Larson, R.K. 1990 *Semantics*. From: D. Osherson, ed., *Invitation to Cognitive Science. Vol 1. Language* MIT Press, Cambridge

Chapter 2 Semantics

We have seen that, as speakers of English, for example, we know facts about its syntax: that expressions divide into categories like verb, noun, preposition and adjective, that verbs and prepositions typically precede their objects in English, that words in a sentence cluster into constituents. In addition to these, we know facts about its semantics or meaning structure: that sentences are related a synonymous or contradictory, that they are true under certain circumstances, that certain notions do or do not correspond to possible words.

2.1 Semantical Relations

Like other kinds of linguistic knowledge, knowledge of semantics reveals itself clearly in the form of certain abilities we possess. One such is the ability to judge that various relations hold among sentences. Consider the examples in (1) and (2):

(1) a. John believed that the Earth is flat.
    b. John doubted that the Earth is flat.

(2) a. John claimed that the Earth is flat.
    b. John denied that the Earth is flat.

As speakers of English, we know intuitively, and immediately, that a certain kind of relation holds within the pairs of (1) and (2) - the same one in both cases. Pretheoretically, we grasp it as a relation of "incompatibility" or "exclusion" of some kind.

This exclusion relation does not arise from the grammatical form of the sentences; we know this because other pairs with the same form (subject-verb-complement clause) fail to exhibit the relation:

(3) a. John knew that the Earth is flat.
    b. John dreamed that the Earth is flat.

Likewise it does not arise from the particular phonetic shapes of the words; this is clear because other languages - for instance, German - express the same relation with quite different words ((4) corresponds to (1)):

(4) a. Hans glaubte dass die Erde flach ist.
    b. Hans bezweifelte dass die Erde flach ist.

The relation we detect in (1) and (2) issues from another property of these sentences - their meaning. The members of the pairs "express contrary thoughts", "describe mutually exclusive situations" or "convey opposing information"; they "cannot both be true at the same time", and so on. It is in virtue of the meanings they have that the sentences in (1) and (2) exclude each other. And it is in virtue of knowing these meanings that we judge this relation to hold.

Exclusion is not the only kind of semantic relation we can recognize; (5)-(7) and (8)-(10) illustrate other, analogous forms:

(5) a. John sold a car to Mary.
    b. Mary bought a car from John.

(6) a. John is in front of Mary.
    b. Mary is behind John.

(7) a. John saw Mary.
    b. Mary was seen by John.

(8) a. John is a human.
    b. John is a mammal.

(9) a. Mary was laughing and dancing.
    b. Mary was dancing.

(10) a. Necessarily, apples are Nature's toothbrush.
     b. Apples are Nature's toothbrush.

In (5)-(7) we grasp an identity relation of some kind holding between the members of each pair - a dimension in which the two are fundamentally the same. Likewise in (8)-(10) we detect a relation of "inclusion" or subordination, a respect in which the first member in some sense "implies" the second member. Here again the relevant dimension is that of meaning. The pairs in (5)-(7) are all (largely) identical in meaning or "synonymous", to use the common term. They "express the same thought", "convey the same information", "describe the same situation", and so on. Similarly, the meanings of the first members of (8)-(10) include those of the second members: the thought expressed by *Mary is laughing and dancing* includes that expressed by *Mary is dancing*, the meaning of *human* implies the meaning of *mammal*, and so on. In each case we are able judge a certain relation between sentences, one reducible neither to sound nor form. To account for this ability, we must assume a certain body of knowledge in our possession: knowledge of meaning.

Besides revealing itself in our capacity to judge various kinds of relatedness between sentences, knowledge of linguistic meaning is apparent in our ability to judge relations between language and the world. Consider the example in (11).

(11) The cat is on the mat

As speakers of English, we recognize that a special relation holds between this sentence and the situation depicted in part (a) of figure 2.1 - one that does not hold, for instance, between (11) and situation depicted in part (b). One way of describing this relation is through the familiar notion of *truth*; sentence (11) is true in the (a)-situation, but not in the (b)-situation:

(2.1) (a) (b)

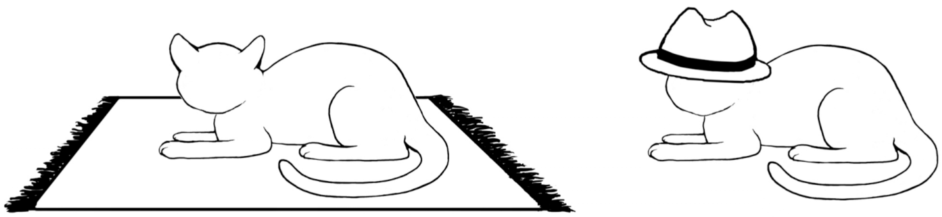

What is it that effects this association between a sentence of English and the world? What is it that we as English speakers know about (11) that allows us to make judgments as to its truth or falsity? Surely not its grammatical form; many sentences with the same grammatical form as (11) (for instance, *the cows are in the corn*) fail to be true in (a)-situation. Not its phonetic properties; the German sentence *Die Katze ist auf der Matte* is also true in the (a)-situation, but is pronounced quite differently. What links the sentence and the situation is meaning. It is in virtue of meaning what it does that (11) is true in the (a)-situation, but not in the (b)-situation. And it is in virtue of knowing this meaning that we can judge its truth or falsity.

2.2 Knowledge of Meaning as Knowledge of Truth-Conditions

These points make clear the reality of our semantical knowledge by showing various things that that knowledge enables us to do. However they do not establish what knowledge of meaning actually is. They do not show precisely what it is we have internalized in acquiring English, German or any other natural language, which grounds our judgements of semantic relatedness or truth and falsity.

To get some insight into this question, consider a simple hypothetical situation. Suppose you're trying to discover whether a foreign friend X knows the meaning of a particular English sentence like (11). You possess a powerful video-display device capable of generating pictures of various conceivable situations, and by using it you find that for any situation presented to X about which you can also make a judgment, X is able to say correctly whether or not sentence (11) is true. That is, for basically the same pictorial situations in which you can give a judgment, he or she is able to say "true" whenever the cat is on the mat, and "false" when it is not. What would you say about X? Does he or she know the meaning of *the cat is on the mat*? Does this kind of evidence settle the matter? Think about this question for a moment.

Intuitively, we are strongly inclined to answer "yes". If X can correctly judge whether the sentence is true or false whenever you can, then X knows the meaning. The evidence seems convincing in the sense that it is difficult to see what further proof we could require of X, or what stronger proof X could provide, that would show that he or she understood the meaning of *the cat is on the mat*. It's hard to see what X could be "missing" that, when combined with this knowledge, would "add up" to knowledge of what (11) means.

This little thought experiment suggests a simple idea about what it is we know when we know the meaning of a sentence, an idea extensively pursued in modern linguistic semantics. It suggests that knowledge of meaning might be fruitfully viewed as knowledge of truth-conditions, i.e., knowledge of something of the form shown in (6), where p gives what the world must be like in order for the sentence in question to be true:

(12) *The cat is on the mat* is true if and only if p

If X has internalized such a piece of knowledge (with p filled in), then X knows the conditions under which *the cat is on the mat* is true. If X knows this, he or she will able judge, for any circumstance, whether (11) is true or not. But we observed above that if X can do this, then, intuitively, he or she knows the meaning of (11). Thus given our thought experiment, "knowing the meaning" seems to be largely captured in terms of "knowing the truth-conditions".

If knowledge of meaning amounts to knowledge of truth- conditions, then we can give a direct account of the semantic abilities discussed above. Recall the examples (1) and (2) (repeated below):

(1) a. John believed that the Earth is flat
b. John doubted that the Earth is flat

(2) a. John claimed that the Earth is flat
b. John denied that the Earth is flat

We said that in virtue of our knowing the semantics of English, we know that the two pairs in (1) and (2) bear a relation of "semantic incompatibility" or "meaning exclusion". Suppose we explicate this relation as follows: two sentences are incompatible if their truth-conditions, together with our real-world knowledge, forbid them from being simultaneously true. Then this will predict incompatibility correctly. The first member of the pair in (1b), for instance, will be true if and only if (iff) John claimed the Earth is flat, and the second will be true iff John denied the Earth is flat. Assuming that we are talking about the same assertion by John, we know that the two cannot be simultaneously true. Any denial of p, is a claim that not-p, and hence not a claim that p. Thus the two exclude each other.

The pretheoretic notions of synonymy ("meaning identity") and hyponymy ("meaning inclusion") can be treated analogously. We can say that two sentences are synonymous if their truth-conditions, taken together with our real world knowledge, entail that they are true in the same circumstances. Likewise, we can say that one sentence implies another if any situation in which the first is true is also one in which the second is true.

Under these proposals the sentence pairs in (5)-(7) will be correctly identified as synonymous. For example, (5a) is true iff John sold a car to Mary, and (5b) is true iff Mary bought a car from John. In virtue of how the world is, any circumstance of the former sort is also one of the latter sort. Hence the two are synonymous. Similarly, the (a)-sentence of each example (8)-(10) will imply the (b)-sentence. Any situation making

(9a) true will make (9b)true as well since any situation in which Mary is laughing and dancing is one in which Mary is dancing. And so on.

Finally, knowledge of truth-conditions will clearly account for our ability to judge that a sentence S is true or false in a given situation. If we know the truth-conditions of S, then knowing whether S is true or false is just a matter of knowing whether these conditions are or are not met.

2.3 Compositionality

Truth-conditions appear to offer a promising approach to sentence meaning and the abilities that flow from knowing them. Let us now consider some facts bearing on the form in which that knowledge is encoded in us. In our exploration of syntax in chapter 1, we saw earlier that the class of well-formed sentences in English - or any other natural language - is essentially boundless. With their grammatical knowledge, human language speakers are able to construct infinite collections of well-formed sentences, such as the following (from Platts (1979)):

(13) a. The horse behind Pegasus is bald
b. The horse behind the horse behind Pegasus is bald
c. The horse behind the horse behind the horse behind Pegasus is bald
.
.

Given this creative aspect of syntactic ability, we know our knowledge of the well-formed sentences of English cannot take the form of a simple list. Since the list is infinite, a finite object with finite storage capacity like our brain simply could not accommodate it. On the basis of this we conclude that syntactic knowledge must be encoded within us in the form of a finite set of rules and principles that allow us to generate the sentences of English from smaller, sub-sentential elements such as words.

Similar issues arise with meaning and knowledge of truth-conditions. The expressions of (13) are not only well-formed sentences of English, they are all meaningful as well. More than that, (13a), (13b), (13c), and so on all have different meanings - or different truth-conditions, as we now say. The first is true only in situations containing at least two horses, the second only in situations containing at least three horses, the third only in situations containing at least four horses, and so on. Since the collection of interpretations associated with (13a-...) is infinite in number, it is clear that our knowledge of truth-conditions for the sentences of English cannot take the form of a simple list like (14). Once again, such a list could not be accommodated in our finite brains:

(14) a. *The horse behind Pegasus is bald* is true iff p1.
b. *The horse behind the horse behind Pegasus is bald* is true iff p2.
c. *The horse behind the horse behind the horse behind Pegasus is bald* is true iff p3.
.
.

Reasoning as above, it seems that our semantic knowledge must also take the form of a set of productive rules or principles that allow us to calculate truth-conditions for sentences from some "smaller" semantic contributions. That is, it appears that the truth-conditions matched with a given sentence of English must be <u>compositionally derived</u>.

How exactly might this go? How are the truth-conditions for a sentence composed, and from what? A plausible hypothesis is that they are calculated using internal syntactic structure. To illustrate, consider (13a), with syntax roughly as shown in figure 2.2, where N' is a nominal constituent intermediate between N and NP. (The phrase structure rules beyond those given in chapter 1 that are needed to generate this structure should be immediately clear.)

(2.2)

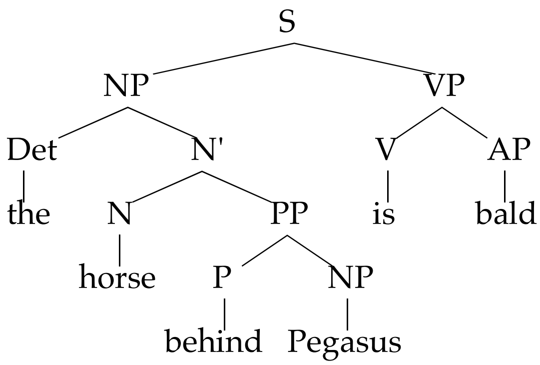

Suppose we could assign some kind of semantic contribution or "value" to each of the "leaves" of this tree structure: a value to *bald*, a value to *Pegasus*, a value to <u>is</u>, and so on. Suppose further that we had a way of combining these values together for each of the branches in figure 2.2 so that, at the top, they yielded the truth-conditions for the sentence: a general way of combining the values of nouns and PPs in the configuration [$_{N}$ N PP], a general way of combining verbs and adjectives in the configuration [$_{VP}$ V AP], a way of combining the values of NPs and VPs in the configuration [$_{S}$ NP VP] to yield the truth-conditions for S, and so on. Then we would in essence be using the syntactic skeleton of figure 2.2 as a guide to figuring out what it means. To borrow a phrase from Quine (1970), semantics would "chase truth up the tree of grammar."

If we could give such initial values and general combination schemes, then we could account very directly for our ability to assign truth-conditions to unbounded sequences of sentences like (13a,b,...). Consider the tree underlying (13b) depicted in figure 2.3:

(2.3)

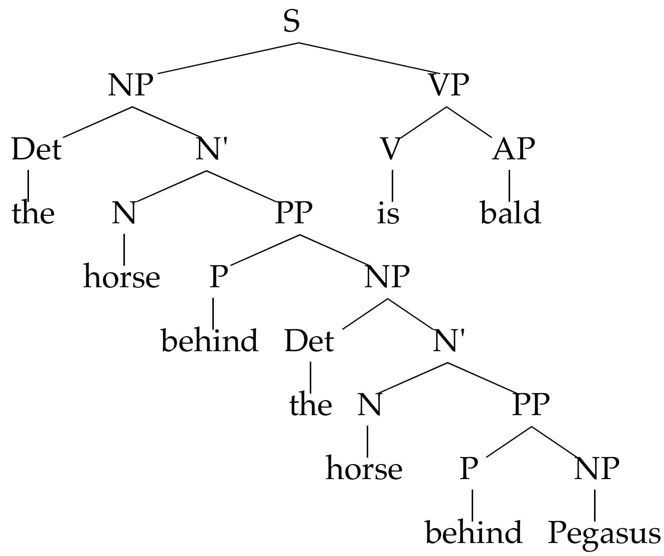

This tree differs from the one underlying (13a) in having extra [$_{NP}$ DET N] and [$_{N}$ N PP] branches involving the lexical items *the, behind* and *horse*. But since all of these elements already occur in (13a), it follows that if we have the semantic resources for computing the truth-conditions of (13a), we will "automatically" have the resources for computing the truth-conditions of (13b). Our semantic values and rules will deliver truth-conditions for (13a), (13b) and indeed all the sentences in the sequence.

2.4 The Study of Meaning: Model-theoretic Semantics

Let us now attempt to implement one version of the picture sketched above, giving elements of a theory that takes truth-conditions as the basis of our semantical knowledge and attempts to derive the truth-conditions for sentences in a compositional way. In doing this, we will adopt the basic perspective of model theory. As we have seen, truth-conditional theories take the view that meaning is fundamentally a relation between language and the world. Interpretation involves systematically correlating sentences with the world through the notion of truth. Model theory studies the relation between languages and worlds in a formal way. It does this by building mathematical models of worlds using the devices of set theory, and by mapping expressions of language into them:

(15)

Language		Model (of a World)
Lexical items	→	Set-theoretic objects
Syntactic rules for building up phrases	→	Rules for combining set-theoretic objects

As (15) shows, constructing a model-theoretic semantics for a language L involves correlating the basic expressions with appropriate set-theoretic objects, and giving rules that state how these constructs combine for each of the syntactic configurations of L. This yields an interpretation - an object in our model - for each of the subsentential expressions of L, and a set of truth-conditions for every sentence.

We can illustrate how a model-theoretic semantics operates by considering a small "sub-language" of English, which we will call L*. L* contains the determiner elements *no, every* and *some*, the common nouns *man* and *fish*, and the intransitive predicates *walks* and *drinks*. It also involves the syntactic configurations shown in figure 2.4 (where X is DET, N or V, and where α is a lexical item). L* thus includes phrases like *some fish*, *no man*, and sentences like *every fish walks*, *some man drinks*.

(2.4)

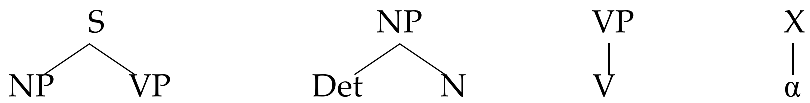

To construct our model-theoretic semantics for L*, we start with some basic set A of individuals - intuitively, the set of entities in our model world M* - and establish a general correlation between the syntactic categories of L* and kinds of set-theoretic objects built out of A. We shall adopt the following mapping:

(16)

Language		Model
Ns, Vs	→	Subsets of A
Dets	→	Binary Relations on Subsets of A

The categories of common noun and intransitive verb are associated with subsets of A. Intuitively, the common noun *fish* is associated with that subset of A that contains the fishes, and the verb *walk* is associated with that subset of A containing the walkers, and so on.

Determiners are matched up with binary relations on sets of As. In particular, we will associate the determiners *the, every*, and *some* with the following specific relations (where '$[\![\alpha]\!]$ ' is to be read 'the interpretation of α'):

(17) Dets:

$[\![\textit{every}]\!]$	=	EVERY, where for any sets X, Y $\subseteq$ A EVERY (X) (Y) iff $X \subseteq Y$
$[\![\textit{some}]\!]$	=	SOME, where for any sets X, Y $\subseteq$ A SOME (X) (Y) iff $X \cap Y \neq \varnothing$
$[\![\textit{no}]\!]$	=	NO, where for any sets X, Y $\subseteq$ A NO (X) (Y) iff $X \cap Y = \varnothing$

The idea behind these assignments descends ultimately from the philosopher Frege, who suggested that determiners correspond to relations between properties or concepts. Thus in examples like *Every whale is a mammal* or *Some whales are mammals,* the determiner serves to relate the properties of whalehood and mammalhood.

In the case of *every*, the relation is one of subordination. Every whale is a mammal if whalehood is a "species" of mammalhood. We have captured the notion of subordination here using the subset relation. In the case of some the relation is one of non-exclusion; some whales are mammals if whalehood and mammalhood are not mutually exclusive properties. Again we capture this using a set-theoretic relation: that of non-empty intersection. The relation expressed by no is similar to the SOME relation. NO holds between two properties like whalehood and mammalhood if the two are mutually exclusive.

Having correlated the syntactic categories of L* with set-theoretic "interpretation spaces", we can construct semantic rules of combination for each of the syntactic configurations in L*. These are as follows:

(18) a. $[\![\, [_{X}\ \alpha\,]\,]\!] = [\![\, \alpha\,]\!]$, where X is Det, N or V

b. $[\![\, [_{VP}\ V]\,]\!] = [\![\, V\,]\!]$

c. $[\![\, [_{NP}\ Det\ \ N]\,]\!] = \{\, Y: [\![\, Det\,]\!]\, ([\![N]\!])\, (Y)\, \}$

d. $[\![\, [_{S}\ NP\ \ VP]\,]\!]$ is true if $[\![VP]\!] \in [\![NP]\!]$ and false otherwise

(18a,b) give the (essentially trivial) interpretation rules for lexical nodes and intransitive VPs. (18c) gives the interpretation of NPs as families of sets - the family of sets that stand in the 'DET-relation' to the set associated with its common noun head. (18d) states that a sentence is true (in our model M*) if and only if the set associated with the verb phrase falls in the family of sets associated with the subject NP.

The initial assignments plus the rules just given determine an interpretation (a set-theoretic counterpart) for every subsentential expression of L*. These in turn determine a set of truth-conditions for every sentence with respect to M*. To see a brief example of how this works, consider the sentence *Some man walks*. In L*, the latter receives the syntax shown in figure 2.5.

(2.5)

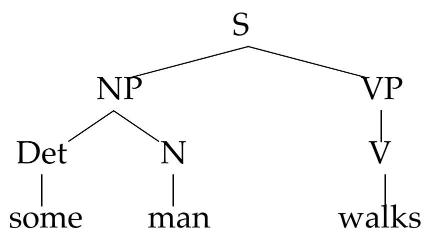

What will its truth-conditions be under our semantics? We compute them compositionally from the interpretations of the parts.

By (18a), the lexical nodes in figure 2.5 all receive the same interpretations as the lexical items they dominate. *Some* corresponds to the binary relation between sets given above:

(19) $[\![[_{Det}\ \textit{some}]]\!] = [\![\textit{some}]\!] =$ SOME

The common noun *man* is interpreted as a set of individuals (intuitively, the set of men in our model):

(20) $[\![[_{\text{N}}\ \textit{man}]]\!] = [\![\textit{man}]\!] = \{\text{ x: x is a man in M* }\}$

Likewise, the intransitive verb *walks* is also mapped to a set of individuals, the set of runners in M*. Taking this result together with (18b) we thus have:

(21) $[\![[_{\text{VP}}[_{\text{V}}\ \textit{walks}]]]\!] = [\![[_{\text{V}}\ \textit{walks}]]\!] = [\![\textit{walks}]\!] = \{\text{ x: x is a walker in M* }\}$

Rule (18c) allows us to combine the results in (19) and (20) and compute the interpretation of NP. *Some man* will correspond to a family of sets - in particular, the family of sets bearing the 'SOME-relation' to the set associated with *man*:

(22) $[\![[_{\text{NP}}\ [_{\text{Det}}\ \textit{some}][_{\text{N}}\ \textit{man}]]]\!] = \{\text{ Y: SOME }([\![\textit{man}]\!])\text{ (Y) }\}$

Given our earlier explication of the 'SOME-relation', this amounts to the following:

(23) $[\![[_{\text{NP}}\ [_{\text{Det}}\ \textit{some}][_{\text{N}}\ \textit{man}]]]\!] = \{\text{ Y: }\{\text{ x: x is a man in M* }\} \cap \text{Y} \neq \varnothing\ \}$

That is, *some man* corresponds to the family of sets having a non-empty intersection with the set of men. Or, more informally, *some man* maps to the family of sets containing at least one man.

Finally, rule (18d) establishes the truth-conditions for the whole sentence. It says that *Some man walks* is true iff the family of sets corresponding to the subject NP contains the set corresponding to the VP. Given the results in (21) and (22) this comes to (24):

(24) $[_{\text{S}}$ *Some man walks*] is true if
$\{\text{ walkers in M* }\} \in \{\text{ Y: }\{\text{men in M*}\} \cap \text{Y} \neq \varnothing\ \}$,
and false otherwise

which is just to say:

(25) $[_{\text{S}}$ *Some man walks*] is true if
$\{\text{ men in M* }\} \cap \{\text{ walkers in M* }\} \neq \varnothing$, and false otherwise

That is, *Some man walks* is true in M* if and only if there is at least one individual who is both a man and a walker in M*. This is intuitively the correct result.

The sample language L* is a very simple one, but it shows how the model-theoretic approach attempts to formalize the basic line of thinking sketched in section 2.3. Using some resources from set theory, we can assign semantic values to basic lexical items and give rules for calculating the values of complex expressions on the basis of their syntax. This ultimately yields truth-conditions for each sentence. To carry this analysis further we would expand the class of lexical items and syntactic configurations in our language L*, but the basic procedure would remain the same.

2.5 Semantical Properties

Modeling of some domain by mathematical constructs has an important consequence, one that is exploited in all science: it allows us to study the properties of the domain through the mathematical properties of the constructs that model it. Since mathematical properties can be stated and manipulated precisely, our understanding gains depth and precision as a result.

These virtues hold in the domain of semantics as well. By modeling aspects of meaning formally, we can capture and study important linguistic properties in a precise way. We will illustrate this briefly with two semantical properties for the category of natural language determiners: directional entailingness and conservativity. For further exploration of the semantics of individual words, see chapter 3, *Lexical Meaning* by Barbara Partee.

2.5.1 Directional Entailingness

In the simple language L* presented above, we modeled determiner meanings with relations between sets and we associated English determiners with certain particular relations. These relations show a number of interesting differences. Consider the inference paradigms for *every* in (26) below (where '#' indicates an invalid inference):

(26) a. Every man runs
____
Every tall man runs

b. Every tall man runs
____
#Every man runs

c. Every man likes a green vegetable
____
#Every man likes spinach

d. Every man likes spinach
____
Every man likes a green vegetable

With sentences involving *every* we get a valid inference whenever we substitute a more specific common noun (*tall man*) for a less specific one (*man*), but not vice versa. On the other hand, we get a valid inference whenever we substitute a less specific VP (*likes a green vegetable*) for a more specific one (*likes spinach*), but again not vice versa.

Rather different patterns of inference emerge with the determiners *some* and *no*:

(27) a. Some man runs
____
#Some tall man runs.

b. Some tall man runs.
———————
Some man runs.

c. Some man likes a green vegetable.
———————
#Some man likes spinach.

d. Some man likes spinach.
———————
Some man likes a green vegetable.

(28) a. No man runs.
———————
No tall man runs.

b. No tall man runs.
———————
#No man runs.

c. No man likes a green vegetable.
———————
No man likes spinach.

d. No man likes spinach.
———————
#No man likes a green vegetable.

Evidently with *some* we must always infer from a more specific to a less specific phrase, whether it is the common noun or the VP. With *no* the situation is just the opposite: we must always infer from a less specific to a more specific phrase.

How can we state the semantical property behind these inference patterns? Taking determiners to correspond semantically to binary relations D between sets, where the common noun supplies the first argument (X) of the relation, and the VP supplies the second argument (Y), the relevant properties can captured as follows:

(29) **Downward Entailingness**: A determiner relation D is

a. **downward entailing in its first argument** if for any X,Y, where $X' \subseteq X$, D(X)(Y) only if D(X')(Y)

b. **downward entailing in its second argument** if for any X,Y, where $Y' \subseteq Y$, D(X)(Y) only if D(X)(Y')

(30) **Upward Entailingness**: A determiner relation D is

a. **upward entailing in its first argument** if for any X,Y where X ⊆ X', D(X)(Y) only if D(X')(Y)

b. **upward entailing in its second argument** if for any X,Y, where Y ⊆ Y', D(X)(Y) only if D(X)(Y')

Thus downwardly entailing environments are ones in which substitution of a set with a subset (from less specific to more specific) yields a valid inference, whereas upward entailing environments are ones in which substitution of a set with a superset (from more specific to less specific) yields a valid inference.

(26a-d) show that EVERY is downward entailing in its first argument, the one corresponding to the common noun, but upward entailing in its second argument, the one corresponding to the VP. Similarly, (27a-d) and (28a-d) show (respectively) that SOME is upwardly entailing in both of its arguments whereas NO is downwardly entailing in both its arguments.

Directional entailingness is a rather simple property of determiners, but one holds considerable interest for linguists: it seems to shed light on certain puzzling facts of English grammar. Consider the distribution of words like *ever* and *anyone*, and phrases like *give a damn*, *budge an inch*. These forms can occur smoothly only in certain rather restricted environments - typically the sort provided by a negative element (a word like *no, not*, or *never*):

(31) a. *John saw **anything**.
b. John didn't see **anything**.

(32) a. *I believe that she will **budge an inch**.
b. I don't believe that she will **budge an inch**.

(33) a. *Max said that he had **ever** been there.
b. Max never said that he had **ever** been there.
c. Max said that he hadn't **ever** been there.

Because of this property, expressions like *ever, anyone, anything, until, give a red cent, lift a finger* are often referred to as <u>negative polarity items</u>.

One interesting question for the study of grammar is, What precisely are the environments in which negative polarity items are licensed? How are they to be characterized? The answer is not self-evident. Note that the licensing environments are not simply those involving negative words like *no, not*, or *nothing*. Negative polarity items are also sanctioned by *every* when they occur in its restrictive term (the bracketed portion in (34a)). They are not permitted in the VP however (34b):

(34) a. Every [person who has **ever** visited Boston] has returned to it.
b. *Every [person who has visited Boston] has **ever** returned to it

This behavior contrasts with that of other determiners such as *no* and *some*:

(35) a. No [person who have **ever** visited Boston] has returned to it.
b. No [person who have visited Boston] has **ever** returned to it.

(36) a. *Some [person who has **ever** visited Boston] has returned to it.
b. *Some [person who has visited Boston] has **ever** returned to it.

The former licenses negative polarity items in both its restrictive term and in the VP. The latter licenses negative polarity items in neither the restrictive term nor the VP.

What is the generalization here? If we consider the directional entailingness properties discussed above, a simple answer suggests itself. Recall that EVERY is downward entailing in its first argument but upward entailing in its second argument. SOME is upwardly entailing in both arguments and NO is downwardly entailing in both:

(37) EVERY (X) (Y) ↓ ↑
SOME (X) (Y) ↑ ↑
NO (X) (Y) ↓ ↓

Recall also that in sentences like (34)-(36), the restrictive term (the bracketed portion) corresponds to the X argument of DET and the VP corresponds to the Y argument. The generalization is clearly the following (from Ladusaw (1980)):

(38) A negative polarity item is licensed in a downward entailing environment.

That is, whenever the phrase containing *anyone, budge an inch*, and so on corresponds to a downwardly entailed argument, the negative polarity item is licensed; otherwise it isn't.

These facts argue strongly that the semantic property of directional entailingness has reality for speakers of English. It is the property to which we are "attuned" in judging the acceptability of sentences containing negative polarity items.

2.5.2 Conservativity

Directional entailingness is a semantical property that distinguishes determiners like *every, some* and *no* in different argument positions. It is of interest because it appears to shed light on specific facts about the grammar of English. There are other semantical properties, however, that *every, some* and *no* share, and these are of interest because they seem to tell us something about human language generally. They appear to give insight into what constitutes a "possible human determiner concept."

One such property that has been studied in some detail is that of **conservativity**. Consider our three determiner relations again:

(39) a. EVERY (X) (Y) iff $X \subseteq Y$
b. SOME (X) (Y) iff $X \cap Y \neq \varnothing$
c. NO (X) (Y) iff $X \cap Y = \varnothing$

Given (39a) we can evaluate the truth of a sentence containing *every* by sorting through the set corresponding to the common noun (X), checking to see if all of its members are in the set corresponding to the VP (Y). Similarly with *some* (39b), we can sort through the common noun set checking to see that some of its members are in the VP set. Finally, with *no* (39c), we can sort through the common noun set and check to see that none of its members are in the VP set. There is an important regularity here. Notice that in each case we can work always with members of the common noun set X in checking whether the given relation holds. The common noun uniformly "sets the scene"; it limits the collection of individuals over which we must range in making the evaluation.

This regularity observed with *every*, *some* and *no* is not found in all quantifier-like relations. Consider the relations expressed by *all-but*, and *everyone-except* as they occur in examples like (40a,b):

(40) a. **All but** boys received a prize.
b. **Everyone except** mothers attended.

Intuitively, to evaluate whether (40a) is true, we do not sort through the set of boys and see if some quantity of them are prize-recipients; rather we must look precisely at non-boys. Similarly, we do not look at mothers but rather at non- mothers in evaluating (40b).

This notion of "setting the scene" or "fixing the collection over which we quantify", which characterizes *every, some* and *no* but not *all-but* and *everyone-except*, is, in essence, the property of conservativity. We may define it more precisely as follows:

(41) A determiner relation D is **conservative** if for any X,Y, D (X) (Y) iff D (X) ($X \cap Y$).

A conservative determiner relation is one that holds between two sets X,Y just in case it holds between the first and the intersection of the first with the second. Since X and $X \cap Y$ are both subsets of X, this means that we always range over members of the common noun denotation in evaluating whether a conservative determiner relation holds. X sets the scene.

Conservativity is a property that appears to characterize all human language determiner concepts - not just *every, some* and *no*, but also *few, many, most, two, three, several*, and so on, and their many counterparts in world languages. It is what might be called a **semantic universal**. This result is quite surprising on reflection, since there is no clear a priori reason why things should be so. There is no sense in which nonconservative determiner relations are "conceptually inaccessible"; nor are they somehow "unnatural" or "unuseful". We have noted informally that *all-but* and *everyone-except* do not share the conservativity property because their common noun does not specify the range of quantification. Notice now that although these expressions are not themselves determiners in English (indeed they are not even syntactic constituents), there is no difficulty in defining a hypothetical determiner relation 'NALL' having exactly their semantics:

(42) NALL (X) (Y) iff $(A - X) \subseteq Y$

(where X and Y are subsets of A, our model's universe of things). Under this definition, a sentence like *Nall squares are striped* would be true, and presumably useful, in exactly the same situations as the sentences *All but squares are striped* or *Everything except squares are striped*:

(2.6)

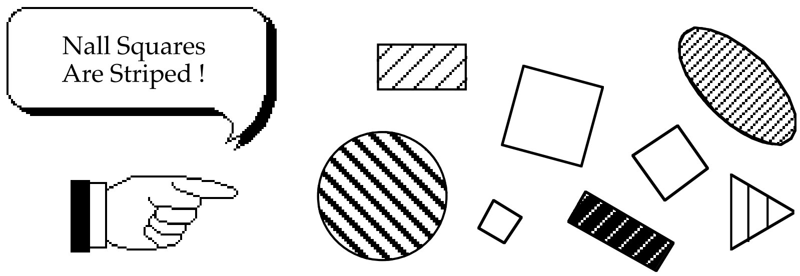

Nall is thus a perfectly reasonable candidate for a natural language determiner relation on general grounds. Nonetheless no such relation occurs in English or in any other human language as far as we know. Nonconservative determiners like *nall* simply seem to be disallowed.

Why conservative determiners should be singled out by natural language is an interesting question that we cannot pursue in detail here. However results by Keenan and Stavi (1986) suggest that this may arise from natural language determiner meanings being typically composed out of certain basic, "atomic" meanings. It can be shown formally that if one begins with elementary determiner concepts such as EVERY, THE (ONE) and POSS(essor), and augments this set with more complex determiner meanings constructed by elementary operations like intersection and complementation, then the result will include only conservative determiners. This is because the "atomic" determiners are all conservative and elementary set-theoretic operations preserve conservativity. The ubiquity of conservative determiners may thus reflect a deep fact about the way our space of determiner concepts is structured: that it forms a so-called Boolean algebra over certain elementary determiner meanings.

Suggestions for Further Reading

The truth-conditional approach to semantics has been developed along a number of different lines. An excellent, extended introduction to the general model-theoretic framework discussed in this chapter is Chierchia and McConnel-Ginet (1990). A good follow-up volume in the same framework, but at a slightly higher level of technical complexity, is Dowty, Wall and Peters (1992).

A classic paper setting out a somewhat different approach to truth-conditions and their relation to meaning is Davidson (1967). Davidson's views are united with Chomskian theory (which emphasizes knowledge of language as the object of linguistic theory) in Larson and Segal (1995). The latter is a textbook that extends Davidson's truth-conditional analysis to many natural language constructions.

A clear, nontechnical introduction to the relational analysis of natural language determiners is Bach (1989), lectures III and IV. See also the chapter on “Generalized Quantifiers” in Partee, ter Meulen and Wall (1990) for an extended introductory discussion of quantifiers and quantifier properties. Technical works in the relational analysis of determiners include: Barwise, J. and R. Cooper (1981), Higginbotham and May (1981) and Keenan, E. and Y. Stavi (1986). For a conceptual, philosophical discussion of the relational account see Wiggins (1980).

The phenomenon of negative polarity is discussed in more detail in Ladusaw (1980) and Linebarger (1987). The first argues for the analysis given in this chapter and the second argues against it.

The truth-conditional approach to natural language meaning is not universally accepted. For criticisms and an alternative line of investigation see Jackendoff (1983,1990).

Study Questions

2.1 Show how the appropriate truth-conditions *Some fish drinks* are derived using our semantic interpretation rules for L*

2.2 The language L* contains *every, some* and *no.* Extend it to include the determiners *two* and *most*. What determiner relations should be associated with these items?

2.3 Extending L* to include proper names like *John* and *Eunice* raises an interesting question. In our semantic theory for L*, NPs are interpreted as sets of sets. But intuitively, it seems the semantic value of <u>John</u> should be an individual j (a person). Can you see a way to interpret proper names that reconciles the view that NPs denote sets of sets with our intuition that *John* ought to be associated with an individual?

2.4 Examine the determiners *all*, *at least two*, *few*, and *not many* with respect to upward and downward entailingness, and check whether the distribution of negative polarity items conforms to these results.

2.5 Do the rules in (17) determine that EVERY is downwardly entailing in its first argument position and upwardly entailing in its second argument position? Similarly do the interpretation rules for SOME and NO allow us to predict their directional entailingness properties?

2.6 In view of the definition of Conservativity, one simple way to check whether a given determiner Det is conservative is to consider the validity of sentences of the following general form, for any nominal A and VP B:

Det A B iff Det A is an A/are As that B

(For instance, *Both men run iff both men are men that run*). If the scheme always yields a true sentence, then the determiner is conservative. If it yields a false sentence, then it isn't. Using this scheme, investigate the conservativity of *every, some* and *exactly two*.

2.7 The expression *only* appears determiner-like in sentences like *Only cats miaow*. However consider the following simple instance of the scheme in 2.6:

Only men run iff only men are men who run

Is this sentence always true or are there situations in which it is false? If the latter, does this overthrow the claim that natural language determiners are conservative, or can you think of a way of defending the latter claim?

2.8 In the chapter, it was claimed that natural language determiners are uniformly conservative. But consider the following sentence due to Westerstahl.

Many Scandinavians have won the Noble Prize.

On reflection this example is ambiguous. What are its truth-conditions? Do any of its readings raise problems for conservativity?

References

Bach, E, (189) *Informal Lectures on Formal Semantics*. SUNY Press, Albany, NY.

Barwise, J. and R. Cooper (1981) "Generalized Quantifiers and Natural Language," *Linguistics and Philosophy, 4*, 159-219.

Chierchia, G. and S. McConnell-Ginet (1990) *Meaning and Grammar*. MIT Press, Cambridge, MA.

Davidson, D. (1967) "Truth and Meaning," in *Inquiries into Truth and Interpretation*. Oxford Univ. Press, Oxford, U.K.

Dowty, D., R. Wall and S. Peters (1992) *An Introduction to Montague's Semantic Theory* (second edition). Kluwer, Dordrecht.

Higginbotham, J. and R. May (1981) "Questions, Quantifiers and Crossing," *The Linguistic Review 1*, 41-79.

Jackendoff, R. (1983) *Semantics and Cognition*. Bradford Books/MIT Press, Cambridge, MA.

Jackendoff, R. (1991) *Semantic Structures*. MIT Press, Cambridge, MA.

Keenan, E. and Y. Stavi (1986) "A Semantic Characterization of Natural Language Determiners," *Linguistics and Philosophy, 9*, 253-326.

Ladusaw, W. (1980) "On the Notion 'Affective' in the Analysis of Negative Polarity Items," *J. of Linguistic Research, 1* .

Larson, R. and G. Segal (1995) *Knowledge of Meaning*. Bradford Books/MIT Press, Cambridge, MA.

Linebarger, M. (1987) "Negative Polarity and Grammatical Representation," *Linguistics and Philosophy, 10*, 325-387.

Partee, B. A. ter Meulen and R. Wall (1990) *Mathematical Methods in Linguistics*. D. Reidel, Dordrecht.

Platts, M. (1979) *Ways of Meaning*. Routledge and Kegan Paul, London.

Quine, W. V. O. (1970) *Philosophy of Logic*. Prentice-Hall, Englewood Cliffs, NJ.

Wiggins, D. (1980) "'Most' and 'All': Some comments on a familiar program, and on the logical form of quantified sentences," in Platts, M. (ed.) *Reference, Truth and Reality*. Routledge and Kegan Paul, London, 318-346.

Answers to Questions

2.1 Show how the appropriate truth-conditions Every fish drinks are derived using our semantic interpretation rules for L*

ANSWER:

In L*, *Every fish drinks* receives the syntax shown in below:

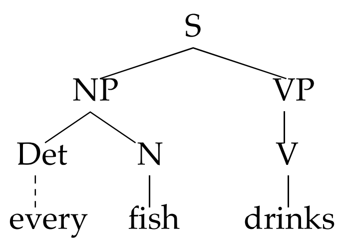

The relevant steps in interpreting this tree are as follows:

(i) $[\![\, [_{DET}\ \text{every}] \,]\!] = [\![\textit{every}]\!] = \text{EVERY}$ (by (18a))

(ii) $[\![\, [_{N}\ \text{fish}] \,]\!] = [\![\underline{\text{fish}}]\!] = \{\, x\text{: } x \text{ is a fish in } M^* \,\}$ (by (18a))

(iii) $[\![[_{VP}\ \text{drinks}]]\!] = [\![[_{V}\ \text{drinks}]]\!] = [\![\underline{\text{drinks}}]\!] = \{\, x\text{: } x \text{ is a drinker in } M^* \,\}$ (by (18a,b))

(iv) $[\![[_{NP}\ \text{every fish}]]\!] = \{\, Y\text{: EVERY } ([\![\underline{\text{fish}}]\!])\ (Y) \,\}$ (by (18c)

(v) $[\![[_{NP}\ \text{every fish}]]\!] = \{\, Y\text{: } \{\, x\text{: } x \text{ is a fish in } M^* \,\} \subseteq Y \,\}$ (by (17))

(vi) $[_{S}$ Every fish drinks] is true if
$\{\, \text{drinkers in } M^* \,\} \in \{Y\text{: } \{\, \text{fish in } M^* \,\} \subseteq Y \,\}$,
and false otherwise (by (18d))

(vii) $[_{S}$ Every fish drinks] is true if
$\{\, \text{fish in } M^* \,\} \subseteq \{\, \text{drinkers in } M^* \,\}$,
and false otherwise (by (vi))

2.2 The language L* contains *every, some* and *no*. Consider extending it to include the determiners *two* and *most*. What determiner relations should be associated with these elements?

ANSWER:

Two. If two is understood as meaning 'exactly two', then the appropriate determiner relation is (ia). If two is understood as 'at least two' the determiner relation is (ib).

(i) a. TWO(X)(Y) iff $|X \cap Y| = 2$
b. TWO(X)(Y) iff $|X \cap Y| \geq 2$

Most. Under the usual understanding of most as 'more than half', the appropriate determiner relation is (ii):

(ii) MOST(X)(Y) iff $|X \cap Y| > |X - Y|$

For the example Most fish drink, (ii) would specify this sentence to be true if the number of individuals in M* that are both fish and drinkers ($|X \cap Y|$) is greater than number of individuals in M* that are fish but not drinkers ($|X - Y|$).

2.3 Extending L* to include proper names like John and Eunice raises an interesting question. In our semantic theory for L*, NPs are interpreted as sets of sets. But intuitively, it seems the semantic value of John should be an individual j (a person). Can you see a way to interpret proper names that reconciles the view that NPs denote sets of sets with our intuition that John ought to be associated with an individual?

ANSWER:

The correct result can be obtained by letting John correspond to the set of sets that contain j, the person. (Ans similarly for other proper names) That is:

(i) a. $[\![[_{NP} \text{John}]]\!] = \{ X: j \in X \}$
b. $[\![[_{NP} \text{Eunice}]]\!] = \{ X: e \in X \}$

The reader can easily show that under this interpretation, we derive the truth-conditions in (iia) for John walks. The latter further reduces to (iib):

(ii) a. $[_S$ John walks] is true if $\{$ walkers in M* $\} \subseteq \{ Y: j \in Y \}$, and false otherwise (by (18d))

(ii) b. $[_S$ John walks] is true if $j \in \{$ walkers in M* $\}$, and false otherwise

2.4 Examine the determiners all, at least two, few, and not many with respect to upward and downward entailingness, and check whether the distribution of negative polarity items conforms to these results.

ANSWER:

	1st Arg	2nd Arg
All	↓	↑
At least two	↑	↑
Few	↓	↓
Not Many	↓	↓

2.5 Do the rules in (17) determine that EVERY is downwardly entailing in its first argument position and upwardly entailing in its second argument position? Similarly do the interpretation rules for SOME and NO allow us to predict their directional entailingness properties?

ANSWER:

The directional entailingness properties of all three determiners can be derived from the interpretations in (17) given some basic results of set theory:

(i) EVERY (X) (Y) iff $X \subseteq Y$

a. For all X,Y,X', if $X \subseteq Y$ and $X' \subseteq X$, then $X' \subseteq Y$
$X' \subseteq Y$ iff EVERY (X') (Y)
Therefore, EVERY (X) (Y) and $X' \subseteq X$ implies EVERY (X') (Y)
That is, EVERY is downward entailing in its first argument position.

b. For all X,Y,Y', if $X \subseteq Y$ and $Y \subseteq Y'$, then $X \subseteq Y'$
$X \subseteq Y'$ iff EVERY (X) (Y').
Therefore, EVERY (X) (Y) and $Y \subseteq Y'$ implies EVERY (X') (Y)
That is, EVERY is upward entailing in its second argument position.

(ii) SOME (X) (Y) iff $X \cap Y \neq \emptyset$

a. For all X,Y,X', if $X \cap Y \neq \emptyset$ and $X \subseteq X'$, then $X' \cap Y \neq \emptyset$
$X' \cap Y \neq \emptyset$ iff SOME (X') (Y).
Therefore, SOME (X) (Y) and $X \subseteq X'$ implies SOME (X') (Y)
That is, SOME is upward entailing in its first argument position.

b. For all X,Y,Y', if $X \cap Y \neq \varnothing$ and $Y \subseteq Y'$, then $X \cap Y' \neq \varnothing$
$X \cap Y' \neq \varnothing$ iff SOME (X) (Y')
Therefore, SOME(X)(Y) and $Y \subseteq Y'$ implies SOME (X) (Y')
That is, SOME is upward entailing in its second argument position.

(iii) NO (X) (Y) iff $X \cap Y = \varnothing$

a. For all X,Y,X', if $X \cap Y = \varnothing$ and $X' \subseteq X$, then $X' \cap Y = \varnothing$
$X' \cap Y = \varnothing$ iff NO (X') (Y)
Therefore, NO (X) (Y) and $X' \subseteq X$ implies NO (X') (Y)
That is, NO is downward entailing in its first argument position.

b. For all X,Y,Y', if $X \cap Y = \varnothing$ and $Y' \subseteq Y$, then $X \cap Y' = \varnothing$
$X \cap Y' = \varnothing$ iff NO (X) (Y')
Therefore, NO (X) (Y) and $Y' \subseteq Y$ implies NO (X) (Y')
That is, NO is downward entailing in its second argument position.

2.6 In view of the definition of Conservativity, one simple way to check whether a given determiner Det is conservative is to consider the validity of sentences of the following general form, for any nominal A and VP B:

$$\text{Det A B iff Det A} \left\{ \begin{array}{l} \text{is an A} \\ \text{are As} \end{array} \right\} \text{that B}$$

(For instance, <u>Both men run iff both men are men that run</u>). If the scheme always yields a true sentence, then the determiner is conservative. If the scheme yields a false sentence, then it isn't. Using this scheme, investigate the conservativity of *every*, *some* and *exactly two*.

ANSWER:

Under Conservativity, the following equivalences are predicted to be true:

Every bird flies iff every bird is a bird that flies
Some bird flies iff some bird is a bird that flies
Exactly two birds fly iff exactly two birds are birds that fly

The equivalences are in fact judged to be true. Furthermore, the same results appear to hold regardless of our choice of nominal (here *bird*) or predicate (here *flies*). Hence *every*, *some* and *exactly two* are conservative.

2.7 The expression *only* appears determiner-like in sentences like *Only cats miaow*. However consider the following simple instance of the scheme in 2.6:

Only men run iff only men are men who run

Is this sentence always true or are there situations in which it is false? If the latter, does this overthrow the claim that natural language determiners are conservative, or can you think of a way of defending the latter claim?

ANSWER:

The sentence "Only men run iff only men are men who run" is not always true. This can be seen from the fact the sentence on the left hand side of "iff" is false (since, for example, ostriches run) whereas the part on the right hand side is always true only men can be running men, etc.). hence the equivalence does not hold for *only* – it does <u>not</u> express a conservative relation.

However, only also does not appear to be a determiner! Consider the following distributional differences between *only* and every:

Only/***All** every man was present
The man **only**/***all** stared.
The contamination was observed **only**/***all** in the park

Only occurs in many places forbidden to determiners. Grammatically, it has the distribution of an adverb like *even*, not a determiner. Hence the non-conservativity of *only* does not appear to be a threat to our generalization about determiners.

2.8 In the chapter, it was claimed that natural language determiners are uniformly conservative. Consider sentence (1) however (due to Westerstahl).

(1) Many Scandinavians have won the Noble Prize.

On reflection this example is ambiguous. What are its readings? Do any of them raise problems for conservativity?

ANSWER:

Sentence (1) has the two very different readings in (i) and (ii):

(i) With respect to Scandinavians, the number of Nobel Prize winners among them is many.

(ii) With respect to Nobel Prize winners, the number of Scandinavians among them is many.

Reading (i) accords with conservativity in that the sense that the quantification is over the individuals given by the common noun <u>Scandinavian</u>. By contrast, reading (ii) is not conservative since the quantification is over individuals given by the predicate have won the Nobel Prize, and not by the common noun. Reading (ii) thus represents a problem for the claim that natural language determiners are uniformly conservative.