Semantics

[Encyclopedia Article]

P. Schlenker

(Institut Jean-Nicod and NYU)

First draft (non-final, June 30, 2008).

Written for the *Linguistics Encyclopedia,* ed. K. Malmkjaer (Routledge)

Semantics is the study of meaning. In contemporary linguistics, it has generally taken the form of a *theory of truth*, which borrows its technical tools from mathematical logic. This connection between meaning and truth is motivated by the observation that a speaker who knows the meaning of a sentence knows, at the very least, under what conditions it is true or false. To put it differently, if we provide any speaker with a syntactically well-formed sentence *S*, together with a sufficiently detailed situation (i.e. a description of the way the world is), the speaker should in principle be in a position to determine whether *S* is true or false. It is clear that speakers have this ability for infinitely many distinct sentences and situations, which they could not all have memorized. Therefore they must have access to certain *rules* that allow them to compute the truth conditions of complex sentences on the basis of memorized facts about their smallest component parts. *Lexical semantics* is concerned with the meaning of these smallest parts - words or morphemes. In this chapter we will be concerned with *compositional semantics*, which seeks to uncover the rules by which complex meanings can be formed.

1 Meaning and Truth

By its very nature, semantics establishes a relation between well-formed sentences and the world. An older tradition viewed semantics as *translational* in nature. The idea was that sentences are interpreted by way of translation, typically into a 'language of thought'. The problem, however, is that such a translational approach does not explain how speakers have knowledge of truth conditions, *unless* some semantic rules have already been provided for the language of thought. It is in the end an empirical question whether such a translational process is indeed at work in the speaker's mind. But what is clear is that at some point some (non-translational) semantic rules must be provided to relate (some) language to the facts of the world. In the following paragraphs we will follow much contemporary research in taking these rules to be specified for some level of syntactic representation of natural language sentences (rather than for a hypothetical language of thought).

To obtain explicit and predictive theories, semantics has generally borrowed its technical tools from mathematical logic. This is a somewhat ironic historical development; for modern logic was viewed by many of its pioneers as a way to *redress* the shortcomings of natural language, which was deemed too vague and ambiguous to be suitable for complex scientific argumentation. But after formal languages were studied with great rigor in the first half of the 20th century, two pioneers of modern linguistics had the idea of treating *English* (or for that matter other natural languages) as if it were a formal language: Noam Chomsky created the field of formal syntax in the 1950's, while Richard Montague founded 'model-theoretic semantics' in the 1960's (the name 'model-theoretic' comes from a branch of logic that studies the interpretation of formal languages). Montague built on key insights of the

Polish logician Alfred Tarski, who had shown in the first half of the century how to give a rigorous definition of truth for formal languages. As soon as English was itself treated as a formal language, it became natural to extend Tarski's program to natural language so as to account for meaning. Chomsky and Montague both engaged in a kind of 'reverse engineering': instead of stipulating formal languages whose syntactic or semantic properties they then studied, they started from the observed properties of English sentences and tried to infer by which syntactic or semantic rules they were created. The Chomskyan and the Montagovian traditions were largely unified in the 1980's and 1990's with the advent of rigorous studies of the 'syntax/semantics interface'. Both traditions contributed to a broad investigation of the universal properties of language, its parameters of variation, its acquisition by children, its impairment after brain lesions, and more generally for its implementation in the brain. Gradually, then, formal semantics has become integrated into the general program of cognitive science - a development which is only at its early stages.

Minimally, a semantic theory should specify rules by which the truth conditions of complex sentences are computed on the basis of memorized properties of words or morphemes, together with a specification of the syntax (derivation tree) of the sentence at hand. Semantic theories must thus satisfy the following condition:

(1) The meaning of any expression is determined from the meaning of its smallest parts and the way they are put together.

Often, however, semanticists have attempted to meet a more stringent requirement, which demands that the meaning of a sentence be determined by the meaning of its *immediate parts* and the way they are put together. This has the effect of limiting the information accessible to semantic rules (example: the sentence *[Mary [saw John]]* has two immediate parts, the Noun Phrase *Mary* and the Verb Phrase *saw John*; but of course it contains three ultimate parts). This principle, which has been the object of sophisticated formal discussions, is called the Principle of (Strong) Compositionality, stated in (2) (the principle in (1) is sometimes called principle of Weak Compositionality):

(2) The meaning of any expression is determined from the meaning of its *immediate* parts and the way in which they are put together.

Contemporary semanticists often implement their theories within an even more stringent framework, called 'type theory', in which the meaning of any complex expression is obtained by applying the meaning of one of its immediate parts, seen as a function, to the meaning of its other immediate part, seen as argument. While it is by no means the only possible framework for semantics, it has the advantage of brevity, and is thus worth discussing in greater detail. Type theory is developed by choosing an inventory of elementary types, which are sets of objects of a particular sort - for instance, one generally takes t to be the type of truth values, assimilated to {0, 1} (with the convention that 0 represents falsity, and 1 represents truth); and e is the type of individuals, assimilated to a domain D of objects such as persons, things, etc. From elementary types, one builds complex types recursively, using the following rule:

(3) If τ_1 and τ_2 are types, $<\tau_1, \tau_2>$ is a type.
(Notational variant: one also sees the notation $\tau_1 \rightarrow \tau_2$ for $<\tau_1, \tau_2>$)

$<\tau_1, \tau_2>$ is intended to denote the set of functions that take objects of type τ_1 as input and return objects of type τ_2 as outputs. In the simplest framework, *it is raining* has type t (because a clause has a truth value), while the meaning of negation (for instance the expression *it's to the case that*, analyzed for simplicity as a single lexical item) is an example of a function of type <t, t>: it takes a clause, of type t, as argument, and forms with it another clause, which also has type t. The case of conjunction is more interesting. When one studies formal languages, one can perfectly well decide that *F and G* is a well-formed formula, which has truth conditions that are specified by a logical rule, but that *and G* on its own has no meaning. Syntacticians have often argued that in English *[F [and G]]* - for instance *[it is raining [and it is cold]]* has an asymmetric structure, in which *and G* forms a natural unit, called a 'constituent'. When this assumption is adopted, the type-theoretic framework makes it possible to assign a meaning to *and G*: it is a function that takes a truth value as argument (in our example, the value of *F*) and returns a truth value as output (in our example, the truth value of the entire conjunction *[F [and G]]*); in other words, it is a function of type <t, t>. From this, we can infer that *and* must itself have a meaning of type <t, <t, t>>: it takes as argument a function of type t (here, the meaning of *G*) and returns a function of type <t, t>, as is summarized in the derivation tree in (4), in which each constituent is annotated with its type.

(4)

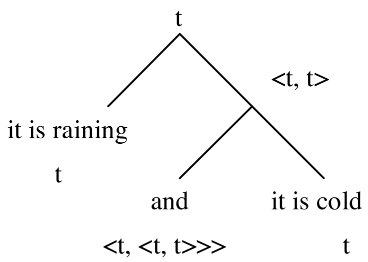

The type-theoretic approach can in principle be applied to all other expressions. For instance, proper names are plausibly taken to denote individuals, and they are thus of type e. What about intranstivie predicates (e.g. intransitive Verb Phrases)? Well, they combine with object-denoting expressions to produce truth values; so they must be functions of type <e, t>: they take as argument an individual, and return a truth value. The same analysis can be applied to Transitive Verb Phrases: *saw John* can be seen as an intransitive predicate, so it is of type <e, t>. Since *John* is of type e, it follows that *saw* is of type <e, <e, t>>. This analysis is generally extended to Noun Phrases as well, with the (considerably slimmer) argument that one can say things like *John is Dean*, and that *be* is plausibly vacuous semantically.

(5) Types of some common expressions

a. Clauses	it-is-raining	t
b. Proper Names	John	e
c. Intransitive Verb Phrases	smokes	<e, t>
d. Noun Phrases	Dean	<e, t>
e. Transitive Verb Phrases	saw	<e, <e, t>>
f. Negation	not, it-is-not-the-case-that	<t, t>
g. Conjunction and disjunction	and, or	<t, <t, t>>

Of course types only specify the sort of function denoted by each expression; to obtain truth conditions, we need to specify the *precise* function in question. One often writes the denotation of an expression *E* as [[E]]. Using this convention, we can for instance posit the following lexical entries, where we define functions explicitly using arrows (we assume that there are two individuals in the domain of discourse, j and m):

(6) [[John]] = j, [[Mary]] = m

[[smokes]] = $\begin{bmatrix} j \rightarrow 1 \\ m \rightarrow 0 \end{bmatrix}$

$\begin{bmatrix} j \rightarrow 0 \\ m \rightarrow 0 \end{bmatrix}$

[[drinks]] =

To these lexical entries, we can add a unique rule, called 'function application', which specifies how any binary subtree (i.e. any subtree with two immediate parts) should be interpreted:

(7) $\left[\!\left[\begin{array}{cc} & \wedge & \\ a & & b \end{array} \right]\!\right]$ = [[a]] ([[b]]) or [[b]] ([[a]]), whichever one is admissible given the rules of type theory.

This makes it possible to compute the value of *John smokes*, which turns out to be true:

(8) [[$\begin{array}{cc} & \wedge & \\ \text{John} & & \text{smokes} \end{array}$]] = [[smokes]] ([[John]]) = $\begin{bmatrix} j \rightarrow 1 \\ m \rightarrow 0 \end{bmatrix}$(j) = 1

Importantly, it is *derivation trees* produced by the syntax rather than strings of words that must be interpreted. This is essential to deal with structurally ambiguous sentences, i.e. sentences that can be assigned distinct structures, as is illustrated in (9):

(9) It will rain and it will be cold or it will be windy.
a. Reading 1: [It will rain and [it will be cold or it will be windy]]
b. Reading 2: [[It will rain and it will be cold] or it will be windy]

The sentence is ambiguous: on Reading 1, it entails that it will rain; on Reading 2, it doesn't. If the input to the semantic component were a string of words without structure, we just wouldn't know which truth conditions to assign to this sentence. The problem disappears if derivation trees rather than strings are the input to semantic interpretation.

With this framework in place, we can ask the basic questions of semantics:
(i) What are the primitive objects that must be postulated to interpret sentences? All standard semantics posit a domain of individuals, and some truth values. But the nature of the truth values is already a matter of debate: besides *true* and *false*, some add a third truth value, #, to handle certain cases of semantic failure (see below). Furthermore, beyond the individual domain, some researchers posit times and so-called "possible worlds", some posit events or situations, and some posit a combination of those. We survey some of these possibilities below.
(ii) What are the rules of interpretation? In the simplest type-theoretic framework, there is a single rule of interpretation, which is function application (defined in (7)). Typically, however, such a system needs to be constrained by additional semantic rules, which are often of a different nature; we will give some examples of this sort.
(iii) How does semantics interact with other modules of the mind - especially syntax and pragmatics? The latter question is by no means trivial. In many cases, the *data* the linguist studies could be analyzed along syntactic, semantic or pragmatic lines. To take an example, there are at least three ways in which one could explain why the examples in (10) are odd:

(10) a. #John is pregnant.
b. #An idea is sleeping.

First, one could claim that the deviance is *syntactic*, because predicates come in the syntax with conditions on the features of their subjects (so-called 'sub-categorization frames'). Thus *John is pregnant* might be deviant because *pregnant* requires a subject with +feminine features, whereas *John* is masculine. Similarly, *An idea is sleeping* might be ungrammatical because *sleeping* demands a subject with +animate features, which is not the case of the Noun Phrase *idea*. An alternative account is to take the deviance to be purely semantic in nature. According to this analysis, there are not two but three truth values: 1, 0 and # - the latter of which encodes semantic failure. Under this view, the function denoted by *be pregnant* yields the value # when it is applied to an argument which is not female. Finally, we could take these sentences to be pragmatically deviant, in the sense that general rules of communicative exchange make them infelicitous; for instance, one could posit that these statements are semantically false, but that they are *trivially* so, and hence systematically uninformative and thus useless. We leave this question open, but we will encounter below several other cases in which the precise boundary between syntax, semantics and pragmatics is a topic of considerable contemporary interest.

2 Individuals: Quantifiers and Pronouns

The *lingua franca* of semantics is the theory of pronouns and quantifiers, which is easiest to develop with respect to individual talk, although it turns out to have important applications beyond it.

2.1 Quantifiers

Let us start with quantifiers, i.e. expressions such as *some student*, *every professor*, *most Frenchmen*, etc. A crucial insight, due to the German philosopher Frege, was that quantifiers are 'second-order properties': they do not denote objects, but rather they say something about the extension of a predicate. For instance, *A student is sick* says that the extension of the predicate *is sick* contains a student; similarly, *Every professor is sick* says that the extension of *is sick* contains every professor. Frege's logic was primarily intended as a tool to study mathematics. It gave rise, among others, to first-order logic, which includes the universal quantifier $\forall$ and the existential quantifier $\exists$, as in the formulas $\forall x\, S(x)$ and $\exists x\, S(x)$ (the logic is called 'first-order' because the quantifiers range over individuals, not over properties of individuals). But there are two crucial respects in which natural language quantifiers differ from these. First, natural language quantifiers are *restricted*: even if we read $S(x)$ as *x is sick*, the formula $\forall x\, S(x)$ ends up making a claim about every object in the universe of discourse, whereas the sentence *Every student is sick* only makes a claim about *students*; in other words, to evaluate the truth of the claim, we may restrict attention to those individuals in the universe that are students. But there is a second respect in which natural language quantifiers differ from their counterparts in logic: they include a variety of numerical quantifiers which do not exist in first-order logic, and furthermore could not even be defined within it. For instance, *most professors are sick* is a statement whose truth conditions could not be defined even if we gave ourselves all of first-order logic, together with a predicate P for *professor* and a predicate S for sick (in fact, a stronger result holds - undefinability would still hold if we gave ourselves an unrestricted quantifier *most things in the universe*).

Semanticists have thus generalized the notion of quantifier used in logic to handle these cases. To do so within the type-theoretic framework which was sketched above, we can reason as follows: *professor* and *smoke* are both expressions of type <e, t>. The syntax of *every professor smokes* suggests that *every professor* is a syntactic constituent, as is *smokes*, of course. For the sentence to return a truth value, *every professor* must have type <<e, t>, t>: it takes a predicative expression as an argument (here: *smokes*, of type <e, t>), and returns a truth value. This can be seen to implement Frege's intuition that quantifiers are predicates of predicates: if we write P = <e, t> for the type of predicates, we see that *every professor* has type <P, t> - which is analogous to the type <e, t> of intransitive predicates, except that the type of individuals e has been replaced with the type P of predicates. Going one step further, this means that *every* must itself have the complex type <<e, t>, <<e, t>, t>, which we can write more legibly as <P, <P, t>>. In effect, we can view *every* (also called a generalized quantifier, or less ambiguously a 'determiner') as a transitive predicate of predicates (just like *see* was a transitive predicate of individuals, of type <e, <e, t>>). This is illustrated in (11).

(11)

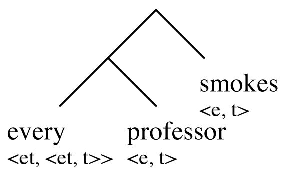

Let us now call EVERY, A, MOST etc. the denotations of the relevant determiners, for which we will now provide truth conditions. For perspicuity, we assimilate PROFESSOR and SMOKE, which are technically functions from individuals to truth values, to sets - the set of individuals that are professors or smokers respectively. Writing ^{C}SMOKE for the complement of SMOKE, i.e. the set of non-smokers, and using $|\ ...\ |$ to refer to the size of a set, we can give the following truth conditions ('iff' abbreviates 'if and only if', and $\cap$ represents set-theoretic intersection):

(12) a. (EVERY(PROFESSOR))(SMOKE) = 1 iff $|\text{PROFESSOR} \cap {}^{C}\text{SMOKE}| = 0$
b. (A(PROFESSOR))(SMOKE) = 1 iff $|\text{PROFESSOR} \cap \text{SMOKE}| \geq 1$
c. NO(PROFESSOR))(SMOKE) = 1 iff $|\text{PROFESSOR} \cap \text{SMOKE}| = 0$.
d. (MOST(PROFESSOR))(SMOKE) = 1 iff $|\text{PROFESSOR} \cap \text{SMOKE}| > |\text{PROFESSOR} \cap {}^{C}\text{SMOKE}|$
d. (AT LEAST THREE(PROFESSOR))(SMOKE) = 1 iff $|\text{PROFESSOR} \cap \text{SMOKE}| \geq 3$
e. (LESS THAN SEVEN(PROFESSOR))(SMOKE) = 1 iff $|\text{PROFESSOR} \cap \text{SMOKE}| < 7$
f. (EXACTLY FIVE(PROFESSOR))(SMOKE) = 1 iff $|\text{PROFESSOR} \cap \text{SMOKE}| = 5$

To take an example, (12)a means that (EVERY (PROFESSOR))(SMOKE) has value 1 (for 'true') just in case the intersection of the set of professors with the set of non-smokers has size 0 - in other words, all the professors are smokers. Similarly, (12)d means that (MOST (PROFESSOR))(SMOKE) has value 1 just in case the number of professors who smoke is greater than the number of professors who don't smoke - which seems like a reasonable approximation of the meaning of this determiner.

With this framework in mind, it is natural to ask which determiner meanings are instantiated in the world's languages. Researchers have found that several semantic constraints are generally satisfied. Two are worth mentioning:

-Natural language determiners are *numerical*: they only 'count' elements that satisfy certain properties (here: PROFESSOR and SMOKE), without discriminating on the basis of their particular identity. So there is no determiner that could crucially depend on the fact that, say, John as opposed to Bill is a professor[1]. In the lexical entries in (12), this property is reflected by the fact that it is only the *size* of certain sets that matters, and not the particular objects they contain.

-Natural language determiners are *conservative*: they only 'care' about those individuals that satisfy their nominal argument. So for instance to determine whether *most professor smokes*, we only need to consider individuals that are professors, and do not need to worry about non-professors. To be more precise, *most* is conservative because no matter who the professors and who the smokers are, *no professor smokes* is true just in case *no professor is a professor*

[1] This property turns out to be directly connected to one that Tarski used when he sought to define what is a logical operation. His idea was that, in essence, an operation is logical just in case it never discriminates among objects on the basis of their identity (technically, this property is called 'isomorphism-invariance'). Interestingly, however, when Tarski's notion is applied to characterize the class of quantifiers (i.e. expressions of type <<e, t>, t>) which count as logical, one obtains a much richer system than first-order logic - in fact, the resulting class is closer to the quantifiers that are in fact instantiated in the world's languages.

that smokes: in evaluating the predicate argument of *no*, we can repeat the nominal argument without modifying the truth conditions. Conservativity is a constraint that certainly has some 'bite': one could plausibly analyze *only* as a determiner that fails to obey it, because to check whether it is true that *only professors smoke* we definitely have to check whether some non-professors do (as a result, *only professors are professors that smoke* does not have the same truth conditions as *only professors smoke*). So *if* the word *only* were a determiner, it would not be conservative. As it turns out, the syntactic distribution of *only* strongly suggests that is not a determiner, as witnesses by the fact it can appear in a variety of environments in which determiners never show up (for instance right before predicates, as in *John is only sick, he isn't dying).*

The semantic study of natural language determiners (or 'generalized quantifiers') has led to important insights about phenomena that had traditionally been treated in syntactic terms. One celebrated example concerns the licensing of Negative Polarity Items such as *ever*, *any, at all,* which in simple examples require a negative expression to license them:

(13) a. John has visited Paris.
b. John hasn't ever visited Paris.
c. *A tourist who has been to France has visited Paris
d. *A tourist who has ever been to France has visited Paris.
e. No tourist who has been to France has ever visited Biviers.
f. No tourist who has ever been to France has visited Biviers.
g. *Every tourist who has been to France has ever visited Paris.
h. Every tourist who has ever been to France has visited Paris.

(13)a-b are the initial motivation for positing that *ever* must stand in a close relation to a negative element (a plausible assumption, advocated in syntax, is that *ever* must be "c-commanded by" a negative element; or to use terms more common in logic, it must be 'in the scope of' a negative element). This hypothesis gains further support from the deviance of (13)c-d and the acceptability of (13)e-f. But the contrast in (13)g-h comes as a surprise: no negative element appears in the sentence, and yet *ever* is licensed when it is embedded in the nominal argument of *every*, but not in its verbal argument. Why? The answer, due to Ladusaw and Fauconnier, is that the constraint on the distribution of *ever* is semantic in nature: *ever* is acceptable just in case it appears in an environment which is *semantically* negative. Semantically negative environments are defined in terms of entailment (for this reason, they are also called 'downward-entailing'): if John doesn't have a property P, a fortiori he doesn't a stronger property P'; for instance, if John hasn't been to France, *a fortiori* he hasn't been to Southern France. By contrast, if John has been to France, it doesn't follow that he has been to Southern France. It can be checked that the data in (13)a-f follow from this characterization. But now the facts in (13)g-h follow as well, because *every* creates a negative environment in its nominal but not in its verbal argument. This can be seen by observing that *every tourist who has been to France has visited Paris* entails that *every tourist who has been to Southern France has visited Paris*. By contrast, the same sentence does not entail that *every tourist who has been to France has visited the south of Paris*. Thus the licensing of Negative Polarity Items can thus fruitfully be stated in semantic rather than syntactic terms. Of course this still fails to explain *why* some words should be sensitive to this particular semantic property; this is still a topic of ongoing research.

In the case of Negative Polarity Items, the theory could have been developed entirely in terms of entailment. In other cases, however, generalized quantifier theory is essential to provide adequate generalizations. This is the case of another puzzle, which concerns the surprising patterns of acceptability produced by the existential *there*-construction:

(14) a. *There is every problem.
b. There is a problem.
c. There is no problem.
d. *There are most problems.
e. There are at least three problems.
f. There are exactly five problems.

A highly successful account of this distribution relies on the hypothesis that the *there* construction is only acceptable when the determiner that comes at the tail of the construction is 'symmetric' with respect to its nominal and verbal arguments. Let us consider the determiner *a*: *A professor smokes* has the same truth conditions as *A smoker is a professor*. Similarly for *No professor smokes*, which means the same thing as *No smoker is a professor*. We say that the determiners *a* and *no* are symmetric because their nominal and verbal arguments can be reversed with no truth-conditional change. By contrast, *Every professor smokes* doesn't mean the same thing at all as *Every smoker is a professor*. By going back to the lexical rules in (12), it can be checked that *a*, *no*, *at least three, less than seven* and *exactly five* are the only determiners in the list that are symmetric, in the sense that their two arguments can be reversed without change (this can be ascertained by observing that in each case they only make claims about the size of PROFESSOR ∩ SMOKE, which is of course the same thing as SMOKE ∩ PROFESSOR). This generalization nicely accounts for our data, and here too it is essential that grammatical constraints can be stated in purely semantic terms.

It should be added, however, that the theory of generalized quantifiers as defined only treats part of the logical complexities of natural language. A very rich domain is offered by the analysis of *plurals*, which give rise to numerous problems that are the object of intense contemporary research[2].

2.2 Pronouns and Binding

Let us turn to pronouns. While their analysis is still a topic of considerable debate, the theory that serves as a focal point treats pronouns as variables in predicate logic: pronouns that are 'free' (i.e. do not have an antecedent) are a sort of 'temporary proper names', whose denotation is provided by an assignment function. Technically, the semantic rules we posited earlier are now relativized to an assignment function s, which assigns objects to variables x_1, x_2, x_3, etc. And we add a special rule for pronouns, which says that the denotation of a pronoun pro_i carrying an index i is whatever the assignment function s assigns to x_i (for words which are not pronouns, the rules we posited earlier in (6) and (7) remain unchanged, except that for uniformity the superscript *s* is added everywhere):

(15) $[[pro_i]]^s = s(x_i)$

Our theory is still insufficient, however, because there are numerous constructions in which pronouns have *variable reference*, and do not denote just one given individual. This is for instance the case in the sentences *Every professor admires himself* or *Every professor likes people who admire him*, where *himself* must and *him* can be construed as having 'every professor' as its antecedent; we say in this case that they are 'bound' (which is the opposite of being 'free'). For simplicity, we stick to the first example (but the second example shows that the difficulty is not limited to reflexive pronouns). The solution is to take the quantifier to be

[2] If 'generalized quantifiers' can be seen as a generalization of the quantifiers of first-order logic, plurals are for their part very close to the second-order quantifiers of second-order logic (which also owes much to the work of Frege).

responsible for the formation of a complex predicate, *λi i admires himself$_i$*, which can be paraphrased as: *is an i such that i admires i.* Thus the sentence in (16)a is taken to have the structure in (16)b, where t_i is a pronoun-like element (called in syntax the 'trace' of the quantifier) and λi is called a *lambda-abstractor*, whose purpose is to form a complex predicate:

(16) a. Every professor admires himself
b. [Every professor] [λi t_i admires himself$_i$]

Intuitively, then, (16)a means that every professor has the property of *being an individual i such that i admires i.* All we need to do to incorporate this view into the analysis of quantifiers developed above is to ensure that the expression *λi t_i admires himself$_i$* has the semantic type of an intransitive predicate, i.e. <e, t>. This is achieved by defining a special rule that guarantees that *λi t_i admires himself$_i$* denotes that function which associates to any individual d the value 1 just in case d admires d. Technically, we define f in such a way that for every object x, f(x) has the value of *t_i admires himself$_i$* evaluated under a modified assignment function that assigns *x* to index i. If we write s[i → x] for an assignment function that fully agrees with s, except that it assigns x to i, the rule can be defined as in (17):

(17) $[[\ \lambda i\ F\ ]]^{s}$= that function f of type <e, t> such that $f(x) = [[\ F\ ]]^{s[i \rightarrow x]}$

2.3 Syntax-Semantics Interface

The semantic analysis we have sketched is intimately related to questions that concern the syntax/semantics interface, i.e. the delineation of the precise boundary between syntax and semantics.

First, syntacticians have long known that the analysis of pronouns must be constrained. Thus *John likes him* cannot mean that John likes John; but nothing in what we said precludes a situation in which the sentence *John likes him$_i$* is interpreted under an assignment function s which assigns John to the pronoun *him$_i$*. The problem can be solved in two ways: by adding syntactic constraints on the distribution of indices, so that *John likes him$_i$* comes out as *syntactically* ill-formed under certain conditions (this is the line followed by Chomsky in his 'Binding Theory'); alternatively, we could revise the semantics so as to predict that such an interpretation cannot be obtained in the first place. The debate between the two approaches, which should be settled on empirical grounds, is the object of ongoing research on the syntax/semantics interface.

Second, the semantic analysis of quantifiers interacts in interesting ways with sophisticated questions of syntax. Quantifiers often give rise to ambiguities that appear to be structural, i.e. to be due to the structure of the sentences at hand, as seen in (18):

(18) a. A doctor will interview every new patient.
b. [a doctor] λi [every new patient] λk t_i will interview t_k
c. [every new patient] λk [a doctor] λi t_i will interview t_k

(18)a can be understood to make the strong claim that there is some doctor that will interview every new patient, as is represented in (18)b; or it can be understood to make the weaker claim that for every new patient, there is a (possibly different) doctor that will talk to him, as is represented in (18)c. Why is there such an ambiguity? A bold view would be that the sentence really *does* have two possible syntactic representations, which literally correspond to (18)b and (18)c. Precisely this claim has been made in studies of 'Logical Form' within generative syntax. The view is emphatically not that the *only* way to account for the ambiguity is to posit that quantifiers appear in a position different from the one in which they are pronounced; such a view is certainly incorrect, as there are sophisticated semantic proposals that predict an ambiguity *without* resorting to such abstract levels of syntactic representation (the debate is still entirely open). Rather, the claim is that there is *independent* evidence for operations of movement that predict the right data when they are extended to quantifiers.

In the case at hand, the argument is in two steps. First, it has been argued in syntactic theory that certain expressions, such as interrogative words, are generated in a certain position but then 'move' to the location in they are pronounced, leaving behind a 'trace' in their original position (the expression and the trace are co-indexed to indicate in complex examples which trace corresponds to which expression):

(19) a. [Which patients]$_i$ will a doctor interview t_i?
b. [Which patients]$_i$ will a doctor try to assist t_i personally ? (Reinhart 1998)

Importantly, this movement is not possible out of all positions; certain syntactic configurations are 'islands' because interrogative words cannot move out of them. For instance, an interrogative word cannot be moved out of a complex Noun Phrase (*the possibility that* ...), as is illustrated by (20)a; similarly, interrogative words cannot move out of an *if*-clause, as shown in (20)b:

(20) a. *Which patients will a doctor examine [the possibility [that we give t_i a tranquilizer]]?
b. *Which patients should a doctor worry [if we sedate t_i]? (Reinhart 1998)

Now the crucial observation is that *if* we posit that the ambiguity observed in (18) is the result of an invisible (or 'covert') movement operation, which takes place *after* a sentence is pronounced rather than before, we predict, correctly, that some readings should disappear when one of the quantifiers is embedded with a syntactic island. This appears to be correct:

(21) a. A doctor will interview every new patient
Ok Reading 1: [a doctor]$_1$ [every patient]$_2$ t_1 will interview t_2
Ok Reading 2: [every patient]$_2$ [a doctor]$_1$ t_1 will interview t_2
b. A doctor will try to assist every new patient personally
Ok Reading 1: [a doctor]$_1$ [every patient]$_2$ t_1 will try to assist t_2 personally
Ok Reading 2: [every patient]$_2$ [a doctor]$_1$ t_1 will personally t_2 personally
a'. A doctor will examine the possibility that we give every new patient a tranquilizer
Ok Reading 1: [a doctor]$_1$ t_1 will examine the possibility that [every patient]$_2$ we give t_2 a tranquilizer.
* Reading 2: [every patient]$_2$ [a doctor]$_1$ t_1 will examine the possibility that we give t_2 a tranquilizer.
b'. A doctor should worry if we sedate every new patient
Ok Reading 1: [a doctor]$_1$ t_1 should worry [if [every patient]$_2$ we sedate t_2]
* Reading 2: [every patient]$_2$ [a doctor]$_1$ t_1 should worry [if we sedate t_2] (Reinhart 1998)

These data can then be explained with minimal semantic effort if we recast the relations between traces and quantifiers in terms of the creation of complex predicates illustrated in (16), so that the syntactician's representations in (21)a are minimally revised to look more like (18) (alternatively, slightly different semantic rules may be posited to apply directly to (21)a). In this way (which is just one example of one possible explanation), syntactic and semantic considerations conspire to predict intricate patterns of interpretation[3].

[3] The analysis of Logical Form in syntax interacts with typological considerations. According to syntactic theory, the fact that certain words move 'overtly' vs. 'covertly' is an arbitrary property of a particular language or construction. So there should in principle be languages in which interrogative words move covertly - this has been the argued to the case of Japanese and Chinese. Similarly, there should be languages in which quantifiers move overtly - this has been claimed to be the case of Hungarian.

3 Beyond Individuals: Contexts, Times, Possible Worlds, Events

Beyond individuals, several other types of objects must be integrated into semantic theory if it is to have any plausibility. Which types of objects must be posited is a subject of debate, but we will briefly consider contexts, times, possible worlds, and events.

3.1 Contexts

Expressions such as *I* (or *you*), *here* and *now* denote individuals, locations and times which depend on the context of utterance: *I* uttered by John does not refer to the same individual as *I* uttered by Mary. Since it does not seem that John and Mary speak different languages, it is useful to relativize the interpretation of sentences not just to an assignment function (which was seen to be useful for third person pronouns), but also to a context parameter, written as *c* in what follows. All of our earlier rules can be preserved in this enriched framework, but we can also define rules such as the following:

(22) $[[\mathrm{I}]]^{c,\,s} =$ the speaker of c
$[[\mathrm{you}]]^{c,\,s} =$ the addressee of c

The context parameter turns out to have many applications, among others in tense semantics[4].

3.2 Times

In order to deal with times, one needs some account of tense. Intuitively, *John left* is true when uttered in a context c just in case there is some moment before the time of *c* at which John left. Similarly, *John will leave* is true in *c* just in case there is a moment after the time of *c* at which John leaves.

In the tradition of tense logic, it was thought that linguistic reference to times crucially differs from reference to individuals in that the former is strictly less expressive than the latter. Specifically, it was thought that natural language does not have pronouns that refer to moments, and quantifiers that can bind them, but that it can only make use of 'operators' that do both things at once - which makes the system less flexible and expressive than reference to individuals. It was further thought that the present tense is just the absence of a past or future operator. This lead to an analysis in which a time parameter *t* is added to the context parameter *c* and to the assignment function *s*. A clause that has no tense is evaluated with respect to *t*, which of course means that the interpretation of *all* expressions must be similarly relativized to times (to be concrete, the verb *smokes* will now denote different functions at different moments, because who the smokers are typically changes over time). As for a clause of the form *PAST S* or *FUT S*, where *PAST* and *FUT* are past and future tense operators, they can be interpreted using the rules in (23):

(23) a. $[[\mathrm{PAST\ F}]]^{c,\,s,\,t} = 1$ iff for some time t' before t, $[[\mathrm{F}]]^{c,\,s,\,t'} = 1$
b. $[[\mathrm{FUT\ F}]]^{c,\,s,\,t} = 1$ iff for some time t' after t, $[[\mathrm{F}]]^{c,\,s,\,t'} = 1$

It is straightforward to apply this analysis to the earlier example *John left*, once we specify that the initial value of t is c_t, the time of the context (we follow syntacticians in taking the tense to occur at Logical Form in a position which is to the left of the rest of the sentence):

---

[4] It was traditionally thought that the context parameter differs from other parameters of evaluation in that it remains fixed throughout the evaluation of a sentence (Kaplan 1989). But investigation of other languages than English has recently lead to a re-examination of this assumption (several researchers now have argued that some verbs of speech or thought, such as *say* or *believe*, can in some languages manipulate the context parameter).

(24) [[PAST he_i leave $]]^{c, s, c_t}$ = 1 iff for some time t' before c_t, [[he_i leave $]]^{c, s, t'}$ = 1, iff for some time t' before the time c_t of c the individual denoted by i is leaving at t'

This analysis is simple and attractive, and it can be refined to handle more sophisticated constructions, for instance complex tenses (e.g. the pluperfect), as was done in Reichenbach 1947. But a major finding of contemporary semantics, originally due to Partee 1973, is that this view is largely mistaken. Partee's main insight is that temporal semantics has to a large extent the same resources as individual semantics: tenses often behave like time-denoting pronouns, and they may be bound by quantifiers (which are often unpronounced). There are two important arguments for this analysis: first, the time argument of verbs often behaves like a pronoun; second, nominals, which are semantically predicates, also have time arguments that display a pronominal behavior.

Let us start with Partee's argument that tense can sometimes be read as a pronoun. Suppose that a well-known and elderly character is in discussion with his editor who wants to put his picture on his latest book. Looking at the picture, he utters (25)a:

(25) a. I wasn't young.
b. PAST not be-young
c. not PAST be-young

The analysis offered by the modal semantics we posited in (23) is inadequate to capture the intended truth conditions, because all it can offer is (25)b or (25)c (this is on the assumption that temporal operators, like quantifiers, can move 'covertly'). But (25)a simply asserts that there was some past moment at which the well-known character was not young, which is not informative in the case at hand; for its part, (25)b asserts that there was no moment in the past at which the character was young, which is certainly false. Neither of those readings is what the character has in mind; rather, he wishes to convey that *at the time made salient by the photograph*, he wasn't young (and he might want to imply that this choice is not optimal). The desired reading is easily obtained by treating the past tense as a time-denoting pronoun, which in this case gets its denotation from an assignment function (duly upgraded so as to assign a value not just to individual but also to time variables). This leads to the representation in (26), where the past tense contributes a time variable t_1 (in a more complete treatment, the tense features themselves would contribute a constraint - often treated as a *presupposition* - on the value of that variable):

(26) not t_1 I be young

Partee shows in detail that several other uses of pronouns have counterparts in the temporal domain as well, but for simplicity we disregard this part of the argument.

Interestingly, a related point was made by Enç 1987 about nominals, which she argued must be endowed with time variables that are allowed to refer autonomously. A famous example involves the noun *fugitives* in (27)a:

(27) a. The fugitives are in jail.
b. t_0 the [t_1 fugitives] be-in-jail

Enç's point is that this sentence does not mean that the people who are *currently* fugitive are now in jail, as would be predicted by a simple tense logi; this only has a contradictory reading. Rather, the sentence is understood to mean that *the individuals who were fugitives at a salient time t_1* are now in jail. In our revised view of tense, the present tense will itself contribute a time variable t_0, which must be constrained (by a different part of the theory) to denote the time of utterance. Together, t_0 and t_1 give rise to the desired reading, in which *the individuals that were fugitives at time t_1 are in jail at time t_0.*

By separating the role played by time quantifiers from that of time variables, this analysis endows natural language with more flexibility than is afforded by tense logic. This has important consequences for the syntax/semantics interface, because we obtain in this way a variety of readings that would be very hard to get with basic temporal operators. A simple example is provided in (28) (more sophisticated examples are offered in Cresswell 1990):

(28) a. Some day, all of Dominique's students will be on the Editorial Board of *Linguistic Inquiry* (and he will rule syntax!)
b. Wrong Analysis 1: [every student] $\lambda 1$ FUT t_1 be-on-the-EB
c. Wrong Analysis 2: FUT [every student] $\lambda 1$ t_1 be-on-the-EB

The intended reading is one on which there is some day D in the future such that at D all of Dominique's current students are on the Editorial Board at D (with the addition of *it will happen that*, this may be the only available reading). Within a framework that only countenances temporal operators, this gives rise to a 'scope paradox':
-For the truth conditions to come out right, the quantifier *all of Dominique's students* must be in the scope of the time operator *some day*;
-But this has the consequence that *student* is evaluated with respect to a non-present moment - despite the fact that on the intended reading *students* means: *current students.*
The paradox can be solved if we posit time variables, as shown by the representation in (29)a, paraphrased as in (29)b:

(29) a. [some i_1: $i_0 < i_1$] [every [i_0 student]] 1 [i_1 [t_1 be-on-the-EB]]
b. There is some future moment i_1 such that every individual who is a student at the current moment t_0 is on the Editorial Board at i_1

Thus time pronouns are both semantically and syntactically essential for a proper understanding of time dependency in language.

3.3 Possible Worlds

Just like times are used to provide a semantics for tense, possible worlds have often been used to provide a semantics for mood and modals, as in *John might come*, or *If John were here, Mary would be happy.* Possible worlds are a topic of controversy in metaphysics, but they have proven helpful to define a semantics for modal logic, which is concerned with reasoning about possibilities; and in turn, modal logic proved initially useful to analyze natural language. In essence, one can think of a possible world as an entity that fully determines the way things are *or could have been.*

Equipped with such a notion, we can further relativize the tense logic we introduced in the previous section to a world parameter - with the convention that the initial value of the world parameter is just the world of the context. This make it possible to define interpretive rules for sentences of the form *may S* or *must S* (we follow syntacticians in taking the modal to occur at Logical Form in a position which is to the left of the rest of the sentence):

(30) a. $[[\text{may S}]]^{c, s, t, w} = 1$ iff for some world w'accessible from w, $[[\text{S}]]^{c, s, t, w'} = 1$
b. $[[\text{must S}]]^{c, s, t, w} = 1$ iff for every world w'accessible from w, $[[\text{S}]]^{c, t, w'} = 1$

Intuitively, one reading of *John may be sick* is that there exists a state of affairs (= a possible world) compatible with what we know in which John is sick. By contrast, *John must be sick* makes the stronger claim that *every* possible world compatible with what we know is one in which John is sick. These truth conditions are easily derived if we take *w' is accessible from w* to mean: *w' is compatible with what is known in w* [at the time of evaluation]. With the further specification that the initial values of t and w are the time and world of the context c, respectively, we can derive the desired truth condition:

(31) [[must he_i be-sick]] $^{c, s, t, w}$ = 1 iff for every world w' compatible with what is known in w at t, [[he_i leave]] $^{c, s, t, w'}$ = 1, iff for every world w' compatible with what is known in the world and at the time of the context c, the individual denoted by i is sick in w' at t.

This analysis turns out to have considerable benefits when we consider different *readings* of modals. *John must be sick* means that for all we know he is sick. *John must work* does no mean that in every world compatible with what we know he works, but rather that in every world compatible with moral norms (or some related notion), he works. In these two readings, we see that something remains constant: the quantificational force of *must*, which makes a claim about *every* possible world with certain properties. What changes, on the other hand, is the domain of worlds which is quantified over: in the first case it is the worlds *compatible with what is known*, in the second it is the worlds *compatible with a norm*. The theory can account for this variation by allowing the precise meaning of *accessible* in (30) to be determined by the discourse situation.

Although this analysis has proven quite powerful, it is generally thought that the same arguments that show that tense talk in natural language is richer than tense logic carry over to the world domain. Semanticists now generally work with systems that include explicit world-denoting pronouns, which may be free, or bound by (pronounced or unpronounced) world quantifiers. In fact, almost all of the data we discussed with respect to the tense domain have a counterpart in the world domain, which suggests that the same measures should indeed be applied in both cases.

3.4 Events and Beyond

Finally, many theories make use of other kinds of objects to handle further constructions. Some of these objects may be used *in lieu* of times and possible worlds, and there is currently no consensus on the 'right' ontology (and to some extent one can 'translate' among approaches that posit different ontologies).

Events are a particularly useful category, which was initially posited by the philosopher Davidson to account for the logic of adverbial modification, which is not easily handled in the simple analysis of verbs (analyzed as expressions of type <e, t>) which was sketched above. His basic observation was that *Brutus stabbed Caesar at midnight with a knife* entails that *Brutus stabbed Caesar* and that *Brutus stabbed Caesar with a knife,* although the conjunction of the latter two sentences does not suffice to entail the first because two different stabbings may have occurred (say, with one taking place at midnight and the other being performed with a knife). On the other hand, this asymmetric pattern of entailment is easily derived if each sentence is analyzed as an existential quantification over events: from *∃e (stabbing(e) & agent(e)=Brutus & at_midnight(e) & with_a_knife(e))*, it follows straightforwardly that *∃e (stabbing(e) & agent(e)=Brutus & at_midnight(e))* and that *∃e (stabbing(e) & agent(e)=Brutus & with_a_knife(e))*, although the conjunction of the last two formulas does not suffice to entail the first one - which is the desired result. Importantly, neither times nor worlds are sufficiently fine-grained to allow for such an analysis.

Although adverbial modification was the initial motivation for positing events, these have turned out to be extremely useful in the study of verbal aspect, which lead to the discovery of surprising analogies between the nominal and the verbal domain. In a nutshell, it was observed that the distinction between so-called *telic* verbs *(die, build a house)* and *atelic verbs* (*be happy*, *run)* can be related to the count/mass distinction in the nominal domain. Classically, telic verbs are compatible with the adverbial *in an hour* and not with the adverbial *for an hour (John ran/was happy for two hours / *in two hours),* whereas atelic verbs display the opposite pattern (*John died / built a house *for two hours / in two hours).* Researchers found that this distinction was connected to a logical property reminiscent of the nominal domain (Bach 1986). *Atelic verbs*, like mass terms (e.g. *water)*, satisfy a property of *cumulative reference:* put together, two events that satisfy the predicate *running* still satisfy the same predicate, just like two samples of *water* that have been put together still count as being *water*. By contrast, *telic verbs*, like *count terms* (e.g. *chair*), fail the test: put together, two events of *building a house* may in general amount to an event of *building two houses* but not of *building a house*; and similarly two chairs put together do not fall under the predicate *chair,* but rather under the predicate *chairs*. The details of the analysis are still the object of lively debate (Rothstein 2004), but there is general agreement that some systematic semantic correspondence between the nominal and the verbal domain is indeed real, and that it can be accounted for in a framework that countenances events.

With events in hand, we may endeavor to revisit the other types of objects we postulated, some of which may now become dispensable (if events are strictly more fine-grained than either times or possible worlds, one may for instance try to define all semantic rules in terms of just individuals and events). But it is very likely that the list is by no means closed. Researchers working on adjectives have posited a rich ontology of *degrees*; those working on locatives have posited *locations*; and those working on manner adverbs have sometimes posited - well, manners. In each case the questions become interesting when one gets specific about the syntax and semantics; and issues we raised about times and possible worlds re-emerge in these new domains: how is reference to these objects effected? and how do the details of the syntax/semantics interface work?

4 The Semantics / Pragmatics Interface

4.1 Implicatures

Even in the simplest cases, the information conveyed by a sentence has two sources: its truth conditions, given by the semantics; but also additional inferences that we typically make by reasoning on the speaker's motives for uttering one sentence rather than another. The latter information is the realm of *pragmatics*.

In a famous example, the British philosopher Paul Grice observed that if I write in a letter of recommendation for my student Bill that *he is always on time and is hard-working*, the recipient will likely infer that Bill should not be hired - but not because any of the qualities I attributed to him was negative. Rather, the fact that I *failed* to mention more directly relevant qualities - such as his intellect, suitability for the job, etc. - suggests that I think he lacks those. Grice called 'implicatures' such inferences, which are derived from the assumption that the speaker follows certain rules of cooperative communication. Although in the present case the boundary between semantics and pragmatics is clear enough, in other cases it is the object of lively debates.

As in other domains of cognitive science, whenever one is interested in the boundary between two modules, one can bring different kinds of evidence to bear on the cartography that one seeks to establish. Distinct modules may be expected to give rise to different rules, but also to be processed differently in real time, to develop differently in language acquisition, to be realized differently in the brain, and to be affected differently in patients that have brain lesions. Precisely this convergence of approaches is beginning to take shape in the study of the semantics/pragmatics interface. Let us take the example of the little word *or*. In logic, disjunction is, by convention, *inclusive*: *p or q* is true just in case p or q *or both* are true. But in natural language, one often observes what seem to be *exclusive* readings. If I say that *I'll invite Mary or Ann*, the addressee will infer that not both of them will be invited. One possibility is that contrary to what is posited in logic, the disjunction of natural language is exclusive: *p or q* is true just in case p or q is true, *but not both*. However an alternative theory posits that the exclusive inference is an implicature of a special sort (called a 'scalar implicature' because it involves a comparison between different members of the scale <*or*, *and*>). Specifically, we start from the assumption that *or* is inclusive, but postulate that in simple cases the addressee reasons as follows:

(32) Scalar implicatures

a. <or, and> forms a scale: any use of *or* evokes a possible replacement with *and*, and vice versa.

b. The version of the sentence with *and* (*I'll invite Mary and Ann*) is more informative than the version with *or* (*I'll invite Mary or Ann*.)

c. Since the speaker is cooperative, if he had been in a position to use the more informative sentence was true, he would have done so. This suggests that he was not in a position to utter *I'll invite Mary and Ann* - possibly because he thinks that the conjunction is false.

The comparison between the two theories (exclusive *or* vs. inclusive *or* with implicatures) has yielded considerable evidence for the pragmatic analysis. The first observation is that the inference in question is defeasible - it is no contradiction to say that *I'll invite Mary or Ann - in fact I'll invite them both*. Any theory that countenances exclusive *or* has to posit that *or* can *also* be read as inclusive so as to account for this possibility. Thus an ambiguity must be posited. But even so, this analysis can be refuted.

-First, *or* is normally treated as inclusive (if uttered with a neutral intonation) in those cases in which the step in (32)b fails, for instance in semantically negative environments. *None of my friends will invite Mary or Ann* definitely rules out that any of my friends invites *both* Mary and Ann, which would be unexpected if *or* could be read exclusively. By contrast, the facts follow on the pragmatic theory: it is clear that if *none of my friends will invite Mary or Ann (or possibly both)*, then *a fortiori none of my friends will invite both*, which shows that the sentence with *or* is in this case *more* informative than the sentence with *and*. As a consequence, the use of *or* cannot give rise to an implicature.

-Second, this analysis makes predictions beyond the data that motivated it. As we just saw, if I say that *None of my friends will invite Mary and Ann*, I will have uttered the *less* informative of the two sentences under comparison. Following the logic of our earlier argument, this should give rise to a new implicature, namely that I was not in a position to say the more informative sentence, *None of my friends will invite Mary or Ann*. If the reason for this is that I take this sentence to be false, we get the pragmatic inference that *At least some of my friends will invite Mary or Ann*, which appears to be just right.

-Studies of language acquisition have shown that children acquire implicatures much later than the basic logical properties of connectives such as *and* and *not*. This is fully compatible with the view that implicatures are a different, and possibly more complex, inference than semantic entailments.

-Studies of language processing have shown that subjects who do compute the implicature take more time than those who don't - which appears to confirm the view that an additional step of reasoning is necessary to obtain the implicature, as suggested by the theory.

The implicature-based analysis has been extended to numerous other phenomena: in simple clauses, *some* and *most* are both taken to implicate *not all*; *might* is taken to implicate *not must*; *good* is taken to implicate *not excellent*; etc.

It should be noted that the precise way in which implicatures are computed has been the object of renewed debate in recent years. Various researchers have argued that *despite* the successes of purely pragmatic accounts, scalar implicatures should be seen as computed in tandem with the syntax and semantics. Other researchers have sought to defend a more traditional pragmatic analysis. The debate is currently open.

4.2 Presuppositions

Presuppositions are another domain in which the boundary between semantics and pragmatics is of considerable theoretical interest. Presuppositions are initially characterized by two properties: first, a simple clause *S* with presupposition *P* is odd (neither true nor false) if *P* is false; second, presuppositions give rise to inferences that are inherited by complex sentences differently from normal entailments, or for that matter from implicatures. For instance, presuppositions are preserved when they appear in questions or under negations, which is definitely not the case of entailments: *John is English* entails that *John is European*, but of course *John isn't English* or *Is John English?* have no such entailment. In (33), we see that things are different with presuppositions triggered by *the* (= presupposition that Syldavia has a king), *know* (= presupposition that John is incompetent), and the 'cleft' construction *it is ... who...* (= presupposition that someone stole your watch):

(33) a. John doesn't like the king of Syldavia. a'. Does John like the king of Syldavia?
=> Syldavia has a king.
b. John doesn't know that he is incompetent. b'. Does John know that he is incompetent?
=> John is incompetent.
c. It is not John who stole your watch. c'. Is it John who stole your watch?
=> Someone stole your watch.

Even more characteristically, presuppositions give rise to universal inferences when they are embedded under the determiner *no* or *none*; for instance, *None of these ten students knows that he is incompetent* leads to the strong inference that *each of these ten students is incompetent* - a pattern which is entirely different from that of entailments or implicatures.

There are two questions that can be asked about presuppositions: first, how are they generated to begin with? second, how are the presuppositions of elementary clauses transmitted to complex sentences? The first question is still open, but the second question has been the object of intense scrutiny. The simplest theory would be that a presupposition is satisfied just in case it follows from what the speech act participants take for granted in the context of the conversation. This would predict that presuppositions are *always* inherited by complex sentences. Sometimes this is the case; thus the conjunction *John is realistic and he knows that he is incompetent* presupposes that John is indeed incompetent. But the apparently analogous sentence *John is incompetent and he knows that he is* presupposes no such thing; rather, it asserts it (the same 'disappearance' phenomenon occurs in conditionals: compare *If John is realistic, he knows that he is incompetent* - which implies that John is incompetent - with *If John is incompetent, he knows that he is* - which implies no such thing).

A highly influential proposal, due to the philosopher Stalnaker, suggests that the basic account is almost correct, but that there are more contexts that meets the eye: the presupposition of *he knows that he is (incompetent)* is evaluated, not with respect to the initial context, but rather with respect to the modified context obtained after the speech act participants have revised their beliefs on the basis of the first conjunct. Since in this case the first conjunct entails that John is incompetent, by construction the context of the second conjunct will entail it as well - no matter what the initial context was; this means that the sentence *as a whole* will not presuppose anything. This account proved extremely influential, but to be generalized to other connectives and operators it required a rather radical departure from standard assumptions. In particular, it was assumed in 'dynamic semantics' that the very meaning of every expression is to modify what is taken for granted in a conversation, which led to a radical revision of the semantic framework: instead of being treated in terms of truth conditions, meanings came to be seen as an *instruction to modify beliefs*. This analysis led to a significant modification of the foundations of semantics; but whether this 'dynamic turn' was justified is still the object of lively debates.

Suggestions for further reading

-A very brief introduction to formal semantics is offered in Larson, R. (1995) Semantics. In Osherson, D., Gleitman, L., and Liberman, M. (eds) *An Invitation to Cognitive Science. Volume 1: Language*. MIT Press.

-Two standard textbooks are Heim, I. and Kratzer, A. (1998) *Semantics in Generative Grammar*, Blackwell; and Chierchia, G. and McConnell-Ginet, S. (1990) *Meaning and Grammar*, MIT Press. A slightly different notational framework is developed in Larson, R. and Segal, G. (1995) *Knowledge of Meaning*, MIT Press.

-A selection of the some of the most important papers in formal semantics are found in Portner, P. and Partee, B. (2002) *Formal Semantics: The Essential Readings*, Blackwell.

-For a collection of survey articles in English, one can read Lappin, S. (ed) (1996) *The Handbook of Contemporary Semantic Theory*, Blackwell. A much longer collection with articles in English and German is offered in Stechow, A. von and Wunderlich, D (eds) (1991), *Handbuch Semantik/Handbook Semantics*, de Gruyter, 1991