stability, would make it more difficult to naturally interpret them generically in the past tense, making a non-generic interpretation more likely. In the study, this is what was found. The full results await further work. Thus, it appears that the interpretations of certain types of sentences of natural language can be understood as affected by a distinction arising from work on the nature of concepts.

3.4 “Suitability”. Ariel Cohen (2005) uses the term “suitability” to refer to the use of existential uses of bare plurals vs. existentially quantified NP’s in the attributed property of generic sentences. Cohen notes that in certain instances substitution of a quantified existential for a bare plural will yield a slight degrading of judgement, whereas in other instances there is no similar degradation. For example, (21a) seems more acceptable than the quantified instance of (21b) (substitution of ‘some wheels’ for ‘wheels’):

21. a. This tractor has wheels.
  - b. (?) This tractor has some wheels
- vs.

Whereas there seems no such similar contrast between the examples found in (22):

22. a. These ambulances have dents in them.
- b. These ambulances have some dents in them.

The judgments are subtle but I believe stable. Cohen explains the contrast this way: suppose the tractor in question in (21) had only two wheels. This is not a “suitable” arrangement of wheels for the proper functioning of a tractor. Under these circumstances, the speaker would be more apt to phrase the sentence as (21b) than as (21a). Put roughly, how the wheels on a tractor are arranged, and how many there are, can affect the tractor’s ability to function fully in the way intended. On the other hand, there is no such similar contrast in (22), since ambulances with dents in them, however arranged, however many there are, do not contribute to the functioning of the ambulance in the way intended (though of course huge dents might inhibit their functioning).

Lyn Frazier again has noted to me that if one divides up the predicates according to whether they are t-properties or k-properties, the “suitability” implicature is or is not triggered. The properties that appear to have the characteristics of k-properties accorded them by Prasada and Dillingham are the ones that appear to trigger the implicature, whereas the t-properties do not. Frazier and Clifton again are working on quantitative means of checking on this, with encouraging results thus far. The hypothesis to be evaluated is whether t-properties yield fewer preferences for bare plurals over “some” NP’s than k-properties, which ought to show a distinct preference for the bare plurals.

3.5 A long-standing puzzle. In Carlson (1977) I puzzle about the following example. If you look at cars, you will clearly see that:

23. Cars have tires.

This is true in most people’s experience for 99+% of all cars observed. It is also true to that level of consistency that the tires observed were all black (white-wall tires are black, too, by the way, but they are “out” these days and so we’ll ignore them). So while it seems to be true that cars do indeed have tires, and it is also true to nearly 100% consistency that the tires are black, there is something nonetheless strange about the following sentence:

24. (?) Cars have black tires.

For those who get the intuition (and some don’t), the sense is that the blackness of the tires is somehow irrelevant. But hold on:

25. Cars have air-filled tires

seems quite fine by comparison, yet why might “black” seem irrelevant when “air-filled” is not? After all, exactly as many tires on cars are air-filled as are black (or, in the event

of flat tires, possibly even more). These distinctions cannot be accounted for, obviously, by some appeal to number, frequency, probability, or proportion. However, if we examine these examples from the standpoint of suitability implicatures, and k- and t-properties, we do see a distinction. The blackness of tires does not further the utility of the vehicle, it does not make it function better in the way intended than if it had green tires; on the other hand, the tires being air-filled does enhance the suitability of the car to function appropriately: let the air out and you soon see why. Further, if one thinks about whether “black” as applied to tires is a k- or t- property, one is most likely to conclude that, despite its consistency as a property of tires, it is a t-property; on the other hand, “air-filled” as a property of tires would in most peoples’ judgment be likely to come out as a k-property.

#### 4. Generic reference and concepts, perhaps

The previous section discussed some linguistic phenomena that seemed to correlate, in some cases rather convincingly, with the k- and t- property distinction discussed and motivated in Prasada and Dillingham’s work. In this section I wish to discuss a couple of potential applications of conceptual structures for noun phrase reference. Recall that the general hypothesis is that generic noun phrases are taken to refer to kinds of things, and that these phrases can be of unlimited complexity, and hence, one can refer to an unlimited number of kinds of things. This appears to be so for generics that make use of bare plural expressions in English.

However, as has been pointed out for some time, the alternative definite generic expressions of English are limited. In Krifka et al (1995), this limitation is expressed as a requirement that English definite generics refer only to “well-established kinds”. That is, there is some presupposed familiarity, or something of that character. This suggestion is not more precisely defined. As an illustration, consider the following contrasts:

26.
  - a. *The bottle* has a narrow neck
  - b. *The Coke bottle* has a narrow neck
  - c. ??*The green bottle* has a narrow neck
  - d. OK *Green bottles/ Coke bottles* have narrow necks
  
27.
  - a. *The Indian elephant* has smallish ears and is easily trained.
  - b. ??*The friendly elephant* is easily trained.

The examples in (26) are originally due to Barbara Partee. They illustrate a distinction not only between “the bottle” but also between “Coke bottles” and “green bottles”. Coke bottles are a familiar type of bottle (or presumed to be) whereas green bottles are not so regarded.

The basic problem here is, however, that at least at the time the examples were first introduced, Coke bottles were (light) green in color, and thus everyone’s experience with green bottles would have been a subset of experiences with green bottles. Thus, green bottles ought to have been more, and not less, familiar. Note that this distinction

disappears when abre plurals are used (26d). Or, suppose Indian elephants are all friendly elephants (and some others were, too); in this case one's experience with Indian elephants would have been no more frequent than one's experience with friendly elephants, yet even on this understanding the distinction remains. Though I've been phrasing things here in terms of familiarity (of experience), well-establishedness may hold something for us instead.

Let us return to the question left in abeyance above as to how concepts (as construed here) might participate in the semantics of language. We have already set aside the hypothesis that they serve as the denotata of linguistic expressions in general, but this does not preclude the possibility that they could serve as the denotations of parts of language, or be parts of meanings (just not the whole story). Greg Murphy (2005) explicitly argues what is implicit in most discussions of concepts, that concepts are the denotations (the meanings of) the lexical items of a language. Certainly this is plausible; discussion of lexical semantics is at the root of most linguistic work that invokes concepts as explanatory elements, and in the psychological work on concepts, lexical semantics is by far the most common vehicle used to express a concept (as when one writes, "the concept GREEN"). More specifically, the kinds of lexical items used tend to be almost exclusively drawn from the nouns, verbs, adjectives, and occasionally the prepositions of English and similar languages.

Let us assume, only for the moment, that concepts as meanings are limited to the lexical items of the language. On this supposition, 'green', 'dog' and 'jump' will denote concepts (and hence there must be the concepts GREEN, DOG, and JUMP), but expressions like "greener than your mother-in-law's kitchen wall", "hungry or thirsty dog" and "jump right over the fence" do not denote concepts. This leaves open the possibility that there are such concepts as HUNGRY-OR-THIRSTY-DOG, but there need not be in order for the corresponding phrase to have a meaning, unlike the monomorphemic lexical items. The consequence is, then, that complex phrases would not take concepts as their meanings, if concepts as meanings are limited only to meanings of lexical items. Thus, we would not (contrary to most common views of concepts) expect 'thirsty dog' to designate a concept, if each of its component parts names one.

Put more to the point: the hypothesis under consideration is that concepts do form meanings for certain types of expressions of natural language, but not for others. They might fail to serve as denotations for two reasons: one would be that the semantics of the language does not map certain category meanings, or it might be that there are no such corresponding concepts. Above I briefly suggested that these be limited to only the lexical items of the language, but this might be at once too strong, and/or too weak. Let us consider the possibility that the constraint on the denotation of English definite singular generics is that the whole NP denotes a kind if and only if the common noun it is based on itself denotes a concept (or, rather, it is presupposed that there be a concept denoted by the common noun). Thus "the lion" will be an acceptable generic noun phrase because the "lion" part will itself denote the concept LION. But the generic NP "the friendly elephant" will not be acceptable (as seems correct) because there is no

concept denoted by “friendly elephant”. At this point, the hypothesis is almost entirely circular, of course, let us pursue it just a little further.

As the examples above indicate, the definite generics are not limited to monomorphemic expressions, as “the Coke bottle” and “the Indian elephant” are acceptable. The former is a compound noun, and thus would qualify as a lexical item of the language. However, “Indian elephant” has the structure of an adjective-noun combination, and it is clear that not all phrases of this structure form acceptable definite generics (as the unacceptability of “the green bottle” and “the friendly elephant” would indicate). And, it does not seem entirely clear that anything that has a compound noun structure necessarily forms an acceptable definite generic:

??The actress scandal  
??The alcohol syndrome  
?? The shift boss  
??The orange peel garbage bag

(cf: Actress scandals are more common than they used to be  
A typical alcohol syndrome involves tremors  
Meet Bob, he’s my shift boss  
You categorize your garbage well: is this your orange peel garbage bag?)

So, if it is not structure that determines whether something is an acceptable definite generic of English, then it must have something to do with the meanings expressed. In this, one has the common intuition that there is a perceived difference between “green bottle” and “Coke bottle”, or “Indian elephant” and “friendly elephant”. It goes something like this: an Indian elephant is one of a recognized variety of elephants, and their properties are not simply those of being an elephant, and being from India, but something more (such as disposition, size of ears, etc. etc.)—it’s a (sub)species. In this sense, “Indian elephant” differs from “friendly elephant” because a friendly elephant is no more than an elephant that is friendly, and that’s it. “Indian elephant” also differs from a phrase like “Indian tree” in that an Indian tree is simply something that is a tree and from India, and that’s it, and thus, the phrase “the Indian tree” ought not be a good definite generic, which seems correct.

It is reasonable to think that the constraint might be one on requiring that the denotations be “natural kinds”, but one must then accept artifacts (“the digital computer”, “the fountain pen”) and socially-created things (“the sonnet”, “the jump shot (in basketball)”) as natural kinds as well. I’m suggesting, for further thought, that perhaps the term ‘concept’ should be countenanced as the critical ingredient in place of “natural kind”. In any event, there are clearly further constraints that neither gives proper account of. It is for instance difficult to naturally apply these to many human categories (“??The lawyer is indispensable in our system of justice,” “??The violin player...”), and, as Vendler observed (1971), terms that are “too general” do not seem to sit well as definite generics:

*The parabola* is easily plotted.

??*The curve* is easily plotted.

## 5. One last phenomenon

The last matter I wish to consider does not concern generic sentences specifically, but it is related to genericity by virtue of the involvement of determinerless nominals, one of the main ways of expressing kind-reference across languages. The phenomenon in question is that of noun-incorporation, in which (typically) a direct object determinerless noun is morphologically joined to a verb, becoming a part of it (e.g. Baker (1988), and for two overview articles, see Gerdtz (1998) and van Geenhoven (2001), and references therein). One example of this is given below.

28. a. Enan            qaa-t            qErir-ninet  
      3sg.Erg reindeer.Abs.pl seek-3sg.s/3.pl.o

      b. Etlon            qaa-rer-g'e  
          3sg.Abs reindeer-seek-3sg.s

"He is looking for reindeer" . (Chukchi, Spencer 1995)

In the (26a) the determinerless noun 'qaat' is in "normal" direct object position, whereas in (26b) the nominal, stripped of its plural marking, is morphologically made into a part of the verb. Both sentences mean about the same thing, though with some subtle differences (beyond the issue of plurality) we'll come back to in a moment.

In the substantial literature on incorporation, it is often noted that the process is not fully general (though at times it can be). In particular, it is often noted that the noun-verb combination must describe what is variously characterized as a "familiar", "habitual" or "generic" activity, or that it describes an activity that is "nameworthy"; it is occasionally described as designating a "unitary concept". These are, of course, notional characterizations which, it appears, point to very much the same phenomenon. For example, in Chukchi, if one uses the incorporating form of "reindeer + kill", it does not mean simply to kill a reindeer (as the unincorporated form does), but to kill a reindeer as a ritual part in the preparation of a meal (Dunn, 1999). That is, in the Chukchi culture, it is a "nameworthy" or "habitual" sort of activity.

Nameworthiness is also a means of dealing with lexical gaps. I will illustrate the point using Norwegian, as studied in great detail in Borthen (2003). While Norwegian does not have incorporation proper, it does allow determinerless singular count nouns in object positions (English is much more restricted in this). The semantics of the determinerless nominals, though morphologically not incorporated, appears to be identical to the semantics of incorporated nominals in other languages (e.g. van Geenhoven, 1998; Chung and Ladusaw, 2004; Farkas and de Swart, 2004). They are "semantically incorporated" forms.

Borthen notes that not all combinations of predicates and bare nominals are acceptable. While, for instance, the examples in (27) are natural and acceptable, those in (28) are not:

29. a. Han anbefalte rullestol  
He recommended wheelchair  
'He recommended a wheelchair'
- b. Jeg kan lese bok, jeg  
I can read book I  
'As for me, I can read a book'
- c. Han eier bil  
He owns car  
'He owns a car'
30. a. \*Per slo jente  
Peter hit girl
- b. \*Hun vasket sykkel ren  
She washed bicycle clean
- c. \*Ulven drepte okse  
Wolf-the killed bull

Borthen goes to great lengths to rule out a wide variety of possible accounts, and in the end provides an explanation for this particular class of constructions, which she calls the 'Conventional situation type' construction, along the following lines: The predicate and the noun together must denote a 'conventional situation type', and the construction can be reasonably seen as a multi-word lexical entry. A conventional situation type is "a property, state, or activity that occurs frequently or standardly in a given contextual frame, and has particular importance or relevance in this frame as a recurring property." (p. 153).

I believe that this same sort of characterization is, in other rubrics, more broadly noted in many studies of the nature of noun incorporation in many languages. Note that Borthen's sense that these could be regarded as multi-word lexical entries is made more plausible in incorporating languages, where the noun/verb combination is in fact of the category of a lexical item, namely, a verb.

We have been considering the possibility that the lexical items of a language, at least, may denote or be associated in some grammatically-significant way, with concepts. If we think of concepts as primarily classifications of sense-experiences we might have, and that words of a language may be used to express or encode many of those concepts, we might take Borthen's characterizations and think in terms of the role a concept might play: consider the possibility that the requirement for a V/N combination to be

acceptable, that it presupposes that there is a corresponding concept. So, for instance, car-owning and book-reading would be such potential concepts, but bull-killing and bicycle-washing, would not be. With such a provision, we could provide an account of such lexical gaps as those noted above, and in many other languages for incorporation and incorporation-like constructions. Of course, this is merely the rankest of speculation without there being some independent means of determining if one, or the other, is, or is not a concept, and at the moment I must leave this up to the ingenuity of psychologists. But the gross outlines of a view of concepts that emerges from suggestions like this and those above is one where concepts are not, in general, available as the denotations of all types of natural language expressions. Instead they are “spotty”, can be idiosyncratic, and to a great extent tied up with the nature and structure of a language and a culture.

In closing this section, we do need to recognize that one major issue has been systematically skirted in the above discussions. We took the point of view that concepts were primarily psychological objects, and hence fundamentally unshared. However, the discussions above have relied upon there being a means of sharing them, or at least aligning them, within a community of individuals speaking the same language. While I believe there are ways of accomplishing this, it is an agenda for another venue.

## **6. Conclusion.**

Concepts, at least as studied by psychologists, have not found a natural place in the study of natural language semantics in the formal semantics tradition as practiced in the past thirty or so years. This does not point to any shortcomings in semantic theory, I emphasize, for there are excellent reasons intellectual as well as practical as to why this should be. However, it does leave us with something of a disconnect between what is going on in the mind, and I take concepts to be major features of it, and the way we talk about things.

This has been an attempt to make suggestions about how we might close that gap some. We do not examine truth-conditional phenomena for evidence of concepts in semantics. Rather, we consider how truth-conditional phenomena are expressed in a language, and the restrictions and constraints on that, to uncover our clues. We also require a clearer articulation of what “concepts” in fact are. The extent to which work on natural language semantics can be put together with work on culture and work on the mind, is the extent to which we begin to have a chance of understanding what might be responsible for the phenomena discussed in this paper.

## **References**

- Baker, M. (1988). *Incorporation*. Chicago: University of Chicago Press.
- Barwise, J. and J. Perry (1983). *Situations and Attitudes*. Cambridge, MA: MIT Press.
- Borthen, K. (2003). Norwegian Bare Singulars. Doctoral dissertation, NTNU,



Trondheim.

Carlson, G. (1977). *Reference to Kinds in English*. University of Massachusetts/Amherst doctoral dissertation.

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

Chierchia, G. (1995). *Dynamics of Meaning*. Chicago: University of Chicago Press.

Chung, S. and W. Ladusaw (2004) *Restriction and Saturation*. Cambridge, MA: MIT Press.

Cohen, A. (2005) More than bare existence: An implicature of existential bare plurals. *Journal of Semantics* 23: 389-400.

Dowty, D., R. Wall and S. Peters (1981). *An Introduction to Montague Semantics*. Dordrecht: D. Reidel.

Dunn, M. (1999). A Grammar of Chukchi. Ph. D. Dissertation, ANU.

Farkas, D. and H. de Swart (2004). *The semantics of incorporation: from syntax to discourse transparency*. Stanford, CA: CSLI.

Dayal, V. (1999) Bare NP's, reference to kinds, and incorporation. *Proceedings of Salt IX*

van Geenhoven, V. (1998). *Semantic incorporation and indefinite descriptions: semantic and syntactic aspects of noun incorporation in West Greenlandic*. Stanford: CSLI.

van Geenhoven, V. (2001). Noun incorporation. *GLOT International* 5(8), 261-71.

Gerds, D. (1998). Incorporation. In A. Spencer and A. Zwicky (eds.) *The Handbook of Morphology*. Blackwell. 84-100.

Greenberg, Y. (2003) *Manifestations of Genericity*. New York: Routledge

Groenendijk, J. and M. Stokhof (1991). Dynamic predicate logic. *Linguistics and Philosophy* 14:39-100.

Jackendoff, R. (1989) What is a concept, that a person may grasp it? *Mind and Language* 4, 68-102.

Jackendoff, R. (1990). *Semantic Structures*. Cambridge, MA: MIT Press

- Kamp, H. and E. Reyle (1993). *From Discourse to Logic*. Dordrecht: Kluwer.
- Kratzer, A. (1991). An investigation of the lumps of thought. *Linguistics and Philosophy* 12, 607-53.
- Krifka, M. , F. J. Pelletier, G. Carlson, A. ter Meulen, G. Link, and G. Chierchia (1995). Genericity: An introduction. In Carlson and Pelletier, eds., 1-124.
- Kripke, S. (1972). Naming and necessity. In D. Davidson and G. Harman, eds., *Semantics for Natural Language*. Dordrecht: D. Reidel.
- Lawler, J. (1973). *Studies in English Generics*. University of Michigan Papers in Linguistics 1:1, Ann Arbor: UM Press.
- McNally, L. (1998). Existential sentences without existential quantification. *Linguistics and Philosophy* 21, 353-392.
- Michelson, K. and M. Doxtator (2002). *Oneida-English English-Oneida Dictionary*. Toronto; University of Toronto Press.
- Millikan, R. (2005). *Language: A Biological Model*. Oxford: OUP.
- Mithun, M. (1984). The evolution of noun incorporation. *Language* 62, 847-94.
- Murphy, G. (2002) *The Big Book of Concepts*. Cambridge, MA: MIT Press
- Partee, B. (1974). *Montague Grammar*. New York: Academic Press.
- Partee, B. (1980). Semantics -- mathematics or psychology?, in Bäuerle, Egli, and von Stechow, eds., *Semantics from Different Points of View*, Springer-Verlag, Berlin 1-14.
- Prasada, S. and E. Dillingham (2005) Principled and statistical connections in common sense conception. *Cognition*.
- Putnam, H. (1975). The meaning of "meaning". In K. Gunderson , ed., *Language, Mind, and Knowledge*. Minneapolis: University of Minnesota Press, 131-93.
- Spencer, Andrew. (1995). Incorporation in Chukchi. *Language*. 71.439-89
- Vendler, Z. (1971). Singular terms. In D. Steinberg and L. Jakobovitz (eds.), *Semantics: an Interdisciplinary Reader*. Cambridge: CUP. 115-33.
- Wilkinson, K. (1991). Studies in the semantics of generic noun phrases. University of Massachusetts/Amherst Ph.D. Dissertation.
- Wittgenstein, L. (1952). *Philosophical Investigations*. New York: Macmillan.