

1. What Semantics is About

1.1. By Way of Introduction

This is an introduction to semantics. What is semantics? There is an everyday use of “semantics” that appears in the following passage culled from the World Wide Web¹

A vote on a child custody bill its sponsor says only changes the word “custody” to “parental responsibility” was postponed for a second time in the Colorado Senate on Friday. Critics of HB 1183 say the measure is **not just semantics** but a serious change to Colorado’s child custody laws that could outlaw joint custody. “This will change hundreds of thousands of people’s lives in Colorado in regard to their ex-spouse’s life and their children,” said Senate Minority Leader Mike Feeley, D-Lakewood, in arguing against the measure.

It is not quite clear why “semantics” got such a pejorative meaning in popular culture. Presumably it just reflects an aversion against the language of jurisdiction. But you should be aware of that when you tell people that you study semantics — or, “just semantics”. In any case, even the Denver Post article continues with:

"Something can be just semantics, but semantics can be important," Harhai [a divorce lawyer] said. "Words can make a huge difference in how people view themselves and their kids and the outcome of their divorces."

Sometimes semantics can indeed be very important, and not only in legal settings. The worst accident in the history of aviation happened on March 27, 1977. Two aircraft collided on a runway in Tenerife, 583 people died. One of the several problems that led to this terrible accident was the following misunderstanding:

- (1) Pilot of KLM 4805: Ah - roger sir, we are cleared to the Papa Beacon, flight level nine zero until intercepting the three to five. We are now at takeoff.
Tower: OK... [masked by noise] Stand by for takeoff, I will call you.

The pilot understood *we are now at takeoff* as ‘we are now taking off’, whereas the tower understood it as ‘we are now at the takeoff point’.²

1.2. Semantics: Form and Meaning

The term *Semantics* goes back to the Greek root *sema* ‘sign, feature’, and refers to the study of meaning, especially in natural languages. A related term, *semiotics*, concerns the study of signs in general, including, e.g., the meaning of human habits, traffic signs, or information processing of animals.

Let us look at the spectrum of the various branches of linguistics³. It is a commonplace assumption in linguistics that language relates **physical phenomena** to **meanings**.

- (2) phonetics — phonology — morphology — syntax — semantics
|—————subject area of linguistics—————|
articulatory movements things?
acoustic phenomena meanings?


¹ “Denver Post Online”, March 1998, author Peggy Lowe.

² Reported in Steven Cushing, *Fatal Words: Communication clashes and aircraft crashes*, University of Chicago Press 1994.

³ This is of course a simplified picture, as it does not cover many branches of linguistics, like historical linguistics or sociolinguistics.

The subject matter of **phonetics** is the physical side of linguistic expressions and the physiological abilities of humans to produce and recognize them. This includes the physical properties of acoustic phenomena related to spoken language and the relevant articulatory movements and auditory abilities, but also sign-language gestures and written symbols.

The subject matter of **semantics**, on the other hand, is much more difficult to grasp. It is, roughly, what linguistic signs stand for. In the case of names this seems to be simple enough. *Abraham Lincoln*⁴ stands for a certain person, *San Francisco* stands for a city, *Rigel* for a star, *Koh-i-noor* for a certain diamond now in the British queen's state crown, *Andrew* (among many other things) for a hurricane that hit Florida in 1994. In other cases this is more difficult. What does the noun *dog* stand for? For all the dogs in the world? But imagine that dogs die out; does this mean that *dog* loses its meaning? Not so; even though both dodos and passenger pigeons died out, the nouns *dodo* and *passenger pigeon* mean something, and their meanings are different. Hence we must distinguish between what expressions stand for in the real world, and their meanings. This relation between expressions, meanings and things is often given in form the so-called **semiotic triangle** (cf. Ogden & Richards 1923, *The meaning of meaning*):



That is, linguistic expressions are first related to **meanings**, and meanings may be related to objects in the world. Hence the relation between expressions and things is an indirect one. The task of semantics then is to develop theories about the relation between expressions, meanings, and the objects these meanings stand for. This is expressed in the following quote of the commentary of Aristotle by Muhammad Al-Farabi (870-950) (please read “semantics” for “logic”):

One of the first things anyone taking up logic must know is that there are sense-objects or, more generally, entities outside the soul [= things]; furthermore thoughts, pictures, and representations within the soul [= meanings]; and [finally] speech and scripts [= expressions]. We must know how they relate to one another; for the logician considers thought as relating both sides, namely to the entities outside the soul and speech. He also studies speech by itself, but always in terms of its relation to thought.⁵

In this course we will be mostly interested in meanings and how meanings and natural-language expressions relate to each other.

Now, the entities that semantics deals with, meanings, cannot be observed directly (they are, as Al-Farabi put it, ‘within the soul’). And hence the proper nature of meanings constitute a serious problem, perhaps the most serious problem, for philosophers from ancient times till today.

Meanings certainly may lead to **actions** that are **physically observable**. We can see this most easily with **commands**. For example, if A tells B: *Come here!*, and B moves to A, then we would say that B has grasped the meaning of A’s sentence and, as a consequence, carried out a physically observable act. The case of **assertions** is perhaps more indirect. A relatively clear case is if A informs B:

⁴ I follow here the convention to *italicize* the words and expressions that we talk about when they occur in a text. Another convention is to use quotation marks. See the comments below on object language and metalanguage.

⁵ Cited after F.W. Zimmermann (ed.), *Al-Farabi’s Commentary and Short Treatise on Aristotle’s ‘De Interpretatione’*, Oxford University Press, 1981.

(4) You are standing on my foot.

If B removes his foot, then this is a relatively reliable sign that B understood the meaning of this sentence. But what about if A, a teacher, tells B, a student,

(5) The capital of Iceland is Reykjavik.

It seems that there is no predictable observable phenomenon related to this sentence. And even in the first two examples, the observable reactions are not really uniform: For example, it may be that B understands A's command, but simply decides not to obey. It seems hopeless to relate meanings to directly observable behavior. (Not that it hasn't been tried, in the psychological framework of behaviorism; B. F. Skinner's *Verbal Behavior* is the most notorious example, now considered a spectacular failure).

Another attempt to grasp meanings is to relate them to **physically observable states of the brain**. If we accept the materialistic tenet that all the knowledge and the behavioral patterns of a person is directly expressed in his or her brain, a physical object that can be observed directly, then we should be able to find some physical correlate to the meaning of sentences that is observable. Modern methods of brain imaging indeed have allowed us to have first glimpses at the brain at work when processing and understanding linguistic information:

- PET (Positron Emission Tomography), which traces glucose use in tissue, and fMRI (functional Magnetic Resonance Imaging), which traces blood volume changes, have revealed that nouns and verbs are processed in different regions of the brain. The same holds for different types of nouns (e.g. natural objects, animals, artefacts) are processed in different regions, and for regular and irregular word forms. These findings are largely consonant with studies of brain lesions that impede the linguistic ability of people (aphasias). For example, there are aphasias that selectively affect verbs, or particular types of nouns.
- ERP (Event-Related Potentials), taken with electro-encephalograms, which tracks electromagnetic waves. One type of wave implied in linguistic processing is N400 (a wave that reaches maximal negative amplitude 400 ms after the onset of a stimulus). This can be observed for words that do not fit semantically into what is expected at a particular point in the sentence, for example, *She took the book and put it in the shelf* vs. *She took the book and put it in the shoe*.

However, the direct observation of semantic processing in the brain still has a long way to go. And most likely it will be very difficult to isolate it from other types of processing. One person will react to a given piece of information quite differently from another person. When we inform John, who has invested most of his retirement money in IBM stocks, that IBM stocks went down 50%, we will see a different change in brain states than if we inform Mary, who has not invested in IBM stocks.

Still another way of making sense of meanings is to assume that they are **entities of a particular, immaterial kind**. For example, if A tells B *The capital of Iceland is Reykjavik*, then A expresses a certain idea, and B grasps that very idea. Sounds plausible — but the question is: What are ideas? One answer that is normally associated with the philosopher Plato is that ideas exist somehow outside of the physically accessible world, but can nevertheless be grasped by humble human beings. This is not a very fashionable view nowadays, but for things like numbers and other mathematical objects and perhaps also linguistic meanings it is still as well- or ill-founded as any other.

The bad news for philosophers, then, is: We don't really know what meanings are. But the good news for linguists is: **We don't have to know what meanings really are!** As linguists, we are interested in the **relation between linguistic expressions and meanings**, and we can investigate this without being committed to very specific assumptions about the nature of meaning. (This was already clear to Al-Farabi, who stressed the point that logicians should investigate how thought relates to speech).

This situation should be familiar from other sciences. Take temperature; people were able to formulate very interesting theories about temperature that were on the right track without knowing what temperature really is (namely, kinetic energy). For example, they observed that if we bring a cold body and a warm body in physical contact, the cold body will get warmer, and the warm body will get colder. This was explained by an exchange of temperature stuff, called “phlogiston”. We know now that this theory was wrong, but it went a long way to explain phenomena related to heat.

Let us consider a particular linguistic example. It is unclear what the meaning of the word *sparrow* really is. But we can observe that it is related to the meaning of *bird*. In particular, the meanings of *finch* and *bird* stand in the same relation as the meanings of *tiger* and *cat*, or of *assassinate* and *kill*. (We will say that *finch* expresses a subconcept of *bird*). This is because everything that is a *finch* is a *bird*, everything that is a *tiger* is a *cat*, and every act of *assassinating* someone is an act of *killing* someone.

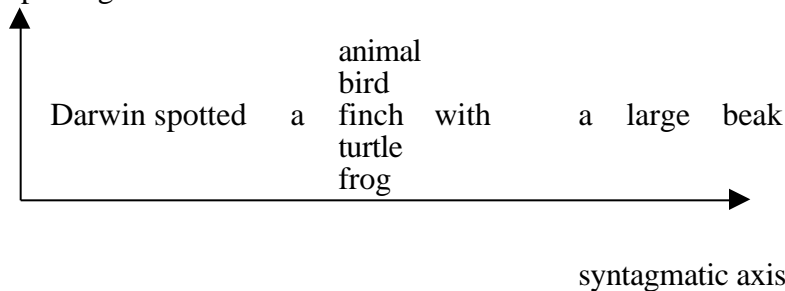
Similarly, we don’t know what the meanings of the words *a*, *beak*, *Darwin*, *finch*, *large*, or *spotted* are. We also don’t know what the meaning of the following sentence is:

(6) Darwin spotted a finch with a large beak.

But we know that the meanings of these words must be somehow related to the meanings of the sentence. And in general we can study the relation between the meaning of complex expressions and the meanings of their parts.

These two examples illustrate two basic aspects of the investigation of linguistic meanings that can be called the **paradigmatic** and the **syntagmatic** aspect. The question of how the meanings of words like *finch* and *bird* are related to each other illustrates the paradigmatic aspect of semantics.

(7) paradigmatic axis




The question of how the meaning of a complex expression is related to the meanings of the parts illustrates the syntagmatic aspect. Let us have a closer look at these two types of questions and how they can be addressed in linguistics.

1.3. Paradigmatic and Syntagmatic Meaning Relations

One important paradigmatic meaning relation that was mentioned already is the one that was illustrated by the pair *sparrow*, *bird*. We said that *finch* expresses a subconcept of *bird*. We also say that the word *finch* is a **hyponym** (literally, ‘undername’) of *bird* (and conversely, *bird* is a **hyponym** (literally, ‘overname’) of *sparrow*).

Another relationship of this kind is illustrated by the words *finch* and *nightingale*. What is a finch cannot be a nightingale, and vice versa. We say that *finch* and *nightingale* express mutually exclusive concepts. And the words themselves are said to be **contraries** of each other. Similarly, *pen* and *telephone* are contraries of each other, and are hyponyms of *artefact*. To illustrate, imagine that we represent objects on a plane, and encircle those objects that belong to the same class. I follow the convention to express the meaning of an expression α in double brackets, $[\alpha]$.



Furthermore, we have cases where two expressions mean exactly the same, for example *woodchuck* and *groundhog*. We say that these two words are **synonyms** of each other.

Another meaning relation can be best illustrated with adjectives. Take *rich* and *poor*, or *large* and *small*. These are contraries allright, as something that is rich is not poor (with respect to the same criteria), but in addition the members of these pairs denote opposites of each other. Such expressions are called **antonyms**.

We find similar meaning relations between larger expressions. Take the following two sentences:

- (9) a. Mary saw a bird.
b. Mary saw a sparrow.

Sentence (b) can be seen as a hyponym of sentence (a). That is, (b) expresses a more detailed piece of information than (a), hence expresses a subconcept of (a). Notice, in particular, that whenever (b) is true, then (a) is true as well. We call this relation between sentences **entailment**: Sentence (b) entails sentence (a). It is clear that this relation between sentences is inherited from the fact that *sparrow* is a hyponym of *bird*, because the two sentences differ only in the choice of these words. Other examples

- (10) a. Mary talked to Bill and John.
b. Mary talked to Bill.

- (11) a. Mary talked to Bill.
b. Mary talked to Bill or John.

- (12) a. Mary ate a hamburger.
b. Mary ate.

Sentences can also stand in the relation of synonymy. In this case, the sentences entail each other. Such sentence pairs are called **equivalent**. Examples:

- (13) a. Mary saw a groundhog.
b. Mary saw a woodchuck.


- (14) a. Mary kissed John.
b. John was kissed by Mary.

The meaning relations discussed so far are fairly simple. In particular, we don't have to understand what meanings really are in order to grasp them. Our natural abilities as speakers of a language allows us to determine these meaning relationships in a clear and unambiguous way.

Let us now turn to the **syntagmatic meaning relationships**, that is, the meaning relation between complex expressions and their parts. Consider a particularly simple example, the meaning of *happy person*. The question is, how is the meaning of

- (15) happy person

related to the meanings of *young* and *person*? Obviously, *young person* is a hyponym to *person*:



We suspect that this may be always the case when we combine an adjective with a noun. That is, whenever we have an adjective A and a noun N, then the combination A + N denotes a hyponym of N. This holds for the great majority of adjectives, but not for all of them — take *alleged murderer* and *fake 100\$ bill*.

But for many syntactic relationships our current tools cannot be applied. Take the sentence *Mary petted the elephant*. What are the meaning relations between, say, *Mary* and this sentence? We cannot really answer this question — a name like *Mary* and a sentence like *Mary petted the elephant* are of a quite different **type** of meaning, so that the question whether one is, say, a hyponym of the other does not make sense at all.

Perhaps it is more promising to first come up with a theory of meaning for sentences, and then derive the notion of meanings of the parts from the theory of sentence meanings?

1.4. Meanings, Truth Conditions, and Possible Worlds

Let us start with sentence meanings and try to derive the meanings of the basic expressions from them. What should we take to be the meaning of a sentence? Let us take simple **declarative sentences** (assertions) as the point of departure; other sentence types, like questions, commands or exclamations seem to be somehow derivative. Imagine a chess board. What does a sentence like the following one mean:

(16) The white queen is on a black field.

In order to say that a person understood the meaning of (16), this person must be, at least in principle, be able to determine under which circumstances (16) is true or false. That is, sentences are taken to express certain assumptions about reality, and if we understand the sentence correctly, we are able to say whether these assumptions are true or false. Hence the notion of **truth** seems to be critical for the meaning of assertion sentences. The version of semantics that takes truth conditions as the basis of all semantic judgements is called **truth-conditional semantics**.

If truth conditions are the way to the meanings of declarative sentences, the question is: How can we express the truth conditions of a declarative sentence? For example (16) the following will do:

(17) The sentence “The white queen is on a black field” is true if, and only if, the white queen is on a black field.

This formulation of truth conditions was proposed by the Polish logician Alfred Tarski in his 1935 paper “The Concept of Truth in Formalized Languages”.⁶ Of course, (17) sounds pretty trivial as it stands. This is because we use the same language, English, to describe the truth conditions of a sentence in this language.

Linguistics is a rather peculiar science in that respect. A physicist, for example, can use English, enriched with technical and mathematical terms, to describe the object of study, for example, subatomic particles. But a linguist has to use English (or another natural language), possibly enriched with some technical and mathematical terms, to describe English (or another natural lan-

⁶ Published in German as “Der Wahrheitsbegriff in den formalisierten Sprachen”, English translation in A. Tarski, *Logic, Semantics, Metamathematics* (Oxford University Press 1956, 152-278). Tarski’s original example was ‘*It snows*’ is true if and only if it snows.

guage). It is very important to make a distinction between the language that is used for the description, and the language that is being described. The first is called the **meta-language**, and the second, the **object language**. In (17), the part in quotation marks is an expression of the object language, and the rest belongs to the meta-language. In linguistics, object language expressions are typically not given in quotation marks, but in italics.

If the metalanguage and the object language are different, then the way how we give truth conditions illustrated in (17) does not sound trivial at all. In the following example, the object language is German.


- (18) The sentence *Die weiße Königin steht auf einem schwarzen Feld*
is true iff⁷ the white queen is on a black field.

You might ask at this point: Do we always have to give the truth conditions, and hence the meaning, of a declarative sentence, in terms of another sentence? That is, is there no other way to render the meaning of an expression by another expression? After all, this sounds as if doing **semantics** is, in essence, doing **translation** from one language into another!

This is a very important question. There is one way that appears to avoid reference to expressions. We can identify the meaning of a sentence with all the circumstances in which this sentence is true. Take our example and imagine all possible configurations of a chess board. Now, for some configurations the white queen is on a black field, for others the white queen is on a white field, and for still others, there is no white queen on the field at all. The meaning of (16) then can be given as follows:

- (19) [*The white queen is on a black field*]
= the chess configurations in which the white queen is on
a black field.

Of course, the number of possible chess configurations is incredibly large. Reducing the chess board to nine fields, and disregarding all the other figures, we can give the following graphical representation for our example. Note that, in contrast to what we had before, the chess boards do not stand for different objects, but for the same objects in different states.



We can generalize this idea for other examples as well. The German philosopher Gottfried Wilhelm Leibniz has introduced the notion of **possible worlds** in philosophy. A possible world is one way how the world could look like — one way of determining every aspect of the world, down to every minute detail. The real world is just one of these possible world (the so-called “actual” world).⁸ Of course, there are incredibly more possible worlds than possible chess configurations, but the basic idea is the same. Hence we can analyze the meaning of the sentence *Austin is the capital of Texas* as follows:

⁷ In logical parlance, *iff* stands for *if and only if*.

⁸ Leibniz wanted to show that the actual world is in fact the **BEST** possible world, in spite of all the evil in it; other possible worlds would have even more evil in them. This idea was ridiculed by the French philosopher Voltaire in his novel “Candide”, which in turn was turned into a musical by Leonard Bernstein, which has a theme song *It's the best of the best of the best of the best of the possible worlds*.

- (21) [*Austin is the capital of Texas*] = the possible worlds in which
Austin is the capital of Texas.

Now, our goal was to find a way of giving the meaning of a declarative sentence without resorting to other expressions. We introduced the notion of possible worlds for that. But when you look closer at the right-hand sides of (19) and (21) you will see that we again have English expressions there. This is because it's simply impossible to list all chess configurations in which the white queen is on a black field, let alone all possible worlds in which Austin is the capital of Texas. We have to describe them in one way or another, and this means, we have to use some language again. Hence, while the notion of possible world gives us a way, in principle, to break out of the "linguistic prison", for practical reasons we have to resort to a linguistic metalanguage once more. In this course we will not make much use of the notion of possible worlds, but simply use English to render the truth conditions of declarative sentences.

Let me add a couple of remarks. First, we are now able to distinguish between the meaning of two sentences even if they both happen to be true, as with the following examples:

- (22) a. Austin is the capital of Texas.
b. Reykjavik is the capital of Iceland.

This is because the possible worlds in which (a) holds are not quite the same as the ones in which (b) holds, even though they both contain the actual world.

Second, there are sentences that have the same meaning because their truth conditions are the same:

- (23) a. Austin is the capital of Texas.
b. The capital of Texas is Austin.
c. Texas has Austin as its capital.

The possible worlds in which these sentences hold are necessarily the same. But we also assume that they are synonyms, hence this is a welcome result. They differ in other respects: (a) is a sentence about Austin, (b) about the capital of Texas (which may be different from Austin!), and (c) is about Texas. This might influence the way how the information represented in these sentences is supposed to be stored. But the information itself is the same in both cases.

Third, we don't only have declarative sentences. There are also questions, commands, permissions, exclamations and other so-called **speech acts**. We will not deal with such sentences here. However, it is presumably wise to start with declarative sentences, because other sentence types presumably can be tackled once we have a theory of declarative sentences in place.

- (24) a. Is Austin the capital of Texas?
b. Give me the salt!
c. How terrible this painting is!

For example, the question (a) can be analyzed as a command of the speaker to the addressee to tell the speaker whether the corresponding declarative sentence *Austin is the capital of Texas* is true or not. The command in (b) can be seen as a command to make the sentence *The addressee gives the salt to the speaker* true. And (c) is an explanation of astonishment about the fact expressed by the sentence *This painting is terrible*.

1.5. Compositionality and Saturation

There is another sense in which the formulation of the truth condition in (17) or (18) is quite trivial, as it stands. To see this, consider a German textbook that consists just of truth conditions of sentences like the following:

- (25) a. *Die weiße Königin steht auf einem schwarzen Feld*
is true iff the white queen is on a black field.
b. *Es ist verboten, den Rasen zu betreten*
is true iff it is forbidden to step on the lawn.
c. *Jakob ist immer quer über die Gleise gegangen*
is true iff Jakob always crossed the tracks orthogonally.
d....

This is a phrase book, and a rather useless one, and not a way to learn German. What is left out in these definitions is how the meaning of a sentence relates to the parts of that sentence. The real task of semantics is to find out how the meanings of the subexpressions of a sentence lead, in a systematic way, to the truth-conditions of the sentence.

Our basic assumption will be, of course, that the constituents of a sentence have similar roles in other sentences. For example, our theory should tell us about the specific contributions of the words *schwarze* and *weißen* and how they lead to a different interpretation of the following sentence:

- (26) *Die schwarze Königin steht auf einem weißen Feld.*
'The black queen is on a white field.'

To see what our theory should achieve, consider the following, much simpler example:

- (27) Austin thrives.

The meaning of (27) is the collection or set of possible worlds in which it is the case that Austin thrives. So what about the meanings of the two constituents, *Austin* and *thrives*?

It seems relatively clear what *Austin* should mean: This is a name that refers to a particular entity, namely the capital of Texas.⁹ But then it is also clear what *thrives* should mean. Its meaning, when combined with the meaning of *Austin*, should give us the meaning of *Austin thrives*, i.e. the set of worlds in which Austin thrives.


- (28) a. [*Austin thrives*]
= all the possible worlds in which Austin thrives.
b. [*Austin*]
= the capital of Texas.
c. [*thrives*]
= something that assigns every entity x to the possible worlds in which x thrives (a so-called **function**).

The basic idea behind this analysis is that the meaning of a complex expression, like *Austin thrives*, should be computable from the meanings of the parts, here *Austin* and *thrives*. This is called the **principle of compositionality**. It was first proposed, although in a rather implicit way, by Frege, and is also known as **Fregean Principle**.

The principle of compositionality can be applied to more complex expressions as well. For example, the following sentence has the indicated phrase structure, which can be independently motivated by syntactic tests:

- (29) Austin is in Texas.

⁹ This is of course a gross oversimplification. It also applies to the Texas hero Stephen Austin, to the English philosopher of language John Austin, to the city of Austin in Minnesota etc. But let us work under the assumption that a particular context of use of a name identifies exactly one referent for that name.



We assume that the meaning of the whole sentence, S, is a function of the meanings of its immediate parts, the noun phrase NP and the verb phrase VP. In turn, the meaning of the VP is a function of its immediate parts, the verb V and the prepositional phrase PP. Also, the meaning of the PP is a function of its immediate parts, the preposition P and the noun phrase NP.

We already see with this simple example that the principle of compositionality will play a major role in linguistic semantics. Independent motivation for it comes from the following consideration. We can express an infinite number of different meanings with a natural language. But the means of a natural language — the set of words and the set of syntactic constructions — is fairly limited. The number of words (or rather morphemes) that a typical speaker knows ranges in the magnitude of tens of thousands, and the number of syntactic and morphological constructions perhaps in the magnitude of hundreds. This limitation is necessary because languages must be **learnable**, and a language with a very large or even infinite inventory of morphemes or rules could not possibly be learned by humans with limited brain and limited time. But due to compositionality the limited building blocks of language do not restrict speakers in their expressiveness — they can express a much greater, actually an infinite number of pieces of information. This was first seen by the philosopher of language Wilhelm von Humboldt (1767-1835), who observed that language allows speakers “to make unlimited use of limited means”.

In particular, we typically do not learn the meanings of complex expressions, like *Austin thrives*. All that we have to learn is the meanings of basic expressions, like *Austin* and *thrives*, and how to compute the meaning of a complex expression from the meanings of its parts.

Frege expressed these ideas in the following words:¹⁰

“It is astonishing what language accomplishes. With a few syllables it expresses a countless number of thoughts, and even for a thought grasped for the first time by a human it provides a clothing in which it can be recognized by another to whom it is entirely new. This would not be possible if we could not distinguish parts in the thought that correspond to parts of the sentence, so that the construction of the sentence can be taken to mirror the construction of the thought. [...]”

If we thus view thoughts as composed of simple parts and take these, in turn, to correspond to simple sentence-parts, we can understand how a few sentence-parts can go to make up a great multitude of sentences to which, in turn, there corresponds a great multitude of thoughts.”

Let’s come back to our original example, *Austin thrives*. We have assumed that the name *Austin* refers to an entity, and the whole sentence refers to a set of possible worlds; Frege would have said, to a thought. In a sense, name meanings and sentence meanings are **complete** or, as Frege termed it, **saturated**. On the other hand, we claimed that the meaning of *thrives* is something that has to be combined with another meaning (e.g. the meaning of *Austin*) to produce a complete meaning (namely, a set of possible worlds). Hence the meaning of *thrives* is **incomplete** or **unsaturated**. Technically, the meaning of *thrives* is a **function** from entities to sets of possible worlds. Frege also claimed that “logical combinations of parts into a whole is always a matter of saturating something unsaturated”.

¹⁰ In “Logische Untersuchungen. Dritter Teil: Gedankengefüge.”, 1918.

To sum up: Even though we still don't know what meanings are, we have already learned a lot about them by the sheer fact that they can be expressed within natural language. We made it plausible that the meaning of complex expressions is a result of specific combinations of the meanings of the parts. And we saw that it is plausible to think of such meaning combinations as involving the combination of unsaturated with saturated meanings.

More specifically, we have argued that a good model for the meanings of assertion sentences are sets of possible worlds, namely the set of worlds in which these sentences are true. Hence, the notion of truth plays a crucial role in this type of semantics.

1.6. Some Elementary Semantic Notions

Before we start with developing more specific notions of truth-conditional semantics, we should discuss a few semantic notions that have not been mentioned so far.

First, the phenomenon of **ambiguity**. It describes cases in which one and the same constituent has more than one meaning. For example, the word *pen* has at least three meanings that seem to be completely unrelated: (a) a writing instrument¹¹, (b) an enclosure for animals, and (c) a prison (not related to (b) but short for *penitentiary*). These meanings just happen to be expressed by the same phonological form is used to express widely divergent meanings. The meanings are so different from each other that it is quite unlikely that any confusion arises.

Ambiguity should be distinguished from **polysemy**. For example, *pen* may stand for a writer (e.g., *a hired pen*), or for a writing style (e.g., *she wrote with a very witty pen*). Also, it may denote the act of writing with a pen, as in *he penned the letter*. These are clearly extensions of the first meaning of *pen*. The border line between true ambiguity and polysemy is often not all that clear. For example, *pen* also is used to denote a female swan (which may be related to the first meaning of *pen*, as feathers were used as pens), and to denote the longish, hard internal part of a squid, probably because of the pen-like shape.


Ambiguity should also be distinguished from **vagueness**. Terms like *several* or *many* or *tall* are vague. We may be unsure about whether to apply a term like *tall* to a person because we do not really know which standards to apply. An American adult male of 5 feet may not be considered tall; a 10 year old boy or a Rwandan pygmy of 5 feet would probably be considered tall. Ambiguity and vagueness may occur together. The adjective *hot* is ambiguous between “high temperature” and “high degree of spicyness” (and of course a range of other meanings!). In either one of these meanings it is vague. It is characteristic for ambiguity that it does not travel in translations. For example, the first meaning of *pen* can be rendered in German as *Schreibfeder* or *Füller*, and the second meaning as *Pferch*. In contrast, German *gross* is as vague as English *tall*.

So far we have discussed cases of ambiguity of single words. Word ambiguity may be inherited by sentences in which these words occur; for example, *Mary saw a pen* is ambiguous in the same way as *pen* is. But we find a new source of ambiguity in sentences that relates to the way how words are combined. For example, the following sentence may mean two quite different things:


- (30) Mary saw that gasoline can explode.
 (i) Mary saw that that container of gasoline exploded.
 (ii) Mary realized that gasoline may explode.

This type of ambiguity results from the fact that one string of words (30) can have two distinct syntactic structures. For our purposes they can be characterized in the following way:

¹¹ By the way, you might think that *pen* and *pencil* are related: After all, they are semantically similar, and their form is similar, too. Not so; *pen* in the meaning ‘writing instrument’ comes from Latin *penna* ‘feather’, whereas *pencil* goes back to Latin *penicillus* ‘brush’, literally ‘little tail’ (which in turn is related to *penis*!) Historical meaning relations like that are investigated in **etymology**. Unfortunately, we will not have the time to talk about the historical aspect of semantics, like the principles behind meaning change, in this course.




'Mary saw that that container of gasoline exploded'




'Mary saw that gasoline has the propensity to explode'

These two ways of combining words lead to distinct meanings. Such ambiguities are called **structural ambiguities**. In our example, a word ambiguity plays an essential part in the structural ambiguity — the ambiguity of the word *can* as “container” and as a modal auxiliary “be able to, may”. But there are also pure cases of structural ambiguities where no word ambiguities seem to be at play, as in the phrase *old men and women* and in the sentence *Mary saw the man with the telescope*.

(31) old men and women

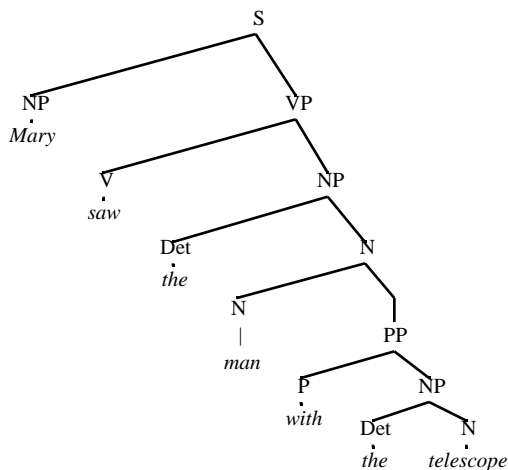


'old men and old woman'




'old men, and women'

(32) Mary saw the man with the telescope.



'Mary saw the man who had a telescope'



'Mary saw the man with the help of the telescope'

There are ambiguities which may not even be directly related to differences in the syntactic structure of sentences (this depends on how these theories see the way how syntax and semantics relate to each other — more on this in the section on syntactic scope and semantic ambiguity).

- (33) All that glitters isn't gold.
 (i) 'Not everything that glitters is gold.'
 (This is the idiomatic meaning of this saying.)
 (ii) 'For everything that glitters it holds that it is not gold.'
 (= nothing that glitters is gold).

Another type of meaning variation that has to be distinguished from ambiguity is the one that we find with expressions like *you*, *here*, or *yesterday*. What *you* means depends on the communication situation; it will denote the addressee of a communication situation, whoever this is. What *yesterday* means will depend on the time at which the communication takes place. Expressions that show this dependence on the context of the communicative situation are called **context-sensitive**. Context-sensitivity is a pervasive phenomenon of natural language. For example, tenses such as past, present and future are context sensitive.

Let us leave ambiguity and related notions and turn to what we may call aspects of meaning. One such aspect has been discussed under the terms **denotation** vs. **connotation**. So far, we were concerned with the truth-conditional or denotational aspect of meaning only. The connotational aspect describes certain attitudes a speaker has towards an entity or a state of affairs. For example, the words *dog* and *mongrel* denote the same objects, namely, canines, but they differ in their connotation — the second is a **pejorative** word that expresses a negative attitude of the speaker towards the object denoted by the expression. Similarly, the following two sentences are true in exactly the same situations, and differ only because (b) indicates a negative attitude of the speaker towards the expressed fact.

- (34) a. It is raining.
 b. Unfortunately, it is raining.

Often certain components of meaning have backgrounded status — they are **presupposed**. For example, a sentence like

- (35) Bill's sister drives a Miata.
 presupposed: 'Bill has a sister.'
 asserted: 'She drives a Miata.'

contains the information that the speaker has a sister, and that she drives a Miata. But the status of these two pieces of information is quite different. In particular, the first piece, that the speaker has a sister, is treated as taken for granted or presupposed, and only the second piece is asserted as new information. Another example:

- (36) John stopped smoking cigarettes.
 presupposed: 'John used to smoke cigarettes.'
 asserted: 'John doesn't smoke cigarettes now.'

This sentence says that John used to smoke cigarettes, and that he does not smoke them anymore. Again, the first piece of information is presupposed. There are various tests that allow us to distinguish presuppositions from other parts (often called assertions). For example, presuppositions survive if we put a sentence into a question, or prefix it with a modal element like *maybe*. The following sentences still carry the information that John used to smoke cigarettes, but not that he does not smoke them anymore:

- (37) a. Did John stop smoking cigarettes?
 b. It may be that John stopped smoking cigarettes.

Another aspect of meaning of an expression is the distinction between the **literal meaning** of an expression and the **implicatures** that are generated by using the expression at a particular occasion. For example, the following text as a recommendation letter will be read in a way that indicates that the writer does not recommend Mr. Jones for the job of a philosophy professor:

To: Department of Philosophy, XX University

To whom it may concern:

Mr. Jones' has visited my classes regularly, and his handwriting is neat.

Yours sincerely,
Professor Paul Grice

Obviously, the fact that the letter leaves out many things that one would expect from a recommendation letter conveys an important meaning in its own right — that the writer probably does not consider Mr. Jones to be suitable for the job of a philosophy professor.

Another type of implicature can be illustrated with the following example:

(38) John has seven cats.

Normally, this sentence is interpreted as saying that John has exactly seven cats (that is, he does not have eight cats). But this is not part of the literal meaning. Observe that the sentence could be true even if John has, as a matter of fact, eight cats — if he has eight cats, then it follows that he has seven (and six, five ...) cats. Observe, in particular, that the following sentence does not express a contradiction:

(39) a. John has seven cats, in fact, he has eight.
b. John has seven cats, perhaps even eight.

This contrasts with cases in which the second clause affects the literal meaning of an expression, which leads to a blatant contradiction:

(40) *John owns a Toyota; in fact, he does not own a car.

It is generally seen as a property of implicatures that they can be **cancelled**. Hence it is only an implicature of (38), and not part of the literal meaning, that John does not have more than seven cats.

This latter type of implicature is very common in natural-language communication. According to one influential theory of implicatures due to the philosopher Paul Grice it arises in the following way. It is assumed that a speaker tries to be maximally informative (without being excessively verbose, of course). Now, the sentence (i) *John has eight cats* is more informative than the sentence (ii) *John has seven cats*, as (i) entails (ii), but is not entailed by (ii). Hence a speaker that utters (ii) instead of (i) must have some reason for this choice. The most plausible reason is, of course, that the speaker assumes that (i) is false. Hence the hearer can conclude that (i) is false, that is, that John does not have more than seven cats. But this is not because it is part of the literal meaning (ii). Rather, the hearer can draw this inference because he or she assumes that the speaker is maximally informative.

1.7. Conclusion

In this section we discussed the subject matter of semantics — the study of meaning of natural language expressions. While it is unclear what meanings actually are, we can build models of meaning that capture the essential understanding of the meaning relations that speakers of natural languages have. It turned out useful to start with the meaning of declarative sentences, which are related to the notions of truth. From that, the meaning of the constituents that make up the meaning of a sentence can be derived, following the principle of compositionality. We also discussed a number of semantic notions, such as ambiguity, polysemy, presupposition and implicature.