IV Aspects of meaning: Semantics and pragmatics

In the two previous segments of this series we have come across several references to meaning. It is high time now to look in more detail at the two subdisciplines of linguistics that look at meaning: semantics and pragmatics. SEMANTICS (from Greek *semaino* (σημαίνω) 'to indicate') is the branch of linguistic enquiry that investigates those aspects of meaning that are conditioned by syntactic structure. This emphasis on structure-based meaning entails that semantics is not concerned with lexical meaning ('what is the meaning of the word *tree*?'), nor with logic or pragmatics. PRAGMATICS (ultimately but uninformatively from Greek *pragma* (πραγμα) 'thing done, fact') studies aspects of meaning which are determined by and in turn influence discourse: it examines what makes a particular utterance a felicitous or infelicitous contribution to a discourse. The dividing line between semantics and pragmatics is not always clear-cut — and some terminology is used in both disciplines (see the discussion of implicatures in §IV.6, for instance). But as a general rule of thumb, *(a)* aspects of the meaning of a sentence conditioned by the structure of the sentence are the province of semantics, and *(b)* aspects of the meaning of an utterance conditioned by the felicity of its use in the context of a discourse are the topic of pragmatics.

IV.1 The subject–predicate distinction, and idiom formation

In our discussion of the *easy to please* and *eager to please* constructions in segment III, we discovered that there are differences between the two constructions that are not directly (or even at all) related to the semantic function of the noun phrase *John*. Syntax and semantics are autonomous subdisciplines of grammar, each with their own responsibilities and rules. But we also discovered that a number of differential properties of these constructions can be related to the fact that in *John is easy to please*, *John* is the semantic object of *please* whereas in *John is eager to please*, *John* serves the semantic role of subject of *please* (and of *eager*).

It is an unfortunate fact that throughout the history of the study of language in the Western tradition (dating back to the philosophers of ancient Greece, esp. Aristotle), there has been a lot of confusion about the terms 'SUBJECT' and 'PREDICATE'. These terms are often used at three very different levels of analysis. For the term 'subject', this is elucidated below:

(i) GRAMMATICAL or SYNTACTIC SUBJECT
the occupant of the structural subject position of the sentence

(ii) SEMANTIC SUBJECT
often equivalent to the 'agent' or 'do-er'

(*iii*) LOGICAL or PRAGMATIC SUBJECT
the topic ('what the sentence is about')

The subject/predicate distinction has both syntactic and interpretive sides. These sometimes converge. In a typical active sentence such as *John kissed Mary*, the occupant of the structural subject position of the clause is the semantic and pragmatic subject of that clause. But the passive counterpart, *Mary was kissed by John*, only the structural subject and pragmatic subject are identical; the semantic subject of the clause, if realised at all, sits in a *by*-phrase.

In sentences such as *nobody kissed Mary*, where the semantic subject occupies the structural subject position, we realise immediately that it would be absurd to call *nobody* (the syntactic and semantic subject) the pragmatic subject of the sentence. The sentence *nobody kissed Mary* cannot be about the referent of *nobody* because *nobody* does not refer to anybody.

So we see that there are reasons for wanting to make careful distinctions between syntactic, semantic and pragmatic subjects. For the predicate, too, such distinctions are called for. Structurally, the predicate is a syntactic constituent separate from the subject. We see the syntactic independence of the predicate from its subject coming to light in sentences such as (1). Whereas the predicate usually follows its subject (and that is exactly where we find it in the *when*-clause of (1)), the *then*-clause shows that the predicate can be displaced into a syntactic position to the left of the subject.

(1) when I promise that I will do the dishes, then $[_{\text{Pred}}$ do the dishes], I_{Subj} will

But the predicate is also clearly a semantic entity that is autonomous *vis-à-vis* the subject. What shows this perhaps particularly clearly is the distribution of idioms. Consider the example in (2):

(2) $\text{John}_{\text{Subj}}$ $[_{\text{Pred}}$ kicked the bucket]
- a. 'John's lower extremity made violent contact with the bucket'
- b. 'John died'

This sentence can of course be interpreted literally, as paraphrased in (2a). But there is also an interpretation for the sentence according to which it is equivalent to (2b). In that second interpretation (called the idiomatic reading of the sentence), *kick the bucket* is a syntactic unit that gets a specialised semantics assigned to it — a semantics that cannot be derived from the meanings of the component parts of the expression.

The availability of an idiomatic interpretation for *kick the bucket* is entirely independent of the properties of the subject of predication. The subject can be *John*, or *Bill*, or any noun phrase whose referent is liable to face death at some point: see (3).

(3) $\{\text{Bill/Mary/the neighbour/the man next door/my boss/...}\}_{\text{Subj}}$ $[_{\text{Pred}}$ kicked the bucket]

But as soon as we replace the verb *kick* with an otherwise semantically equivalent expression (such as *strike with one's foot*, as in (4a)) or substitute for object *the bucket* a noun phrase headed by its synonym *pail* (as in (4b)), the idiomatic interpretation for the predicate immediately disappears: though the sentences in (4a,b) are perfectly grammatical, they cannot be interpreted as 'John died'.

(4)
- a. $\text{John}_{\text{Subj}}$ $[_{\text{Pred}}$ struck the bucket with his foot] ⇒ no idiomatic reading
- b. $\text{John}_{\text{Subj}}$ $[_{\text{Pred}}$ kicked the pail] ⇒ no idiomatic reading

It is relatively easy to form idioms from the combination of a verb and other material that occurs within the verbal predicate — such as an object (as in (3) or *to miss the boat* 'to lose an opportunity') or a combination of an object and a secondary predicate (as in *to paint the town red* 'to have a party'). By contrast, there are virtually no cases of idiomatically frozen expressions consisting of a subject and a verb, with a free choice of object. English has one such idiom: *A little birdie told me that ...* Here, it is specifically the combination of *a little birdie* (the subject) and the

verb (*tell*) that forms the idiom: the pronoun *me* can freely be replaced with some other (pro)noun, and the contents of the *that*-clause are entirely open to whatever one might wish to put in it. But as is so often the case, this really is the exception that proves the rule: the generalisation is that the predicate (the verb plus its object and possibly other predicate-internal material, such as secondary predicates) can readily be 'frozen' into an idiom whereas subject+verb idioms with an unfixed object are exceedingly rare.

From (1), we know that the predicate is a constituent. The verb does not form a constituent together with its subject to the exclusion of the object, however: (1′) is word salad.

(1′) *when I promise that I will do the dishes, then [I do], will the dishes

The fact that verb+object idioms are overwhelmingly more common than subject+verb idioms can be understood from this perspective. Semantic idioms must be syntactic constituents. The subject and the verb do not form a syntactic constituent, hence as a rule the subject and the verb cannot form an idiom together to the exclusion of the object. But because the verb and the object do form a syntactic constituent, it is possible to 'freeze' this constituent into a semantically specialised unit (an idiom) to the exclusion of the subject.

Syntax and semantics thus pattern together in certain ways. The semantics assigns an interpretation to the structures built in syntax. But these structures are built independently of semantic considerations, and do not use semantic notions as their guidelines. The syntax and semantics are autonomous components of linguistic analysis — autonomous but at the same time inextricably related to one another.

IV.2 Compositionality and 'expletives'

A central guideline in the construction of the meanings of sentences is the COMPOSITIONALITY PRINCIPLE, due to the German philosopher Gottlob Frege. The idea behind the compositionality principle is very simple: the meaning of a complex whole is a function of the meanings of its parts; the meanings of the parts compose to deliver the meaning of the whole. So once we know what *the*, *boy*, *girl*, *kiss* and *-ed* (past tense) mean, we should be able to figure out what *the boy kissed the girl* means. And indeed, this is how a language user generally goes about his or her business.

It is important to bear in mind, however, that not everything in language has meaning, and that the meaning of a syntactic construct is not always (straightforwardly) compositional. The latter we already discovered in the previous section when we talked about idioms. The idiomatic meaning of the sentence *John kicked the bucket* (i.e., 'John died') is not fully a function of the meanings of its component parts: while *John* and *-ed* contribute in the familiar way to the meaning of the sentence, the verb and its object combine non-compositionally to form the idiomatic VP predicate.

In *John kicked the bucket*, it may not be transparent what the meaning of the VP is; but we are sure that it does have a meaning, albeit a non-compositional one. But syntactic structures are frequently populated with elements that do not appear to have a contribution to make to the meaning of the sentence at all. Some functional/non-lexical categories have very little meaning. Thus, the complementiser *that*, which introduces finite subordinate clauses in English, does not seem to mean anything at all — indeed, we can usually forgo a complementiser altogether: (5) is grammatical with or without the complementiser *that* (hence the parentheses).

(5) John thinks (that) Mary is smart

Similarly, the use of a definite article in combination with a proper name, and the use of a pronoun in combination with a subordinate clause, both found in (dialects of) Hungarian (as in (6)), seem to make no obvious contribution to the way we interpret the sentences that contain these elements.

(6) a János azt hiszi hogy a Mari okos
the János it thinks that the Mari smart
'János thinks that Mari is smart'

But although it would be difficult to put one's finger on the meaning contributed by the pronoun *azt* 'it' in the Hungarian sentence in (6), it is striking that leaving this pronoun out would produce a sentence that, while certainly grammatical, is not exactly semantically equivalent to (6) as it stands. When one says *János hiszi hogy Mari okos*, one tends to convey that one János truly believes that the proposition 'Mari is smart' is true; but *János azt hiszi hogy Mari okos* does not have this 'truly believe the truth of *x*' interpretation. So it may not actually be the case, after all, that the use of the pronoun *azt* in (6) is genuinely meaningless.

For the pronoun *azt* in (6), the descriptive label 'expletive' or 'pleonastic' (both meaning 'filler') is frequently used. This terminology gives direct expression to the belief that the pronouns in question are used merely to fill a particular syntactic position. The element *there* seen in (7b) occupies the structural subject position apparently just because in (7b) (unlike in (7a)) the notional subject of the sentence (*firemen*) did not move into this position.

(7) a. firemen are available
b. there are firemen available

But as in the case of Hungarian *azt* just discussed, even though the 'expletive' itself means very little, its use can nonetheless have an effect on meaning. The sentences in (7a,b) are not perfectly semantically equivalent: (7a) supports two possible interpretations, while (7b) can have just one. The only reading available for (7b) is one that says that a bunch of firemen are at our disposal at our present location at the present point in time. The sentence in (7a) can express this, too; but in addition, it can also be used to make a statement about firemen in general: firemen have the general property (regardless of any particular location or point in time) of standing at the ready. A generic statement of this sort cannot be made with the aid of (7b). So although it would be difficult to say what the word *there* occupying the structural subject position means by itself, it does nonetheless seem to 'do something' to the meaning of the sentences in which it occurs. The meaning of (7b) might not be entirely compositional — but the big challenge that meaningless elements seemed to pose for the compositionality principle might not be quite so severe as it initially seemed.

Whether there exist truly meaningless elements in natural language is a question whose answer is to some extent in the eye of the beholder. It is unquestionably the case that elements can be found in grammatically well-formed sentences whose contribution to the semantics of these sentences (let alone their pragmatics) is very difficult to put one's finger on. But we may want to be careful not to immediately declare such elements devoid of meaning altogether: we may just not have studied them closely enough; a stronger microscope may be needed.

From this point onwards, we will set aside 'expletives' and the question of whether they do or do not make a contribution to meaning. We will concentrate instead on elements which do have meaning, and will try to determine how it is that they contribute to the meaning of the sentence or discourse that they find themselves in.

IV.3 Type theory

When we look at the sentence from the perspective of meaning, what is fundamental to it is that it expresses a PROPOSITION — a combination of a SUBJECT, a PREDICATE and TENSE to which we can attach a TRUTH VALUE ('1' for true statements, and '0' for false ones). In the theory of semantic enquiry that is called TYPE-THEORETIC SEMANTICS, the semantic type <t> ('truth value') for a particular sentence is derived from the meanings of its parts, in conformity with the compositionality principle. Let us illustrate this first for a simple sentence with an intransitive verb, such as the one in (8).

(8) John smokes

a. a sentence represents a statement that is either true ('1') or false ('0')
→ the sentence must be of semantic type <t>
b. *John* denotes an entity in the extralinguistic universe
→ its semantic type is <e> (where 'e' stands for 'entity')
c. the verb *smokes* constitutes the predicate of the sentence; that predicate combines with *John* to deliver the sentence as a whole, whose semantic type is <t>
→ the predicate must be of type <e,t> — a function of individuals to truth values
→ the intransitive verb *smokes* by itself constitutes the predicate, hence is of type <e,t>

The semantic type of the intransitive predicate *smokes*, viz., <e,t>, should be read as follows: 'the predicate wants something of type <e> to deliver something of type <t> (i.e., a sentence)'. All predicates with subjects of type <e> want to be of type <e,t>: after all, they all want to be able to combine with a subject of type <e> to deliver a truth value for the sentence. So for the transitive sentence in (9), which likewise has an entity of type <e> as its subject, the predicate will once again want to be of type <e,t>. But this time around, the predicate is not just the verb: the predicate contains the verb and its object. Since the object is of type <e>, and since the object and the verb combine to deliver a predicate of type <e,t>, this entails that the transitive verb *like* in this sentence must be represented as something that wants to combine with something of type <e> to yield something of type <e,t> — in other words, the transitive verb must be of type <e, <e,t>>.

(9) John likes Mary

a. a sentence represents a statement that is either true ('1') or false ('0')
→ the sentence must be of semantic type <t>
b. *John* and *Mary* denote individuals, entities in the extralinguistic universe
→ their semantic type is <e> (where 'e' stands for 'entity')
c. the verb *likes* combines with *Mary* to deliver the predicate of the sentence; that predicate in turn combines with *John* to deliver the sentence as a whole, whose semantic type is <t>
→ the predicate must be of type <e,t>
→ the transitive verb is the head of the predicate; it must be of type <e, <e,t>>

For any branching node in the syntactic structure, we compose the two daughters and derive the semantic type of the mother node. This is what COMPOSITIONALITY amounts to: the semantics of a complex whole is transparently composed of the meanings of its parts:

(10)

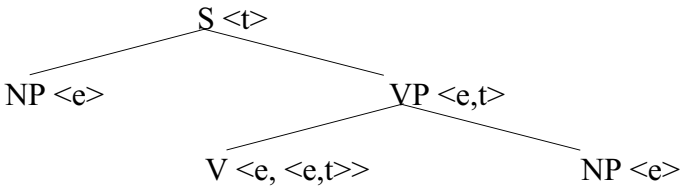

IV.4 Quantifiers

So far we have confined ourselves to sentences whose subjects and objects denote individuals in the extralinguistic universe. Noun phrases that denote individuals are called REFERENTIAL EXPRESSIONS. But not all noun phrases denote individuals. Natural languages contain a variety of QUANTIFICATIONAL noun phrases (or QUANTIFIER PHRASES) that have no (fixed) referent. When a speaker utters (11a), this speaker is not referring to any identifiable individual in the extralinguistic universe to whom the predicate *likes Mary* applies: indeed, the whole point of uttering (11a) is to convey that no such individual exists (in the speaker's mind). And when someone utters (11b), there is no clear and stable referent for the subject *everybody* in the extralinguistic world either: *everybody* might refer to everybody in the speaker's house, or everybody in Hungary, or everybody in Europe, or everybody in the world — only the context will (hopefully) tell. Whereas *John* and *Mary* refer directly to individuals, quantified expressions such as *nobody* and *everybody* never do — not just when they serve as subjects, as in (11a,b), but equally when they serve as objects, as in (11a′,b′).

(11) a. nobody likes Mary — a′. Mary likes nobody
b. everybody likes Mary — b′. Mary likes everybody

A quantified noun phrase such as *everybody* is of type <<e,t>, t>, a function of sets to truth values. It does not denote an entity (i.e., it is not of type <e>), but in combination with a predicate (of type <e, t>) it can deliver a truth condition. When a quantifier of type <<e,t>, t> serves as the object of a transitive verb (which we had previously identified as being of type <e, <e,t>>), as in (11a′,b′), we cannot compose the two to arrive at a predicate of type <e,t> that can subsequently combine with the subject of type <e>:

(12) *

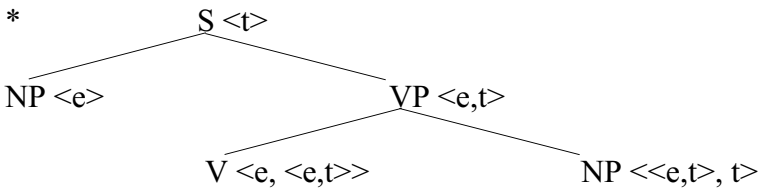

The verb in (12) says 'give me something that is of type <e> so that I can deliver a predicate of type <e,t>', but the object is not of the type that allows the verb to compose with it. So (12) is ill-formed. But the sentence it tries to represent is grammatical. What to do? There are three logical possibilities: *(i)* creating multiple lexical entries for verbs; *(ii)* postulating a semantic operation whereby the type of the verb can be changed; and *(iii)* postulating a syntactic operation affecting quantifiers.

Option *(i)* involves creating different lexical subtypes of transitive verbs, each of which is suitable for combination with a particular type of object. Concretely, transitive verbs that combine with a referential object are, as before, of type <e, <e,t>>; but for verbs that take a quantificational object, one could propose that they lexically belong to the type <<<e,t>, t>, <e,t>>. While an approach of this works, it amounts to listing every verb compatible both with referential and with quantificational objects (i.e., every transitive verb, in principle) twice in the lexicon — and not just the lexicon of English, but also the lexica of all other languages that allow transitive verbs to take quantificational objects. This is obviously not a particularly parsimonious strategy: it hugely increases the size of the lexica of the world's languages.

Instead of lexically listing two different types for each transitive verb, we could alternatively resort to option *(ii)*, listing each verb just once but then performing a semantic operation on the verb that 'shifts' its type, in the course of the derivation, to the one appropriate for combination with a quantificational object. In the particular case of a quantificational object of type <<e,t>, t>, we can shift the verb's type from <e, <e,t>> to the more complex <<<e,t>, t>, <e,t>>. After this TYPE SHIFTING operation applied to the transitive verb, the quantificational object can combine with the verb to form a predicate of type <e, t>.

Option *(iii)*, not surprisingly, is particularly popular in the literature on the syntax–semantics interface. On this approach, the transitive verb's type is left entirely untouched; instead, a syntactic operation is performed on the quantified object that removes it from the object position within the predicate, and leaves behind something of a type that an ordinary transitive verb of type <e, <e,t>> can compose with, i.e., a variable of type <e>. This syntactic operation targeting quantifier phrases is called QUANTIFIER RAISING (QR). It produces representations which are very similar to things which we know syntax can also produce overtly, with the aid of the rule of topicalisation. Thus, compare (13a) (the product of QR performed on (11b′), at the syntax–semantics interface) to (13b,c).

(13) a. $[_S$ everybody$_{<<e,t>,t>}$ $[_S$ Mary likes <e>]]
'for everybody, it is the case that Mary likes them'
b. $[_S$ (as for) John, $[_S$ Mary likes him]]
c. $[_S$ John, $[_S$ Mary likes ~~John~~]] (... but Bill, she doesn't)

QR is very useful not just in dealing with the problem posed by (12) but also in solving a puzzle posed by quantificational objects: SCOPE AMBIGUITY. Consider the sentence in (14). This sentence allows two rather different interpretations, paraphrased in (14a,b). In the first, *someone* scopes over *everyone*: there is a single individual, let's call him 'Bob', who is in a loving relationship with everyone in the universe of discourse. In the second reading, everybody in the discourse is also being loved by someone — but this time around, the responsibility for everyone being loved does not rest on the shoulders of just one specific individual: for everyone, there is someone who loves them, quite possibly a different lover for each lovee. On the surface, (14) is just one string of words; but this single string can give rise to two very different interpretations. Such ambiguity poses a problem for the COMPOSITIONALITY principle — but only if the string in (14) has just a single syntactic structure. If there are two different syntactic structures associable with this single string, and if these structures are such that they compositionally give rise to the meanings in (14a) and (14b), then the problem for the compositionality principle will disappear. QR is a useful tool with which this can be achieved. The structures given below (14a) and (14b) illustrate. These structures are the two LOGICAL FORMS (or LFs) for the sentence in (14)

(14) someone loves everyone

a. 'there is someone such that (s)he loves everyone'
[$_S$ someone$_{<<e,t>,t>}$ [$_S$ everyone$_{<<e,t>,t>}$ [$_S$ <e> loves <e>]]]

b. 'for everyone, there is someone who loves him/her'
[$_S$ everyone$_{<<e,t>,t>}$ [$_S$ someone$_{<<e,t>,t>}$ [$_S$ <e> loves <e>]]]

In its simplest formulation, Quantifier Raising is a syntactic operation that targets every quantifier and takes it to a position outside the sentential core (marked as 'S', short for 'sentence'). Whenever a particular sentence contains two (or more) quantifiers, this means that two (or more) applications of Quantifier Raising are required. Each of these instantiations of QR will attach a quantifier to the left edge of the sentence. So at the end of the day, what we will get is a bunch of quantifiers at the left edge of S, each attached as an adjunct to S. The order in which QR targets these various quantifiers is free in principle. And as a result, the various quantifiers can end up in different positions and scopal relations relative to one another in the finished product. In the derivation reflected by the structure below (14a), *everyone* undergoes QR first, followed by *someone*, the latter ending up structurally higher than *everyone* and hence taking SCOPE over *everyone* (a case of 'linear scope': the relative scopes of the quantifiers correspond to their relative positions in the linear string of the pronounced product). In (14b), by contrast, QR moves *someone* first and then attaches *everyone* above it, resulting in a reading in which *everyone* scopes over *someone* (a case of 'inverse scope': the LF in (14b) turns the pronounced output product upside down, giving the object quantifier a scope that is *wider* than that of the subject quantifier).

Scope ambiguity is not strictly the privilege of sentences with multiple quantified noun phrases. We find it also in sentences with just a single quantifier, as long as the sentence in addition contains some other scope-taking element. Negation is a scope taker. Often negation scopes over an entire sentence (as in *John didn't talk to anyone*), but sometimes its scope is much narrower, as in *John talked not to Mary but to Sue*: here it is not being denied that John talked to someone; it just was not Mary who he talked to. Because negation is a scope-taking element, it is expected that when we combine into a single sentence a negation and a quantified noun phrase, scope ambiguities should be able to arise. Indeed they do: sentences of the type in (15) are ambiguous. On the reading for (15) paraphrased in (16a), the sentence is equivalent to *nobody left* or *everbody stayed*. But if (15b) is the correct paraphrase, it is entirely possible that some people left while others did not, as in *not everybody left*.[1]

(15) everyone did not leave

a. 'for everyone it is the case that they did not leave'

b. 'it is not the case that everyone left'

1 Here we encounter something that we had run into previously in our discussion of 'bracketing paradoxes' of the *unhappier* type in segment II. The morphology combines *un-* and *-er* with the stem *happy* in a particular way — *-er* attaches first (because it resists being attached to a host that has more than two syllables), after which *un-* comes in. But although *un-* is combined with *happier* and would thus be expected to scope over *-er*, the most natural interpretation for *unhappier* is one which is paraphrased as 'more unhappy'. We can take care of this reading by allowing the relative scopes of *un-* and the comparative morpheme *-er* (which is also scopal: cf. *more unhappy*, where *more* is explicitly a quantifier) to be reversed to yield a Logical Form in which *-er* scopes over *un-*. If this is tenable, movement of scope-taking elements can target not just constituents of syntactic structures but also subparts of word-size constructs, which will both enhance the parallelism between morphology and syntax and emphasise the pervasiveness of scope-driven movement in natural language.

In English, a particular surface structure does not always correspond directly to the semantic interpretation that one would like to associate with it. As we have discovered, this can be dealt with quite elegantly by a syntactic movement rule. But there are also languages, such as Hungarian, in which surface strings typically 'wear their semantics on their sleeve', and do not as readily give rise to ambiguity as they do in English. While English (14) is ambiguous, in Hungarian the string in (16) only has the paraphrase in (16a), not (16b):

(16) valaki mindenkit szeret
someone everyone loves
a. 'there is somebody such that (s)he loves everyone'
$[_S$ someone$_{<<e,t>,t>}$ $[_S$ everyone$_{<<e,t>,t>}$ $[_S$ <e> loves <e>]]]
b. *'for everyone, there is someone who loves him/her'
*$[_S$ everyone$_{<<e,t>,t>}$ $[_S$ someone$_{<<e,t>,t>}$ $[_S$ <e> loves <e>]]]

A plausible response to this is to say that already in the overt syntax of Hungarian (16), the quantifiers are in their Logical Form positions, outside the sentential core. The core sentence will then only contain the variables bound by these quantifiers; and the fact that *valaki* is pronounced before *mindenkit* tells us that the subject quantifier is structurally higher than the object quantifier, giving us just the reading in (16a). With Hungarian QR'ing its quantified noun phrases in overt syntax, before the output product is pronounced, we derive the desired result that, in the general case, Hungarian encodes the scope relations between quantifiers in the surface linear string.

English does not perform QR in the overt syntax: it postpones its application until after the sentence has already been given its phonological form. (This is called 'covert movement'.) There are two ways in which (14) can 'feed' QR, with the structures given below (14a,b) as the outputs. Although we do not notice on the sound side the effect of the operations that place the quantifiers outside the core sentence, we do notice their effect on the meaning side of the grammar.

It is the kinds of form/meaning NON-ISOMORPHISM (or mismatch) represented by the b–readings of (14) and (15) that linguistic semantics is very much interested in. Such lack of isomorphism involves aspects of meaning which are conditioned by syntactic structure.[2] We see this clearly from the fact that quantifier-scope ambiguities of the type seen in (14) and (15) do not arise just any time we have two quantifiers: they arise only if the two quantifiers are in the same clause together — the 'clausemate restriction' on QR, illustrated in (17). This is a syntactic fact: it makes reference to syntactic locality.

(17) someone hopes that everyone will get a 5 for this course
a. 'there is someone who hopes that everyone will get a 5 for this course'
b. *'for everyone, there is (a potentially different) someone who hopes that they will get a 5 for this course'

IV.5 Entailment, upward and downward

Sentences can logically entail other sentences, and such entailments can go 'from smaller to larger' (UPWARD ENTAILMENT) or 'from larger to smaller' (DOWNWARD ENTAILMENT):

2 We will discover a different kind of form/meaning non-isomorphism in our discussion of meaning *ex nihilo* in §IV.10.

(18) a. John is a linguist ⇐ ENTAILS ⇐ John is a good linguist
b. John is not a linguist ⇒ ENTAILS ⇒ John is not a good linguist

From the truth of *John is a good linguist*, we deduce that John must be a linguist; but from the truth of *John is a linguist*, we must not jump to the conclusion that he is a good linguist. The entailment relation between the two sentences in (18a) goes from smaller to larger — the set denoted by *a good linguist* is smaller (i.e., contains fewer candidates in the extra-linguistic universe) than the set denoted by *a linguist*; (18a) allows us to reason logically from the subset to the set but not the other way around. This is a case of upward entailment. In (18b), on the other hand, we can reason from the set to the subset but not the other way around: if *John is not a linguist* is true, it must also be true that he is not a good linguist (after all, there are no good linguists who aren't linguists); but from the truth of *John is not a good linguist*, we can draw no conclusions about his being or not being a linguist. The entailment here goes from the set to the subset, a case of downward entailment. Sentential negation systematically reverses the usual entailment relation between sets and subsets.

Questions and conditionals, like negation, affect the entailment relations between propositions. From the previous paragraph, we recall that *John is a good linguist* entails that John is a linguist. But does the question in (19b) give rise to the entailment that John is a linguist? The answer is 'no': *is John a good linguist?* does not give rise to any conclusion regarding John's being or not being a linguist. Unlike the declarative statement in (19a), the question in (19b) is not upward entailing. The same is true for the conditional in (19b): again, we cannot reason from the subset to the set; it may very well be the case that only good linguists, and not linguists in general, can speak many languages.

(19) a. is John a good linguist?
b. if John is a good linguist, he must be able to speak many languages

Certain elements of natural language are sensitive in their distribution to the question of whether the environment that they occur in is upward entailing or not. So-called NEGATIVE POLARITY ITEMS (NPIs, for short) do not want to occur in upward-entailing environments. Thus, (20a) is quite impossible if the NPI *any* is included in it; but (20b–d) do allow *any* to be used.

(20) a. John knows (*any) Latin
b. John doesn't know (any) Latin
c. does John know (any) Latin
d. if John knows (any) Latin, he will be eligible for this job

IV.6 Implicatures: On the threshold of semantics and pragmatics

Though in everyday speech we might treat the verbs *to entail* and *to imply* quite interchangeably, there is an important terminological difference to be made between entailments and IMPLICATURES. The discussion of implicatures in this section will take us on the road from semantics to pragmatics.

The British philosopher of language Paul Grice is known in linguistics for introducing implicatures, and making a key distinction between two types of them:

(21) a. CONVENTIONAL implicatures
b. CONVERSATIONAL implicatures

These have become mainstays of work in semantics and pragmatics. Conventional implicatures are semantic in nature; they cannot be 'cancelled' because they are, by convention, inherent in the semantic properties of a particular element in the sentence. Conversational implicatures, on the other hand, are pragmatic, arising in the conversation, and are 'cancellable'.

Consider the sentences in (22):

(22) a. among the first-year students, only John (*among others) likes syntax
b. among the first-year students, it is John (#among others) who likes syntax

For (22a), which uses the focus particle *only*, the implicature of EXHAUSTIVITY is conventional — it is conventionally associated with the meaning of the lexical item *only*. Because of its conventional, lexically associated nature, the exhaustivity implicature of (22a) is not cancellable: insertion of *among others* is semantically incongruous, hence the asterisk.

In (22b), on the other hand, exhaustive (aka identificational) focus is syntactically expressed with the aid of a so-called cleft construction (*it is x that S*). Here, there is also an implicature of exhaustivity, but this time around, the implicature is merely conversational: in a typical discourse, the interlocutor would indeed expect the focus of (22b) to identify exhaustively the relevant argument of the predicate of the relative clause, and therefore to conclude that John is the only first-year student in the universe of discourse who likes syntax. But because the exhaustivity implicature of (22b) is not associated with the lexical meaning of any particular element in the sentence (*it* is not lexically exhaustive, nor is *is* or the relative operator *who*), it is cancellable in principle: insertion of *among others* is pragmatically awkward but not ungrammatical in (22b).

IV.7 Information structure

In our brief discussion of implicatures in the previous section, we came across the notion of exhaustive/identificational focus. This is one of several DISCOURSE roles that the referents of syntactic constituents can play in the INFORMATION STRUCTURE of a sentence. Broadly speaking, the following information-structural functions can be distinguished:

(23) a. TOPIC *vs* COMMENT
b. FOCUS *vs* PRESUPPOSITION

There are a wide variety of definitions of 'topic' and 'focus' available in the literature — a consequence of the fact that there are a wide variety of different topics and foci (where 'foci' is the plural of 'focus'). We are not going to be able to do justice to the complexities of the notions of 'topic' and 'focus' here. For our purposes in this discussion, we will understand 'topic' as 'familiar, old information' against the background of which the 'comment' is made, and 'focus' as 'new information' updating or correcting the 'presupposition' (i.e., what is supposed to be the common ground shared by the speaker and the listener).

Communications of the following sort (from the official English translation of the briefing distributed by the ELTE Epidemiological Operative Coordinating Body on 11 March 2020) may help illustrate the (occasionally rather disconcerting) presuppositions that 'lie beneath' utterances:

(24) a. 'In order to maintain the normal operation of the University, especially with regard to tasks related to the organising of distance learning activities, teachers, researchers, and members of the administrative staff are obliged to continue working during the ordered Rectorial break and the Spring break, as well as in the period following these.'

b. 'Doctoral students are obliged to continue their research activities and to participate in distance learning.'

The focus of both (24a) and (24b) is the predicate headed by *obliged*. The presupposition is that employees and doctoral students would *not* continue working during the relevant periods: if it had been taken for granted that work would continue, the focus of these statements would not have been sensible.

The English sentences in (25a) and (25b) both allow either of the two noun phrases to have the information-structural role of focus, with the other one being a topic. But in the context of a constituent question, we can bring the information-structural roles of these noun phrases to light: in answer to the question *Who does Adam love?*, the sentence *Adam loves Eve* gives *Eve* the role of focus, and *Adam* the role of topic; but in answer to the question *Who loves Eve?*, the sentence *Adam loves Eve* assigns the role of focus to *Adam*, with *Eve* being a topic. The same applies, *mutatis mutandis*, to the example in (25b).

(25)	a.	Adam loves Eve	Q1: who does Adam love?	⇒ *A*=TOPIC, *E*=FOCUS
			Q2: who loves Eve?	⇒ *E*=TOPIC, *A*=FOCUS
	b.	Eve is loved by Adam	Q1: who is Eve loved by?	⇒ *E*=TOPIC, *A*=FOCUS
			Q2: who is loved by Adam?	⇒ *A*=TOPIC, *E*=FOCUS

Though there is a tendency for the grammatical subject to be the topic of the sentence, in English this is by no means a hard-and-fast rule: both sentences in (25) are information-structurally ambiguous in principle. But the sentences in (26)–(29) are each unambiguous — these are syntactic constructions (so-called pseudo-cleft and *it*-cleft sentences, resp.) with a special information-structural profile.

(26)	a.	the woman who Adam loves is Eve	⇒ *A*=TOPIC, *E*=FOCUS
	b.	the man who loves Eve is Adam	⇒ *E*=TOPIC, *A*=FOCUS
(27)	a.	the man who Eve is loved by is Adam	⇒ *E*=TOPIC, *A*=FOCUS
	b.	the woman who is loved by Adam is Eve	⇒ *A*=TOPIC, *E*=FOCUS
(28)	a.	it is EVE who Adam loves	⇒ *A*=TOPIC, *E*=FOCUS
	b.	it is ADAM who loves Eve	⇒ *E*=TOPIC, *A*=FOCUS
(29)	a.	it is ADAM who Eve is loved by	⇒ *E*=TOPIC, *A*=FOCUS
	b.	it is EVE who is loved by Adam	⇒ *A*=TOPIC, *E*=FOCUS

The use of particles such as *only* and *even* also yields information-structural disambiguation, as we see in (30) and (31):

(30) a. Adam only/even loves EVE ⇒ *A*=TOPIC, *E*=FOCUS
b. only/even ADAM loves Eve ⇒ *E*=TOPIC, *A*=FOCUS
(31) a. Eve is only/even loved by ADAM ⇒ *E*=TOPIC, *A*=FOCUS
b. only/even EVE is loved by Adam ⇒ *A*=TOPIC, *E*=FOCUS

Hungarian was a major catalyst of the research in this area, thanks to the fact that Hungarian syntax is thought not only to 'wear its LF on its sleeve' but also to 'wear its information structure on its sleeve'. In Hungarian, topics and foci tend to be in designated structural positions the clause, with the natural order of constituents being 'topic before focus'. We see this in the examples in (32) (where, as before, small capitals indicate focus).

(32)	a.	Ádám csak ÉVÁT szereti Adam only Eve.ACC loves	*csak ÉVÁT Ádám szereti only Eve.ACC Adam loves	⇒ *A*=TOPIC, *E*=FOCUS
	b.	Évát csak ÁDÁM szereti Eve.ACC only Adam loves	*csak ÁDÁM Évát szereti only Adam Eve.ACC loves	⇒ *E*=TOPIC, *A*=FOCUS

The English syntax/information-structure interface is rather less 'rigid' than its counterpart in Hungarian. English has (pseudo-)cleft constructions of the type in (26)–(29), but uses them only under special circumstances. Also, English tends not to place constituents at the left edge of the sentence purely for information-structural purposes — sentences such as *HER I like, not HIM* (contrastive focus fronting) and *Her, I like; him, I don't* (contrastive topic fronting) are relatively uncommon in English (though they are indubitably grammatical).

Although English often exhibit a rather loose relationship between syntax and information structure, *wh*-questions are much stricter: in general, these MUST front the *wh*-constituent, and in root contexts *wh*-fronting also MUST bring about subject–auxiliary inversion (SAI):

(33)	a.	who does Adam love?	*Adam loves who?
	b.	I wonder who Adam loves	*I wonder Adam loves who?

The rule of *wh*-fronting applies generally to all questions, both root and non-root ones — but there are two notable exceptions: *(a)* 'echo questions' and *(b)* 'quiz master questions'. We can dismiss *(a)* quickly: in a situation in which someone tells you that John was eating sushi, you can respond by echoing back the statement in the form given in (34B), with the *wh*-word *in situ* (i.e., in its normal place, to the right of the verb); but echo questions are not genuine requests for information — rather, they express disbelief at a statement previously uttered in the discourse.

(34) A: John was eating sushi
B: John was eating WHAT?

So-called 'quiz master questions', illustrated by (35), do serve as requests for information, hence are *bona fide wh*-questions. The fact that they do not feature *wh*-fronting is in all likelihood a direct reflex of the special type of discourse in which they are used. Under normal circumstances, a questioner needs to signal to his/her interlocutor(s) that (s)he is asking a question, and *wh*-fronting, in conjunction with SAI, has the function of signalling precisely this. But in the context of a quiz, it is clear to all the participants in the discourse that the quiz master is asking questions — that is the quiz master's express role. So the fact that (s)he is asking the contestants a question is not something

that the quiz master needs to signal: it is plainly understood. This is exactly what allows us to explain the fact that quiz master questions forgo *wh*-fronting and concomitant SAI: these operations are *redundant* here, and therefore they are not performed. It is only in discourses in which the status of a person's utterance as a question is not a given that the utterance must be explicitly marked as a question; and in English *wh*-fronting plus SAI plays this role.

(35) for $500, John F. Kennedy was shot in which American city?

A different kind of exception to the general rule that English *wh*-questions must involve *wh*-fronting comes from questions in which the identity of more than a single argument in the event is questioned at the same time:

(36) a. who kissed who?
b. who ate what?
c. who gave what to whom?

Unlike (35), such sentences do not counterexemplify the need to have a *wh*-constituent at the left edge of a question: there *is*, after all, a *wh*-phrase in initial position in the sentences in (36). The exceptional character of (36) lies in the fact that not all of the *wh*-phrases in the question undergo *wh*-fronting: in fact, only one *wh*-phrase is allowed to front in English (though there are languages, especially in the Balkan *Sprachbund*, which front all *wh*-expressions *en bloc*).

Multiple *wh*-questions such as (36) are of interest to students of linguistic meaning because they are not all alike as far as their interpretation and discourse function are concerned. For (36a), a natural condition of use would be a situation in which the speaker is familiar with the participants in the kissing event but would like to know who instigated the kissing. Stricly speaking, (36a) is not a request towards the interlocutor to learn the names of the participants: the questioner already knows that, say, John and Mary were the ones involved. What the questioner would really like to know is the direction of the exchange: whether John planted a kiss on Mary, or *vice versa*. In a discourse situation such as the one just described, the multiple *wh*-question in (36a) can be replied to with just a single pair of NPs.

For questions such as those in (36b), such a reply typically would not do. The natural discourse context in which such questions are used is one in which a PAIR-LIST ANSWER is solicited: John ate sushi, Donald ate steak, Bill ate natto, etc. etc. To reply to (36b) with just *John ate sushi* would be distinctly awkward. For (36c) it would likewise be inappropriate to provide a single triple in response. But here it is legitimate for the responder to enumerate for each individual choice of giver a list of pairs of gifts and beneficiaries: *John gave a book to Sue, a CD to Bill, a plant to Mary*, etc. etc., followed by *Bob gave a picture to Sue, a postcard to Bill, a gift card to Mary*, etc. etc. The pragmatics of double and triple *wh*-questions will likely continue to be under close scrutiny for some time. The last word certainly has not been said about these constructions.

Just like it is not the case that every *wh*-phrase must front in an English root *wh*-question, it also not the case that the finite auxiliary must show up to the left of the subject in every English root question. Indeed, when the initial *wh*-element is itself the subject, the auxiliary never occurs to its left: *who is eating sushi?* is grammatical while **is who eating sushi?* is ill-formed. The distribution of subject–auxiliary inversion in *wh*-questions appears to be strictly grammatically governed in English: it does not appear to have any effect on meaning.

Things are different in the case of SAI in tandem with the fronting of a negative constituent to sentence-initial position. Only if the fronted negative-marked constituent is a focus does the syntax give rise to SAI; when the negative constituent is a topic, no SAI takes place. This gives rise to minimal pairs of the type in (37):

(37) a. with no job would she be happy ⇒ *with no job* = FOCUS
b. with no job, she would be happy ⇒ *with no job* = TOPIC

The sentence in (37a) says that there is no job that could possibly make her happy. The message conveyed by (37b), on the other hand, is that if she were out of a job, she would be happy.

IV.8 Referential dependencies

In our discussion of syntax, we came across sentences (see §III.9) for which postulating a silent pronoun ('PRO') in the subject or object position of an infinitival clause was helpful in accounting for their ambiguity:

(38) there is always someone to help out
a. 'there is always someone that can help people out'
someone$_i$ [$_S$ PRO$_i$ to help out]
b. 'there is always someone that people need to help out'
someone$_i$ [$_S$ to help PRO$_i$ out]

Once we have allowed ourselves the liberty of postulating these abstract silent pronouns, there is no further indeterminacy regarding the interpretation of these pronouns. The silent pronouns in the representations given below the paraphrases in (38a) and (38b) are strict with respect to their choice of antecedent: they are inextricably linked to *someone* (as marked by the coindexation).

In many cases, however, it is rather more difficult to make out what a pronoun is supposed to make reference to, even in connected discourse. Consider, for instance, the mini-discourse in (39):

(39) A: 'Have you heard anything about John?'
B: 'Yes, at the office the other day I saw that John was engaged in a conversation with Bill. And I heard that he told him that he might fire him.'

B's second sentence contains four pronouns, all masculine singular. If we exclude (for simplicity) the possibility that any of these pronouns refers to someone not identified in the mini-discourse, there is still a choice of antecedents for these pronouns: *John* or *Bill*. But without any further knowledge about the roles played by John and Bill at the office, it will be hard to figure out who told whom about the firing, and who was firing whom. Imagine that we know that John is the boss and Bill one of his employees. Then our pragmatic knowledge (in particular, our knowledge about firings) will tell us that the pronoun *he* in *he was going to fire him* must refer to *John*. But that still does not necessarily entail that the pronoun *he* in *he told him* will also refer to *John*: it is imaginable, after all, that Bill was telling his boss that he (i.e., the boss) might fire him. The only thing that we can be absolutely certain of in this mini-discourse is that the pronouns *he* and *him* in *he told him* cannot refer to the same person, and that the same is true for the pronouns *he* and *him* in *he might fire him*.

That the pronoun *him* that serves as the object of *fire* cannot refer to the same person as the pronoun *he* that serves as the subject of *fire* is something we know not just on the basis of world knowledge (bosses are not in the habit of firing themselves): it is our syntactic knowledge regarding referential dependencies that tells us that a personal pronoun that serves as an argument of a certain predicate can never take as its antecedent another argument of that same predicate. This holds throughout — also, therefore, of the predication of which *told* is the head. So in *he told him*, too, it is impossible for *he* and *him* to be coreferential — not because pragmatic factors rule this out (people do actually tell themselves all sorts of things) but because the sentence grammar (i.e., the syntax) does not tolerate such local coreference. If we want two co-argument pronouns to be coreferential (i.e., to refer to the same individual), we must build a REFLEXIVE predication, which requires the use of a reflexive ANAPHOR as the object. And when we do indeed choose a reflexive anaphor as the object, as in (40), it *must* refer back to the pronoun in subject position.

(40) he told himself that...

The distribution of personal pronouns and reflexive anaphors generally involves a rigorous division of labour, such that in syntactic contexts where a pronoun can be used, it cannot be replaced with an anaphor cannot be used, and in contexts where an anaphor can be used, replacing it with a pronoun is impossible. In (41a) the two pronouns cannot be coreferential; if we want coreference of the subject and object of *like*, the object needs to be a reflexive anaphor, as in (41b); switching the pronoun and the anaphor such that the latter is the subject and the former the object results in ungrammaticality, which persists even if we replace *him* with another reflexive anaphor. Reflexive anaphors require a co-argument as their antecedent; and the subject of a certain predicate can never serve as the antecedent of the object of the same predicate.

(41) a. I consider him_1 to like him_2 $1 \neq 2$
b. I consider him_1 to like himself_2 $1 = 2$
c. *I consider himself_1 to like him_2
d. *I consider himself_1 to like himself_2

All of the things mentioned in the previous paragraph are pieces of syntax-based knowledge which speakers of English share, and for which no pragmatic skills or information about the extra-linguistic world is required. But beyond this knowledge, determining the antecedent of a pronoun is very often a matter in which discourse participants need to rely heavily on pragmatics and world knowledge.

Prosody sometimes comes to the rescue to disambiguate. Take, for instance, the following sentence in (42). Pronounced with an intonation pattern in which little or no prosodic prominence is given to the pronouns *he* and *him*, the sentence in (42) will be understood in such a way that it was John who hit Bill, after first having gently touched Bill on the back. But now imagine that we pronounce both *he* and *him* with heavy stress, as in (43) (where the small capitals mark emphasis). This time, we understand the second conjunct in such a way that it was Bill who violently reciprocated John's gentle gesture. Syntactically, (42) and (43) look exactly the same; but their interpretations are exactly each other's opposites.

(42) first John_1 gently touched Bill_2 on the back, and then he_1 hit him_2
(43) first John_1 gently touched Bill_2 on the back, and then HE_2 hit HIM_1

The speaker of (43) would not have resorted to the special prosodic effects employed here unless (s)he wanted to signal what these prosodic effects give rise to. From the point of view of speech production, it takes much less effort to pronounce utterances with neutral intonation than with these special prosodic effects. So such effects are used sparingly; and when they do get used, this is a red flag to the interlocutor that some interpretive effect is desired which could not (as easily) have been conveyed if neutral intonation had been chosen.

IV.9 Grice's Cooperative Principle and the Gricean Maxims

The way in which at the end of the previous section we saw prosody being exploited for interpretive effect is a powerful illustration of the dynamism of the interaction between interlocutors in discourse. The philosopher Grice, whose work was already referred to in §IV.6 in connection with implicatures, formulated (44) as the cornerstone of such interaction, and on (44) he based the four maxims given in (45)–(48).

(44) COOPERATIVE PRINCIPLE
'make your contribution such as it is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged'

(45) MAXIM OF QUALITY
make your contribution one that is true
a. do not say what you believe to be false
b. do not say that for which you lack adequate evidence

(46) MAXIM OF QUANTITY
a. make your contribution as informative as is required for the current purposes of the exchange
b. do not make your contribution more informative than is required

(47) MAXIM OF RELEVANCE
make your contribution relevant

(48) MAXIM OF MANNER
make your contribution in a perspicacious manner

It is interesting to note that the way in which the Cooperative Principle is defined in (44) actually flouts Grice's own Maxim of Quantity ('be brief'): just the first part of (44) ('make your contribution such as it is required') would have sufficed — after all, every discourse contribution inevitably occurs at *some* stage, and can reasonably be expected to have *some* purpose in or give *some* direction to the discourse in question. The material following 'make your contribution such as it is required' is strictly speaking redundant given the knowledge that this statement is about cooperative discourse (knowledge that is straightforwardly derivable from the name of the principle).

The Maxims of Quality ('do not lie'), Quantity ('be concise and to the point'), Relevance ('be relevant') and Manner ('be clear') are guidelines for proper participation in a conversation. They can be flouted at will — and often are, for a particular purpose. The language of politicians and civil servants frequently stands out as particularly disrespectful of the Gricean Cooperative Principle and its maxims: in politics it is often considered desirable to hedge, use woolly language, or flat-out lie:

(49) a. 'If you ask me for a straight answer, then I shall say that, as far as we can see, looking at it by and large, taking one thing with another in terms of the average of departments, then in the final analysis it is probably true to say, that at the end of the day, in general terms, you would probably find that, not to put too fine a point on it, there probably wasn't very much in it one way or the other, as far as one can see, at this stage.'

b. 'Unfortunately, although the answer was indeed clear, simple, and straightforward, there is some difficulty in justifiably assigning to it the fourth of the epithets you applied to the statement, inasmuch as the precise correlation between the information you communicated and the facts, insofar as they can be determined and demonstrated, is such as to cause epistemological problems of sufficient magnitude as to lay upon the logical and semantic resources of the English language a heavier burden than they can reasonably be expected to bear.'

These statements (made by the fictional senior civil servant Sir Humphrey Appleby in the BBC Television sitcom *Yes Minister*) are particularly non-cooperative discourse contributions, illustrating the ways in which the Gricean Maxims can be flouted for a particular purpose (here obfuscation or evasion).

IV.10 Meaning *ex nihilo*: Silent operators

In §IV.4, it was pointed out that there can be a lack of isomorphism between form and meaning in sentences with multiple scope-taking elements — thus, recall the fact that a sentence such as *someone loves everyone* supports a so-called inverse-scope reading, with *everyone* taking scope over *someone*. A more extreme kind of form/meaning non-isomorphism presents itself in cases in which we get some sort of meaning without there apparently being anything in the overt signal that is responsible for contributing it — meaning *ex nihilo* ('out of nothing, out of the blue'), if you will.

One such case is the use of silent pronouns, as in (38) (repeated below). Here we have silence with meaning — but this is at least syntactically represented silence, so there is something in the syntactic representation that the semantics can interpret (albeit that this 'something' happens to lack a phonological realisation).

(38) there is always someone to help out

a. 'there is always someone that can help people out'
someone$_i$ [$_S$ PRO$_i$ to help out]

b. 'there is always someone that people need to help out'
someone$_i$ [$_S$ to help PRO$_i$ out]

There is another covert, 'under the radar' aspect to the interpretation of (38): the MODAL SEMANTICS of the relative clause. A close look at the paraphrases in (38a) and (38b) makes it apparent that they contain an element of modality, brought out by *can* and *need*. There appears to be nothing about the meanings of the individual ingredients of the relative clause that predictably gives rise to this. We might perhaps be inclined to attribute the modal semantics of the relative clause in (38) to the infinitival marker *to*. But this is unlikely to be tenable. There are infinitival clauses that have *to* but no element of modal semantics. Take, for instance, the sentence *everyone considers Mary to be*

smart. There is no modality involved here at all: Mary simply *is* smart, and everyone agrees on this. Even though the infinitival marker *to* has a syntactic relationship with modals such as *can* (they arguably are housed by the same node in the syntactic structure of the clause), they do not have the same meaning; indeed, *to* appears not to have any meaning at all. So the modality of (38) really seems to come *ex nihilo*.

A possible response to meaning out of the blue is to assume that there is a silent OPERATOR in the structure of the sentence on which the mysterious semantic contribution can be pinned. For (38), we could, in this vein, postulate a null modal operator in the infinitival clause. This approach makes good sense, but it faces a serious question: under what circumstances are we allowed to postulate abstract operators? Imagine that we could do this freely: then any utterance could mean anything that could be contributed by an abstract operator, in addition to the meanings contributed by the overt ingredients of the utterance. That would plainly be a stretch. So we have to be careful in our deployment of silent semantic operators. We do not want something like *you are smart* to be able to mean exactly the opposite.

Or do we? Some forms of irony and sarcasm hinge precisely on this. If someone were to say to you something along the lines of (50), you probably would not take it as a compliment — especially if it was pronounced with the kind of prosody that is typical of ironic/sarcastic statements.

(50) a. 'Oh, you're so smart!'
b. 'Oh, so you're an expert, are you?'

Perhaps (50) works pragmatically as a negative evaluation because the structure of (50) can harbour a silent negative operator. What may jibe with this is the use of so-called constant polarity question tags in combination with ironic positives-turned-negative, as in (50b). In the typical case, English tag questions work in such a way that a statement is accompanied by a tag of opposite polarity — so if the statement is positive, the tag is negative (*she is very smart, isn't she?*), and *vice versa* (*she isn't very smart, is she?*). But ironic positives are followed by positive tags: a speaker saying (50b) is using the positive tag precisely to bring out the irony/sarcasm of his/her ascription of expertise to the addressee. The use of constant polarity tags may be understood against the background of the general rule of polarity reversal in tag questions if the 'false positives' in (50) are represented in such a way that they contain a silent negative operator: then the host for the tag in (50b) is in fact negative, and the use of the positive tag is no exception to the polarity reversal seen elsewhere in English tags.

Though constant polarity tags may thus support the postulation of a silent negative operator in (50), there is a strong indication that this operator cannot be assumed to be present in the syntactico-semantic representation of these ironic/sarcastic positives. From §IV.5, you will recall that negative polarity items are expressions which flourish in downward-entailing environments (incl. negative statements) but cannot occur in upward-entailing ones: while *John doesn't know any Latin* is grammatical, **John knows any Latin* is not; similarly, *John doesn't know Latin at all* works but **John knows Latin at all* does not. The NPIs *any* and *at all* in the grammatical examples are licensed by the negative operator, realised here as *-n't*. Now, if (50), in its ironic/sarcastic use, features a silent negative operator, *and* if this operator is present in the syntactico-semantic representation of these utterances, the expectation arises that such utterances could harbour negative polarity items. But they cannot:

(50′) a. 'Oh, you're so good at (*any) Latin (*at all)!'
b. 'Oh, so you're an expert on (*any) languages (*at all), are you?'

The ungrammaticality of (50′) with *any* and *at all* included need not defeat the hypothesis that in their ironic/sarcastic use, utterances such as (50) have a silent negative operator in them. This hypothesis can be successful if *(a)* the silent negative operator is introduced only in the pragmatic representation of these sentences, and *(b)* negative polarity items must be be licensed in the course of the syntactico-semantic derivation (which integrally precedes pragmatic analysis). Of course, such a hypothesis makes it incumbent on the analyst to determine precisely under which circumstances we are allowed to add meaning to an utterance via the inclusion of a silent operator in the pragmatic representation of the utterance. This is a far from simple matter. We will not have an opportunity to pursue this any further here: we have reached our limit for this segment, so we must leave this, as they say, for future research.