Reports of the Keio Institute of Cultural and Linguistic Studies 53 (2022), 223~237

Chapter 2
Semantics as Science
Connecting Sentences
Introduction to Semantics for Non-native Speakers of English

Christopher Tancredi

2.1 Introduction

The study of semantics is one part of the study of language, also known as **linguistics**. The goal of linguistics is to understand human language. The goal of semantics is more limited: to understand how literal meaning works in language. There are many ways of studying language and meaning. In this textbook we approach linguistics as a science.

What does it mean to be a science? A science **observes**, or looks at, how things behave in the world, and tries to give a **theory**, or explanation, of that behavior. Importantly, the theory has to make **predictions**, expectations about how things will behave that can be tested and possibly shown to be wrong. If a theory's predictions match observations, the theory is **supported** by the observations. If the observations go against the predictions, the theory is **falsified**, or shown to be wrong, by the observations. When a theory is falsified, it needs to be changed or replaced by a better theory.

Two types of observation are widely used to test the predictions of semantic theories. The first is observations about whether a sentence is true or false in some situation. The second is observations about how two sentences relate to each other. For two declarative sentences, the most important relations are entailment and contradiction:

Observed relations:

Entailment: A sentence S_1 **entails** another sentence S_2 just in case the truth of S_2

follows from the truth of S_1.

Contradiction: A sentence S_1 **contradicts** a sentence S_2 just in case it is impossible for both S_1 and S_2 to be true together.

Examples of these relations are given below:

John came and Mary left	**entails**	*John came*
John came	**does not entail**	*John came and Mary left*
John came and Mary left	**does not entail**	*Mary does not like John*
John came	**does not entail**	*John didn't come*

John came and Mary left	**does not contradict**	*John came*
John came	**does not contradict**	*John came and Mary left*
John came and Mary left	**does not contradict**	*Mary does not like John*
John came	**contradicts**	*John didn't come*

When two sentences do not contradict one another, we also say that they are **compatible**.

In Chapter 1, we said that expressions have a sense, and that the sense picks out a denotation, or extension, at every possible world. The pairing of worlds with extensions we called the intension of an expression. Intensions are part of the theory of semantics. They make it possible to **analyze**, or represent, observed entailments and contradictions:

Analysis of entailment:
A sentence S_1 *entails* a sentence S_2 just in case there are no worlds *w* such that the intension of S_1 contains <*w*,True> and the intension of S_2 contains <*w*,False>.

This analysis uses truth values at possible worlds to analyze the intuitive notion of one sentence following from the other. According to this analysis, S_1 *entails* S_2 just in case the intensions of S_1 and S_2 contain only the patterns in A, B or C, not the pattern in D:

	Intension of S_1 contains:	Intension of S_2 contains:
Pattern A	<*w*,True>	<*w*,True>

Pattern B	<*w*,False>	<*w*,True>
Pattern C	<*w*,False>	<*w*,False>
Pattern D	<*w*,True>	<*w*,False>

A similar analysis can be given of contradiction:

Analysis of contradiction:
S_1 *contradicts* S_2 just in case there are no worlds w such that the intensions of S_1 and of S_2 both contain <*w*,True>.

According to this analysis, S_1 *contradicts* S_2 just in case the intensions of S_1 and S_2 contain only the patterns in B, C or D, not the pattern in A:

	Intension of S_1 contains:	Intension of S_2 contains:
Pattern B	<*w*,False>	<*w*,True>
Pattern C	<*w*,False>	<*w*,False>
Pattern D	<*w*,True>	<*w*,False>
Pattern A	<*w*,True>	<*w*,True>

The definition of entailment does not tell us why *John came* follows from *John came and Mary left*. It does, however, show an effect that holds just in case the one sentence follows from the other: that every world in which *John came and Mary left* is true is a world in which *John came* is true. The definition of contradiction also does not tell us why *John came* contradicts *John didn't come*. It only shows an effect that holds just in case the sentences are contradictory: that there is no world in which both *John came* and *John didn't come* are true. To answer the *why* question, a theory of semantics is needed.

2.2 Logic: The Formal Basis for Semantics

Before we start to build a theory of semantics, we need to first deal with an important issue: natural language can be unclear. If we try to describe the meaning of a sentence using natural language, then the description itself can end up equally unclear. To illustrate one problem that can arise, consider the following sentence:

Visiting relatives can be boring.

This sentence is **ambiguous**: it has two different meanings. It can say something either about the activity of going to visit relatives or about relatives who visit people. How do we describe this difference in meaning? Saying that the sentence can either mean that visiting relatives can be boring or that visiting relatives can be boring obviously doesn't help. This just uses the same ambiguous sentence twice. It does not help us to see what the two meanings are.

In practice we get around this problem using two methods. The first is to **paraphrase** the intended meaning in English using sentences that are unambiguous. We could say, for example, that *Visiting relatives can be boring* means either that it can be boring to visit relatives, or that relatives who visit can be boring. Though useful, giving English paraphrases does not help answer why a sentence has the paraphrases it has.

The second method for showing ambiguity is to translate the sentence into a language that contains no ambiguity. If the sentence is ambiguous, it will have two different translations in one of these no-ambiguity languages. No natural language lacks ambiguity. However, many languages of **logic** do.

How do we translate an English sentence into a sentence, or **formula**, of logic? We could do the translation intuitively, finding a formula of logic that is close enough in meaning to the English sentence. This would be similar to the process of giving an English paraphrase of the sentence. It would not answer the question of why that translation is a good translation, though. Fortunately, we can do better. In Chapter 1, we introduced the idea of compositionality. That was the idea that the meaning of a complex expression is composed out of the meanings of its parts. If we do our translation compositionally, then this will start to answer the *why* question.

2.2.1 Propositional Logic

A logic for doing semantics should give us the tools we need to interpret any expression of English. We build toward such a logic in steps. We start with a very simple logic, **propositional logic**, that will only be useful for interpreting a small number of English expressions. Later chapters will add to that logic in order to interpret more and more English expressions.

In propositional logic, the basic expressions -- the expressions whose interpretation is taken as given -- are **propositions**. Propositions are the meanings of **declarative** sentences. If a sentence is ambiguous, it has two or more propositions as meanings. If two sentences are **synonymous**, they have a single proposition as their shared meaning. An **active** sentence like *John saw Mary* and its **passive** version *Mary was seen by John*, for example, are generally taken to be synonymous.

Propositional logic divides propositions into basic propositions and complex propositions. A complex composition is one that has one or more propositions as parts. A basic one does not. The meaning of the sentence *John showed Mary a picture* is basic in this sense. That meaning is composed of different parts, the meanings of the words and phrases. However, none of those parts is a proposition: they are not things that can be true or false. The meaning of the sentence *John came and Mary left*, in contrast, is complex. It has two parts that are propositions: the meaning of *John came*, and the meaning of *Mary left*.

Propositional logic aims to explain how the meaning of a complex proposition depends on the meanings of its parts. The parts that matter to us here are propositions and **logical connectives**, or **logical operators**. We will introduce four logical operators below, together with the English words they are used to analyze.

2.3 The Semantics of *and*, *or*, *if*, and *not*

Syntactically, the English expressions *and*, *or*, and *if* combine two sentences into one. Given that the meaning of a sentence is a proposition, this means that semantically these expressions combine two propositions into one. Examples of these expressions are given below.

> *John came* ***and*** *Mary left*.
> *John came* ***or*** *Mary left*.
> ***If*** *John came, Mary left*.

English *not*, expressing **negation**, does not combine syntactically with a sentence. It rather combines with a verb phrase, or perhaps with an **auxiliary** verb like *do*, *be* or *have*. However, the effect of negation is to negate a sentence. The sentence below, for example, negates the sentence *John will come*.

John will ***not*** *come.*

Below we analyze each of these expressions.

2.3.1 Conjunction

The English expression *and* expresses **conjunction**. The standard logical analysis of *and* is as **logical conjunction**, represented with the symbol **&**. Logical conjunction joins two propositions together to create a new, complex proposition. Let p and q be two propositions. Logical conjunction creates from these the new proposition $p \,\&\, q$.

The interpretation of logical & can be described entirely in terms of **truth values**:

> If p is true and q is true, then $p \,\&\, q$ is true. Otherwise, $p \,\&\, q$ is false.

This interpretation can be represented by a **truth table**, where each line represents one way things could possibly be.

	p	q	$p \,\&\, q$
A	T	T	T
B	T	F	F
C	F	T	F
D	F	F	F

Analyzing English *and* as logical *&* explains our intuitions about the role of *and* in the sentence *John came and Mary left*. If *John came* is true and *Mary left* is true, then the sentence as a whole is true. In all other cases, the sentence as a whole is false. This is exactly what the analysis predicts.

Differences between (logical) *&* and (English) *and*

Analyzing *and* as *&* works whenever *and* connects two declarative sentences. However, English *and* has many uses that do not connect sentences. As seen below, *and* can connect two names, two verb phrases, and two prepositions among other things:

(A) *John saw [Mary and Bill]*
John [saw Mary and heard Bill]
John walked [to and from] the store

Logical & cannot do the same. It was defined specifically to connect propositions. The interpretations of names, verb phrases and prepositions are not propositions. That is, they are not the kinds of things that can be true or false. Using x' to mean *the meaning of x*, none of the following expressions is well formed in a logic that defines & as it was defined above. (An asterisk, *, is used to mark something as ill-formed.)

**Mary' & Bill'*
**[saw Mary]' & [heard Bill]'*
**to' & from'*

This is a challenge to the analysis proposed. How can the challenge be met?

There are three general approaches that can be taken to meet the challenge. The first is to analyze *and* as ambiguous, having & as one of its meanings and something else as another. The second approach is to give *and* a single meaning that can apply to other kinds of meanings in addition to propositions. The third approach is to leave the meaning of *and* unchanged and to argue that *and* is actually connecting sentences in examples like those in (A).

Which of these approaches should we take? There is no right answer to this question. However, there are some principles that we should try to follow when trying to answer it. One of these is Occam's Razor:

Occam's Razor
Entities should not be multiplied without necessity.

Occam's Razor suggests that between the first two options, we should prefer the second to the first since it uses one meaning instead of two. Occam's Razor does not choose between the second option and the third.

Another principle does not have a standard name. I call it *stick to your guns*:

Stick to your Guns:
Do not give up an analysis until you are required to do so.

In this case, the analysis we were given always interprets *and* as *&*. Following this principle, we should prefer the third option over the first two. The first option only interprets *and* as *&* sometimes, not always. The second option never interprets *and* as *&*. Only the third option interprets *and* always as *&*. Of course, this third option may end up not working in the end. In that case, we will have to try a different option. Until we can show that the third option fails, however, it is our best option to pursue.

The third approach has been given a name. It is called **conjunction reduction**. On this approach, the sentences in (A) are more complicated than they appear. In particular, they have structures like the following.

(B) *[John saw Mary and ~~John saw~~ Bill]*
[John saw Mary and ~~John~~ heard Bill]
[John walked to ~~the store~~ and ~~John walked~~ from the store]

The parts that are crossed out are taken to be parts of the sentence. They are interpreted but not spoken. Adopting the structures in (B) for the sentences in (A) makes it possible to keep our analysis of *and* as *&* for these sentences. The full analysis of the first sentence in (A) ends up as follows, for example:

Spoken form:	*John saw Mary and Bill*
Interpreted form:	*John saw Mary and John saw Bill*
Logical interpretation:	*[John saw Mary]' & [John saw Bill]'*

A second challenge to analyzing *and* as *&* comes from ordering. Often, when we connect two sentences with *and*, we imply a fixed order of events. This can be seen in the following examples.

John drank beer and he drove home.
John drove home and he drank beer.

The first sentence implies that John drank beer before driving home. The second sentence implies the opposite order: that John went home before drinking beer. Logical *&*, however, says nothing about order. Translating *and* as *&*, then, does not directly explain this ordering.

Here again we have several options for how to account for the observations. For example, we could analyze *and* as ambiguous between *&* and a meaning that includes ordering. Or we could argue that sometimes *and* comes with an unspoken *then*. Or we could argue that the ordering comes from the order in which the sentences are spoken. If either the second or the third option works, the observation about ordering cannot be used to argue against interpreting *and* as *&*.

A third challenge to analyzing *and* as *&* comes from sentences like the following:

> *John, Mary and Sue surrounded Bill.*
> *Mary, Sue and Bill gathered in the park.*

These sentences look similar to the sentences analyzed as conjunction reduction. However, here that analysis does not work. It predicts that the first sentence should mean that John surrounded Bill, Mary surrounded Bill and Sue surrounded Bill. This is a possible way of understanding the sentence, but it is not the way most people will understand it. On the normal understanding of the sentence, no single person surrounds Bill. Only the three of them together do so. For the second sentence, a conjunction reduction interpretation is unacceptable. It would imply, for example, that Mary gathered in the park. However, gathering in the park is not something that a single person can do. It takes two or more to gather.

This last challenge is generally taken to show that *and* cannot be interpreted only as *&*. An interpretation is needed that allows *and* to directly combine individuals as well as propositions, and *&* cannot do that. This new interpretation could be an additional interpretation of *and*, or it could replace *&* altogether.

It would be nice to be able to give an interpretation for *and* that works for all of the sentences we have looked at. However, interpretations that do so all require tools beyond propositional logic. In fact, they generally require tools of logic beyond those that will be introduced in this textbook. To avoid having to develop these advanced tools, we will treat *and* as ambiguous between & and some other interpretation. For the rest of this textbook, the only examples we look at that use *and* will be examples where *and* can be interpreted as &.

2.3.2 Disjunction

English *or* expresses **disjunction**. It is standardly analyzed as logical disjunction, represented by the symbol $\vee$. Logical disjunction works like logical conjunction: it combines two propositions p and q into a new proposition, $p \vee q$, and the meaning of $\vee$ can be given completely by a truth table:

	p	q	p ∨ q
A	T	T	T
B	T	F	T
C	F	T	T
D	F	F	F

As with *and*, there are challenges to interpreting English *or* as logical $\vee$. The first challenge is parallel to the first challenge facing *and*: English *or* can combine with expressions having many different kinds of meanings, while logical $\vee$ can only combine with propositions.

John saw Mary or John saw Bill
John saw Mary or Bill
John saw Mary or heard Bill
John walked to or from Tokyo

The solutions to this challenge are the same as with *and*: change the meaning of *or*, or analyze the sentences as examples of conjunction reduction.

A second challenge to the interpretation of *or* as $\vee$ comes from cases in which the sentences connected by *or* are both true. To see the challenge, imagine a multiple-choice test where the teacher gives a hint:

The answer to 1 is a or the answer to 2 is b.

Suppose it turns out that the answer to 1 is a *and* the answer to 2 is b. A student would be right to complain that the teacher misled them. If *or* is interpreted as $\vee$, however, the sentence should be true. Why would it be misleading?

As before, there are different options for meeting this challenge. One option is to change the meaning of $\vee$ so that $p \vee q$ is false when p is true and q is true. This would be a change from an **inclusive** meaning -- including the case when both p and q are both true -- to an **exclusive** meaning -- excluding that case. Then what the teacher said would be predicted to be false. But is it false? Intuitions are not clear. The sentence is clearly misleading, but that is not the same as being false. There are, however, other sentences where an exclusive meaning clearly makes a wrong prediction, like the following.

> If a person has long eyelashes, at least one of their parents does too. Since John has long eyelashes, *his mother has them or his father has them*.

If *or* is interpreted exclusively, then the sentence in italics above entails that only one of John's parents has long eyelashes. That is, it is false if both of John's parents have long eyelashes. We do not understand the sentence in this way, though. We rather understand it to mean that at least one of John's parents has long eyelashes. We accept it is possible that they both have long eyelashes.

A second option for facing this challenge is to analyze *or* as ambiguous. One of its meanings could be *inclusive* and the other *exclusive*. A misleading use of *or* could then be one where the speaker means one of these where most people would expect them to mean the other.

The *stick-to-your-guns* option is to take or to mean inclusive $\vee$ and to give a non-semantic explanation for cases where *or* is understood exclusively. In many cases, a common sense explanation is possible. Suppose John was at a stoplight. I tell you *he went straight or he turned right*. This is a case where only one option can be taken. Going straight means not turning right, and turning right means not going strait. Even if *or* means $\vee$, we know it is not possible for both to happen. The situation allows only the two possibilities allowed by exclusive *or*, even if the semantics allows for an extra possibility.

2.3.3 Conditionals

English *if* is used to make **conditional** sentences. Its standard analysis is as **material implication**, represented by the symbol →. Material implication works like & and ∨: it combines two propositions into a new proposition, and its meaning can be given by a truth table:

	p	q	p → q
A	T	T	T
B	T	F	F
C	F	T	T
D	F	F	T

An example where this looks like the right analysis of *if* is the following. Imagine students taking a multiple-choice test. The teacher gives the students a hint. She says, *If the answer to 1 is a, the answer to 2 is b*. Did the teacher say something true or false? That depends on the situation. If the answer to 1 is *a* but the answer to 2 is something other than *b*, then what she said is false. In all other situations, though, what she said is true. This is exactly what we should expect if *if* is interpreted as →.

Even though interpreting *if* as → works in this example, many people feel that this interpretation is wrong. Intuitively, *if* usually suggests a stronger connection between two things. The sentence *If it's raining, the streets are wet*, for example, suggests that rain *causes* the streets to be wet. Similarly, the sentence *if John's umbrella is wet, it's raining* suggests that John's umbrella being wet is a reason to conclude that it is raining. Material implication, though, says nothing about these connections. Do we need a different interpretation of *if* to explain this connection? Let's try sticking to our guns first.

In the truth tables above, p & q, $p \vee q$, and $p \rightarrow q$ are all true when both p and q are true. If a speaker believes that p and q are both true, then, what should they say? Suppose for example that the speaker believes it's raining and the streets are wet. Should they say *It's raining and the streets are wet*, *It's raining or the streets are wet*, or *If it's raining, the streets are wet*? We have a clear intuition here. The sentence with *and* is clearly the best. Why? Because it describes the situation exactly as the speaker believes it to be. It is true when the sentences it combines are both true and false otherwise. The other two sentences with *or* and *if* leave open other possibilities. This suggests that our choice of what to say depends on more than just whether the sentence we say is true. It also depends on whether what the sentence says matches what we believe about how things are. The sentence has to be true of situations we think exist and it has to be false of situations we think do not exist.

Suppose that the sentence a speaker uses has to match the speaker's beliefs. Then, if *if* means →, in what situations should the speaker use the sentence *if it is raining, the streets are wet?* It should be used in a situation like the following: the speaker believes that either it's raining and the streets are wet, or it's not raining and the streets are wet, or it's not raining and the streets are dry, but he does not know which of the three is correct. How can the speaker be in this state? Not by looking outside. Looking outside he might learn whether it's raining. He might also learn whether the streets are wet. He cannot, however, learn only that it is not the case that it is raining and that the streets are dry.

One very common way to be in this state is by believing that rain causes the streets to be wet. Having this belief is then a good enough reason to say *if it's raining, the streets are wet*. Can we conclude that a speaker who says this sentence has this belief? Not with certainty. There are other possible beliefs the speaker could have that would match what they say just as well. However, in many cases this might be the best conclusion we can come to. In those cases, we have a good reason to think that the speaker is trying to communicate that belief to us.

2.3.4 Negation

In English, negation is expressed with the word *not*. Its standard interpretation is logical negation, represented by the symbol ~. Logical negation differs from &, ⋁, and → in that it only combines with one proposition, not two. Its meaning can be described in a truth table: *~p* has the opposite truth value as *p*:

	p	~p
A	T	F
B	F	T

Unlike *and*, *or*, and *if*, English *not* does not combine directly with a sentence. We say *John did not come*. We do not say *Not John came*. This makes the translation from

English into logic a little complicated. We cannot simply compose the interpretation of *not* with the interpretation of *John did come*. *John did come* is not a **constituent**, that is it is not a single expression, in the sentence *John did not come*.

While this is an important problem, we do not yet have the tools to deal with it here. We will develop those tools in the next two chapters, however. They will allow us to compose *not* with *come*, compose *not come* with *did*, and finally compose *did not come* with *John* to produce the logical translation ~*[John did come]'*. For now, we will ignore this problem and act as if *not* combines with a sentence.

One challenge to analyzing *not* as ~ comes from sentences like the following.

> *John didn't kick the bucket*. *He passed away*.

The English expressions *kick the bucket* and *pass away* both mean the same thing: *die*. If *not* can only be translated as ~, it looks like these sentences should say *John didn't die*. *He died*. This is not how we understand them, though. We instead understand them as saying that *John kicked the bucket* is not an appropriate way to say *John died*: *He passed away* is the appropriate way.

When *not* is used to comment on *how* something is said rather than to negate *what* is said, it is called **metalinguistic negation**. Can metalinguistic negation be translated as ~? If the rest of the sentence is interpreted normally, the answer is clearly "no". However, if the rest of the sentence can be interpreted as indicating appropriateness, then the answer is "yes". For example, if *John kicked the bucket* can be interpreted as meaning *"John kicked the bucket" is the appropriate way of saying John died*, then *not* will negate this meaning. The resulting interpretation will say *"John kicked the bucket" is* ***not*** *the appropriate way of saying John died*. And this is something that we can get by interpreting *not* as ~.

2.4 Summary

Analyzing English *and*, *or*, *if* and *not* as logical &, ∨, → and ~ faces challenges. In some cases, it looks like the challenges can be met. In other cases, some other analysis might be needed. Either way, however, the logical operators serve a useful purpose. They can be used to give unambiguous interpretations as potential interpretations of

English expressions. As we will see in later chapters, this includes expressions that do not involve the words *and*, *or*, *if* or *not*. As symbols of logic, the interpretations given to $\&$, $\vee$, $\rightarrow$ and $\sim$ are neither right nor wrong. They simply are what they are. The question of right or wrong, however, does arise for the semantics. Do the formulas these symbols are used in give the right meaning of an English expression? We cannot give a general answer to this question, though. We have to answer this question on a case-by-case basis.