Perspectives on kind subjects

Ellen Lau
September 18 2024

In this handout, I review some of the basic observations about 'kind' readings of bare plurals in English. For this meeting I may only make it through the undisputed kind subjects and so we can leave the rest as an appendix.

[For context: I want to adopt some version of a sortal approach to noun meaning for psycho/neuro theory, so I'm interested in observations in the semantics literature that seem to put pressure on the traditional set-theoretic treatment of nouns.]

0. Undisputed kind subjects

(1a) Dodos are extinct.
(1b) Computers were invented in the 20th century.

Bare plural subject sentences of English of the type in (1a-1b) seem not to be understood as predications on individual dodos or computers. The predicates EXTINCT and INVENTED seem instead to apply to something more abstract, of which the individual dodos and computers are instances. That more abstract thing, in these and similar examples, is often called a "kind", i.e. the dodo kind and the computer kind. This point is mostly undisputed in the linguistic semantic literature on kinds (Krifka et al. 1995).

Kind subjects as in (1a-1b) do not seem to mean (2a-2b), that is, universal quantification over the set of dodos or the set of computers[1].

(2a) #Every dodo is extinct.
(2b) #Every computer was invented in the 20th century.

While essentially plural predicates also resist universally quantified subjects to some extent (3b), (2a-2b) are worse and their unacceptability seems to have a different source.

(3a) Wolves clustered near the stream.
(3b) ?#Every wolf clustered near the stream.

Intuitively the reason is that, unlike the essentially plural predicate in (3a), the predicates in (1a-1b) are not combining with expressions that simply refer to a set of individuals, even if that set were understood as extending across all possible worlds[2].

Accounting for kind subjects thus requires at least two enrichments to a simple set theoretic semantics:

[1] Note that (2a-2d) are fine on a *taxonomic kind* reading of the noun, in which e.g. 'every' quantifies over dodo species or computer models. That is consistent with these predicates being restricted to kind subjects.
[2] A common intuition in the literature is that kind-referring expressions rather objectify the *essence* or *character* of that set.

- (i) Assume that some of the individuals in the model domain are abstract kinds, not instances or sets of instances

- (ii) Account for the (apparent) systematic lexical ambiguity in whether common nouns denote a kind or a set of instances. This can involve modifying the semantics assumed by the theory for *any* sentence containing a common noun.

Carlson's (1977) dissertation is a classic within modern natural language semantics, as he carefully enumerated the significant enrichments to the model and to the noun semantics that seemed to be needed to maintain a set theoretic approach in the face of these facts. However, today most introductory texts in semantics never mention kinds at all, nor the additions to the theory that are required to account for them (Heim & Kratzer; Larson & Segal; Coppock & Champollion). Perhaps this is a problem.

1. Enrichments required by undisputed kind subjects

As outlined above, undisputed kind subjects require at least two kinds of enrichment[3] to the theory: allowing kinds, and explaining the systematic generalization within the language that a given common noun can be used to refer to both a kind and the set of its instances.

1.1 Carlson (1977)

Carlson's (1977) dissertation adopted several strategies that have been influential since:

- To address the first issue, Carlson simply allows kinds to be a sort of individual in the domain. In more neutral terms, kinds and objects/instances are just two sorts of property-bearing elements in the domain, which can serve as the argument to predicate functions in exactly the same way. This doesn't require any other change to the logic, and maybe we don't need to be concerned about this more abstract notion of individual given that we already had plenty of other elements in the domain of individuals that are either immaterial, or cannot be identified just in terms of material (e.g. 'the party' or 'love').

- To address the second issue—the noun semantics—Carlson proposes introducing a realization relation R(x,k) to the domain, which relates kinds and the objects/instances that realize them. Then one can use R to state, in one way or another, the systematicity in the meaning contribution of the noun across kind-denoting sentences like (1a-1b) as well as instance-denoting sentences.

Most proposals since have assumed something akin to R but have varied in exactly how the two meanings get determined. Carlson (1977b) proposed that the noun itself was unambiguous, and denoted the kind individual only. He localized the ambiguity in the resulting sentence meanings in lexical variation across predicates, where some predicates introduced the R relation and took the kind individual as an argument of R roughly as indicated in (4a-4b).

[3] One might argue that the semantic theory doesn't need to commit to what aspects of the world are individuals and what aspects aren't, so that this is not really an enrichment. As a technical matter this could be true, but given the commonsense usage of 'individual' as well as the common mathematical one (e.g. there is no obvious analog to 'kind' within number theory), most theorists have felt that it is at least worthy of note.

(4a) Dodos are extinct: Extinct(d)
(4b) Dodos are sick: ∃x[R(x,d) & sick(d)]

In (4a), we have simple predication of the individual denoted by 'dodo', which is the kind individual. In (4b), we have a predicate that doesn't apply to kind individuals directly, but which introduces the R relation, and then applies to the individual instances in the domain that are related by R to the kind individual denoted by 'dodo'.

This approach was elegant because it didn't require ambiguity in any particular lexical item. Fully avoiding ambiguity though would mean that each predicate would *either* take instance individuals or kind individuals as arguments (either it's part of the predicate meaning to introduce the R relation or it isn't). However, if one is willing to grant that bare plurals are interpreted as kind subjects in cases beyond the undisputed ones in (1a-1b), then many predicates seem able to accommodate *both* kind and object meanings, and thus the elegance of this approach is diminished.

It is also worth noting that this approach requires us to alter the semantics we assume for many sentences not involving kinds, like 'John is sick', to include the introduction of the R relation by sick() indicated in (4b). Carlson took this to be a virtue, proposing that an instance like John could be conceived as encompassing many stages across time just like a kind like 'dodo' could be conceived as encompassing many dodo instances.

1.2 Chierchia (1998)

Chierchia's (1998) approach is a bit more complicated. Like Carlson (1977), he assumes that kinds are allowed as individuals in the domain. Instead of positing a relation between kind individuals and instance individuals in the domain, Chierchia assumes a correspondence between *properties* and kinds: 'To any natural property, like the property of being a dog, there corresponds a kind, viz. the dog-kind'.

Chierchia (1998) influentially suggested that in some languages, nouns refer to kinds (e.g. Chinese) and therefore could serve as arguments directly, and that in other languages, nouns refer to predicates (e.g. Romance) and therefore require a determiner to serve as an argument. Silent type-shift operators, motivated by the assumed property-kind correspondence in the domain, are assumed to convert back-and-forth between predicate and argument as needed for composition with the rest of the expression. Although Chierchia took the pattern of determiner use and mass/count/kind interpretations across languages as evidence in support of his view, subsequent work from a larger number of languages has suggested a much more complicated empirical picture (e.g. Schmitt & Munn 1999 on Brazilian Portuguese).

It is also interesting to note that Chierchia (1998) seems to take a more plurality-like view of kinds than Carlson. He writes,

> It seems natural to identify a kind in any given world (or situation) with the totality of its instances. Thus, the dog-kind in our world can be identified with the totality of dogs, the scattered entity that comprises all dogs, or the fusion of all dogs around. In our framework this entity is modeled by the set of dogs. This means that we can model kinds as individual concepts of a certain sort: functions from worlds (or situations) into pluralities, the sum of all instances of the kind. A kind may lack

instances in a world/situation (as is the case, presently, for the dodo), in which case the corresponding individual concept will be undefined.

Although this may be meant only as a loose speculation about *why* we have kind-denoting expressions, it is important to hold on to the point that predication in kind-subject sentences cannot amount to predicating a set of instances, because intuitively in 'Computers were invented in the 20th century', we are not predicating 'invented()' for instances of computers.

1.3 Taxonomic kinds

Later work suggested that more than one relation might need to be introduced into the domain in order to capture a third type of noun interpretation, the 'taxonomic kind' reading:

(5a) Only two dogs are native to North America: poodles and daschunds.
(5b) Three frogs were just added to the endangered species list.
(5c) Several appliances were invented here.

In (5a), 'dogs' doesn't seem meant to refer to instances of dogs, but it also doesn't seem to refer to the singular dog kind as in 'Dogs are extinct'. Here, 'dogs' seems to refer to a plurality of subkinds of dog—not two kinds, not two instances, but two subkinds. Capturing the systematicity of these readings on an approach like Carlson's seems to require now adding subkinds to the domain along with another relation over entities in the domain: a taxonomic subkind relation T(x,y) that relates a subkind y to the kind k. Then one can state that each noun can denote a kind k, or the set of subkinds y that bear the relation T to k, or the set of objects x that bear the relation R to k (Krifka et al., 1995).

1.4 Summary thoughts on undisputed kind subjects

Influential set theoretic accounts of undisputed kind subjects such as (1a-1b) make significant enrichments to the model (positing that relations between elements of the domain play a crucial role in the semantics), and require significant changes to the traditional semantic analyses of sentences containing common nouns that do *not* receive kind interpretations. Given that standard textbooks do not discuss these issues, and given that outside of this subliterature, analyses of sentences containing non-kind noun interpretations are not usually updated in the way required by these accounts, we may not yet appreciate the full ramifications of these enrichments for the larger theory.

2. Disputed kind subjects

Up to now we have focused on undisputed kind subjects only, because these are cases which seem to require enriching the standard theory. I will next briefly acknowledge the much larger range of bare plural subject sentences like (6a-6d) which are consistent with a kind subject analysis along the lines discussed above, but for which other analytical options are also available.

(6a) Tigers have stripes.
(6b) Birds fly.
(6c) Ducks lay eggs.
(6d) Ticks carry Lyme Disease.

Unlike (1a-1b), the predicates in (6a-6d) are compatible with individual objects, as in (7a-7d)

(7a) [pointing] That tiger has stripes.
(7b) [pointing] That bird flies.
(7c) [pointing] That duck lays eggs.
(7d) [pointing] That tick carries Lyme Disease.

Predicates like 'be extinct' are incompatible with instances, so we have assumed that they are predicates of something else, which we've called kinds. But since kinds are abstract, it is difficult to intuit what *other* predicates are compatible with them. The tiger kind can have the 'be extinct' predicate applied to it, but can the tiger kind have the 'have stripes' predicate applied to it? We could take the acceptability of (6a) to suggest that the answer is 'yes', assuming that the bare plural subject in (6a) is mapped to a kind individual just like in (1a) 'Dodos are extinct'.

But if we suspect that the answer is 'no', and that (6a) is really an assertion about individual tigers and not the kind, then we can alternatively seek an analysis in which, unlike (1a), the bare plural subject in (6a) is quantificational, such that the predicate applies to a set of object individuals, not a kind individual.

We already know from sentences like (8) that sentence-initial bare plurals in English sometimes denote sets of individual objects rather than kinds:

(8) Squirrels destroyed my roof last night.

(8) clearly doesn't say that it's a property of the squirrel kind that it destroyed my roof last night. So another way of putting the question is whether sentences like (6a-6d) should be analyzed like (1a) as predications of a kind subject, or more like (8) as involving quantification over a set of individual object instances.

2.1 Universal quantification

Let's consider some versions of a quantificational analysis. First and most simply, the bare plurals could be taken to denote the full set of individuals that bear the property denoted by the noun, e.g. that 'Tigers' in (6a) is taken to denote the set of individuals that bear the tiger property, equivalent to an expression containing a universal quantifier like 'every tiger' or 'all tigers'. However, this idea is widely recognized today to be a non-starter. As Carlson (1977) and many others have established, speakers systematically judge the truth conditions of (6a-6d) to differ from their variants with 'every' and 'all', as in (9a-9d) (given the existence of albino tigers, penguins, male ducks, and harmless ticks). The bare plural subject 'Tigers' is thus not interpreted as containing a universal quantifier, nor as any existing quantifier in English ('most' is out as, e.g. less than 50% of ducks lay eggs).

(9a) Every tiger has stripes.
(9b) All birds fly.
(9c) Every duck lays eggs.
(9d) All ticks carry Lyme Disease.

2.2 Generic [first order] quantification in Carlson (1977)

A second, more seriously entertained possibility is that sentences like (6a-6d) get their interpretations through some *different*, heretoforth unknown type of quantification over individual objects. Such theories analyze these sentences as examples of a broader class of 'generalizing' or 'characterizing' sentences that are not limited to bare plural subjects, like those sentences in (8a-8b).

(8a) John smokes.
(8b) Mary walks the neighbors' dog.

A simple view of these sentence meanings would be that something like (8a) just asserts that the 'smoke' predicate applies to John. But in semantic theories that assume that every assertion must be anchored to a time, then 'John smokes' needs a more complicated analysis, because it does not seem to assert that the 'smoke' predicate applies to John at a particular time.

Influential quantificational theories of (6a-6d) beginning with Carlson (1977) have followed this intuition to suggest a general division between 'particular sentences', making assertions about particular objects or events, and 'generic sentences', asserting generalizations across objects or events. On these views, the difference between the two sentence types is due to different kinds of quantification, where the habitual interpretation of (8a-8b) is due to a silent generic quantifier, in this case over events, which we might call *Gen*. Then, since we have already chosen to assume that the silent *Gen* quantifier drives the generalizing nature of sentences that don't contain bare plural subjects like (8a-8b), we can extend its use to explain the generalizing nature of (6a-6d) by having *Gen* quantify over the set of individual objects denoted by the noun, instead of the kind being predicated directly.

In Carlson's (1977) proposal, the generic was derived in a particular way corresponding to Carlson's theory that predicates were either 'individual-level' or 'stage-level'. For instance individuals, stages correspond to realizations of the instance at different times, and for kind individuals, stages correspond to realizations of the kind in different instances/objects. On this theory something like 'walk the neighbors' dog' is inherently a stage-level predicate, and in a past-perfect version like (9b) would be an assertion about one realization of Mary at a particular point in time. In contrast, in (8b) above, a *Gn* operator applies to the stage-level predicate to change it into a generalizing/characterizing predicate that would apply to the individual object Mary, not stages of Mary.

(9a) John smoked yesterday.
(9b) Mary walked the neighbors' dog.

On this account, if the predicates in (6a-4d) could be considered inherent individual-level predicates, then the genericity of the interpretation could come from the kind subject alone. But for those predicates (like 'fly' in 6b) which seem to have stage-level readings ('That bird flew yesterday but not today'), it must be assumed that the *Gn* operator changed the original stage-level predicate into a predicate of individuals before it could be applied to the kind individual subject.

2.3 Generic [second-order] quantification in Carlson (1989)

By contrast, in most later proposals, *Gen* is an adverbial quantifier relating two (intensional) sets (e.g. Carlson 1989), with a meaning analogous to 'often' or 'usually'. This shift was made to resolve an empirical problem with Carlson's (1977) analysis. The earlier analysis held that generalizing sentences were always predicates of individuals, not stages of individuals. This was

meant to explain why one cannot get an existential interpretation of a bare plural subject together with a generalizing interpretation of the sentence; (4a) cannot mean there exist some tigers with stripes, because in the existential reading the bare plural subject denotes 'stages' of the tiger kind (instances of tigers), and so can't combine with a predicate of individuals. However, Milsark (1974) and others had noted that in fact, some sentences do have existential bare plural subjects together with generalizing interpretations, like (10a). Furthermore, variants with expletive subjects (10d) also have generalizing interpretations.

(10a) Typhoons arise in this part of the Pacific.
(10b) Flowers grow out behind the old shed.
(10c) Guards stand in front of the queen's palace.
(10d) There arise typhoons in this part of the Pacific.

If the syntactic subject is assumed to be the subject of the sentential predication, (10a-10d) clearly violate the predictions of Carlson's (1977) analysis.

There are two related ways out. One, relatively little explored in this literature, is to keep the assumption that generalizaing sentences always involve first-order predication of individuals, but to assume that the subject of the individual-level predication is not the syntactic subject of the sentence, but some other element in the sentence (the clear candidate in 10a-10d is the individual location, e.g. 'this part of the Pacific' in 10a/d).

The other, which is what Carlson (1989) and later work in this literature have pursued, is to drop the first-order subject-predicate analysis and move to a second-order analysis, for example where the *Gen* quantifier is relating the set of locations in this part of the Pacific to the set of locations that have typhoons arising in them, where the bare plural 'typhoons' is interpreted as a propositional function (see Delfitto 2006 for review). The second-order approach is like the first-order subject-predicate approach in that, in place of the asymmetry between subject and predicate, it assumes an asymmetry in the second order relation between sets,where one set is the 'restrictor' of the quantifier resulting in its higher 'scope'. Both approaches also assume that the surface syntax underdetermines which element(s) is taken to have which role in the meaning.

One concern about any version of the covert quantifier approach, is that no known human language appears to have an overt determiner which expresses the posited meaning. Another concern is that it is hard to precisely specify what the *Gen* quantifier does mean with respect to quantity, given that 'often' or 'usually' are event domain analogs to 'most' and thus do not provide the basis for a good gloss for cases like (4c-4d). Both these properties make *Gen* unlike any other quantifier in traditional natural language semantics.

To sum up this section, although the undisputed kind sentences of the (1a-1b) type are the strongest evidence that set theoretic semantics needs to be enriched to handle kinds, there is a much broader range of sentences like (4a-4d) for which kind subjects could be a natural analysis, but where individual-object based analyses are also possible. The most influential such analyses today assume that such sentences contain a covert generic quantifier.

Since undisputed kind sentences like (1a-1b) already require enriching the theory to handle kind individuals, the covert quantifier analyses don't have any particular advantage in terms of parsimony, or simplicity (on both Carlson and Chierchia's approaches, the default interpretation of the noun is as a kind, and the instance-based formulas are more complex). These analyses also require positing quantifiers that are cross-linguistically covert and whose properties seem different

than standard quantifiers. Therefore, we must carefully consider additional facts about scope and modality in determining what is the correct analysis for disputed kind subject sentences. As we have seen, these analyses also depend on broader theoretical choices about whether the habitual sense of simple sentences like 'John smokes' need be explicitly represented in the semantics, and about the semantics of sentences with existential subjects.

3. Instance-of-kind theory

It is to its credit that strictly compositional set theoretic natural language semantics makes strong predictions about what meanings should and should not be possible in natural language. As its domain is sets of individuals, the theory predicts that all natural language sentences should have meanings expressible in terms of sets of individuals, that derive straightforwardly from its parts. In Section 1, we saw that the existence and behavior of undisputed kind subjects requires enrichments to this theory. In Section 2, we saw further evidence of other bare plural subjects which might also be interpreted as kinds, although this is disputed in the literature.

In developing the cognitive, internalist theory of natural language semantics needed for psycholinguistics, we need to have these enrichments in mind if we are to capture phenomena like the undisputed kind subjects in (1a-1b). Sentences with kind subjects seem not to assert something about an individual object (however idealized) or a set of individual objects (however unspecified), but rather about a kind—that is, a mental abstraction of the world structure that results in all individuals not being arbitrarily different from each other. In an internalist semantic theory in which sentences are interpreted against a *mental* model of the world, not the world itself, that model should therefore somehow include not just mental representations of individuals, but also mental representations of kinds.

Crucially, human languages seem to be sensitive to the relationship between individuals and kinds (in that the same noun is used in sentences about both), and therefore in an internalist model we would like these relations to be represented, either in the linguistic semantics itself, or in the conceptual data structures that the language interfaces with.

Several alternative approaches to noun semantics can be considered. In the 20th century philosophy literature, authors like Geach and Wiggins have defended 'substantive' or 'sortal' theories of noun meanings. In these approaches, nouns don't name bare properties or sets of properties or individuals in the domain. Rather they name something different, which we could call a kind, that crucially includes criteria for individuating instances of the kind and determining their identity (e.g. what counts as the 'same' book vs. what counts as the 'same' dog). Syntacticians like Mark Baker (1998) have argued that this approach to noun meaning can help motivate the differing syntactic properties of nouns and adjectives. In the psychology literature, Sandeep Prasada's instance-of-kind theory proposes a corresponding model of the non-linguistic conceptual data structures that nouns are assumed to name.