Lecture 2

Connectives

In the previous lecture, we started using capital letters—e.g. P, Q, A, B, C, etc.—to represent arbitrary true/false sentences. We'll call these **atomic sentences**, or **atoms**, since we will not try to break them up into parts. When building truth tables, the atomic sentences were grouped to the left and used to build up the space of possible worlds. On the right we listed complex sentences, and noted that if we were going to be able to fill in the truth table for these complex sentences, their truth or falsity had to depend entirely on the truth or falsity of the considered atomic sentences. In the main example, we took P and Q to be our atoms, and considered the truth tables for the complex sentences "P and Q", whose truth table did depend solely on the truths of P and Q, and "P caused Q", whose truth table could not be constructed.

We say that "and" is a *truth-functional connective*. It's a *connective* because (quite literally) it connects the two sentences P and Q. It's *truth-functional* because its truth is entirely dependent upon, or *a function of*, the truth of the connected sentences. "P caused Q", as previously noted, is a *non-truth-functional connective*, since, while it is a connective, it's truth is not a function of the truth of the sentences it connects.

Below are introduced five truth functional connectives, along with their formal logical symbolizations, and a discussion of their relationships with natural English.

2.1 Conjunction

We start, of course, with *and*—or, in fancy logic parlance, *conjunction*. We represent conjunction with the "carrot" symbol: '$\wedge$'. So, rather than write "P and Q", our formal logical notation will have us write $P \wedge Q$. We typically read $P \wedge Q$ as "P and Q" (exactly as it originally was).

The truth table for $\wedge$ is as we constructed it in the previous lecture:

P	Q	P and Q
1	1	1
1	0	0
0	1	0
0	0	0

That this is the truth table for a symbol meaning "and" is perhaps obvious, but that obviousness is misleading. It bears repeating: $P \wedge Q$ is true if P and Q are true, and false otherwise. But that doesn't mean that $P \wedge Q$ has the truth table that it does *because* it is a formalization of the English connective "and". Rather, $P \wedge Q$ has the truth table that it does, *and* it happens to be the case that the use of "and" in natural English is quite similar.

One disanalogy between the two—$\wedge$ and "and", that is—is that "and" in English often implies a temporal order. If I say "I sat down and ordered a coffee.", it is implied that I sat down and *then* ordered a coffee. Similarly, "I ordered a coffee and sat down." implies that I ordered the coffee first and then sat down. But as much as this implication of temporal ordering appears in the natural use of "and", there is no formalization of temporal ordering in our logic and so a sentence like $C \wedge S$, where C is "I ordered a coffee." and S is "I sat down." has a truth that is indifferent to the order in which the events happened.

When we start to set aside these sorts of non-truth-related implications of a connective, many more words are modeled by $\wedge$, such as "but", "however", "moreover", "nonetheless", and "yet". Consider the sentence: "I ordered coffee but I was served orange juice." What the "but" tells you is that I ordered coffee *and* I was served orange juice. What the "but" *implies* is that there is some sort of conflict between my having ordered coffee and having been served orange juice (namely, that I was served the wrong thing). But that implication has no bearing on the *truth* of the sentence, and so "but" ends up with the same truth table as "and"/$\wedge$, because, absent the implication of conflict, a sentence of the form "P but Q" just informs you that both P and Q are true.

2.2 Disjunction

Like conjunction, *disjunction*—symbolized with '$\vee$'—has a clear companion in natural English—"or"—and a sensible truth table to match:

P	Q	$P \vee Q$
1	1	1
1	0	1
0	1	1
0	0	0

$P \vee Q$, read as "P or Q", is true as long as at least one of P or Q is true.

The main point of departure between our formal $\vee$ and the natural "or" of English is that "or" in English is vague as to whether it is *inclusive* or *exclusive*. An inclusive "or", which is what our $\vee$ models, is true when one of the connected sentences is true or both of them are true. An exclusive "or" is one that is true when one *and only one* of the connected sentences is true. Exclusivity is often implied with a natural English "or", but on further probing is ultimately not there. Consider, for example, the sentence "Solomon is at work or Solomon is at school." Implicitly, Solomon is at *one* of school or work, not both and not neither. But Solomon might be a teacher—making work and school the same place—and so, albeit somewhat pedantically, we'll want to say that the "or" in "Solomon is at work or Solomon is at school." is an inclusive one.

More natural (less obnoxious) cases of inclusive "or" in English tend to appear in more complex sentences. For example: "An absence from work will be excused if you or a dependent are ill." Here the "or" is definitely inclusive. The sentences it connects are "You are ill." and "A dependent of yours is ill.". Clearly the policy intends the following:

- If you are ill, then your absence will be excused.
- If a dependent of yours is ill, then your absence will be excused.
- If neither you nor a dependent are ill, then your absence is *not* excused (barring the relevance of some other policy).

But what if both you *and* your dependent are ill? If the "or" in the policy is exclusive, then your absence wouldn't be excused. But that's clearly absurd. Either your illness or a dependent's is *sufficient* for an excused absence, and both of you being ill doesn't do anything to undermine the sufficiency of either.

2.3 Negation

Unlike the other connectives that we are discussing, negation does not actually *connect* two sentences, but rather *operates* on one. Negation, which we will symbolize with '$\neg$', corresponds to "not" in natural English. Here's the truth table:

P	$\neg P$
1	0
0	1

We read $\neg P$ as "not P". So, for example, if P is "Paula is celebrating her 90th birthday this year.", then $\neg P$ could standardly read in any of the following ways:

- Paula is *not* celebrating her 90th birthday this year.
- *It's not the case that* Paula is celebrating her 90th birthday this year.
- *It's false that* Paula is celebrating her 90th birthday this year.

That the truth table is what it is comes as no surprise if you think about what the entries on the truth table mean. The "1" under P indicates that it's true that P, so "It is false that P.", i.e. $\neg P$, is not true. Similarly, the "0" indicates that P is false, so "It is false that P.", i.e. $\neg P$, *is* true.

2.4 Material Conditional

The *material conditional*, or *implication*, is by far the strangest of the connectives we are discussing, but also one of the most important and powerful (especially when we start thinking about arguments). Implication is symbolized with '$\rightarrow$', and is standardly associated with "If —, then —" sentences. It's truth table is

P	Q	$P \rightarrow Q$
1	1	1
1	0	0
0	1	1
0	0	1

The sentence that comes first in the implication—in this case, P—is called the *antecedent* (literally, "coming before"), and the sentence that comes second—Q—is called the *consequent* (since its truth is, in some sense, a consequence of the truth of the antecedent).

For an example, consider a sentence from earlier: "If you are ill, then your absence will be excused." Let P be "You are ill.", and let "Your absence is excused." be represented by Q. One way to think about the material conditional is that it's true if the rule it describes is not clearly broken in the given world/row of the truth table. So, in the first world, it's true that you are ill and it's true that your absence is excused. The company who has the policy "If you are ill, then your absence will be excused." is not in violation of the policy, and so we say that the conditional is true. In the second world, it's true that you are ill, but its false that your absence is excused. Either the policy is not in force at all, or it is being violated, and in either case we say that the conditional is false. So far, so good (and, hopefully, sensible).

Now (and this is where things get iffy) what do we do if it's false that you are ill? In the third world, it's false that you are sick, but your absence was excused nonetheless. It's not clear what the policy of the company might be at this point. But we do know that the company didn't *violate* the "If you are ill, then your absence will be excused." policy, and so we say that the material conditional is true. Similarly in the fourth world. The only way for the company to violate the policy is when you are ill and they don't excuse your absence. So, when you *aren't* ill, we can't say that the company *has violated* the policy, and so the material conditional comes out true.

This is weird for two reasons (at least). The first, which might already be apparent, is that the company might not have the policy of "If you are ill, then your absence will be excused." at all in worlds three and four, and yet we still mark the conditional as true (since the policy, in force or not, wasn't clearly violated). The second is that this policy/rule framework for understanding the truth table doesn't actually have to be present in the meaning of the connected sentences. Remember that the material conditional is supposed to be a *truth functional connective*, so its truth should depend only on the bare truth of the connected sentences. This means that some very strange sounding conditionals will come out true. For example, letting P be "It's Tuesday." and Q be "It's raining.", a rainy Tuesday (corresponding to the first line of the truth table) will make it true that "If it's Tuesday, then it's raining." when, intuitively, "If it's Tuesday, then it's raining." is false. Even worse, since a conditional is true if its antecedent is false, "If it's Tuesday, then it's raining." will also be true on any day other than Tuesday, whether it's raining or not. There are two (related) things that can be said to mitigate the strangeness of "If it's Tuesday, then it's raining." being true on a rainy Tuesday (or ever).

One is to note that conditional statements in natural English come with a lot of vague implications about the relationship between the connected sentences. "If it's Tuesday, then it's raining." sounds false because when are inclined to read it as making the additional claim that it being Tuesday *is the reason that* it's raining. "It being Tuesday is the reason it is raining.", or "It's raining because it's Tuesday." both sound like they are saying the same thing as "If it's Tuesday, then it's raining.". But they aren't, or, at least, they aren't in general. As already discussed with "because" in the previous lecture, and similarly for "is the reason that", the truth of sentences including these words depends on the *meaning* of the connected sentences, and not just on their truth. The material conditional, then, is strange in part because it's rare that we encounter a conditional in natural English that doesn't care about the meaning of the things it's connecting.

The other thing that we can do to mitigate the strangeness of the material conditional is to take it as an opportunity to acknowledge the limitations of our logic. The logical structure of Natural English is enormously complicated, a major topic of research in linguistics, philosophy, cognitive science, machine learning, and more, and it's even an open question whether there is *any* system of logic that can capture all of its complexity. As much as we might be able to accomplish with the system of logic we are learning now, we shouldn't expect it to get everything right.

2.5 Biconditional

The *biconditional*, which we symbolize with '$\leftrightarrow$', is a bit funny because describing its truth table is more straightforward than it's natural English counterparts. Here's that truth table:

P	Q	$P \leftrightarrow Q$
1	1	1
1	0	0
0	1	0
0	0	1

Quite simply, $P \leftrightarrow Q$ is true if P and Q have the same truth value, and false otherwise. $P \leftrightarrow Q$ is commonly read as "P if and only if Q.", or "P just in case Q.", and may also appear in text as "P iff Q.", where "iff" is a shortening of "if and only if".

The biconditional appears most often when giving definitions or truth conditions. For example:

- Someone is a bachelor *if and only if* they are an unmarried man.

- $P \leftrightarrow Q$ is true *if and only if* P and Q have the same truth value.
- Call 911 *just in case* it's an emergency.

For all its strangeness, the material conditional shines when looking at the biconditional. When reading $P \leftrightarrow Q$ as "P if and only if Q.", we can pull apart the conjunction and get "P if Q" and "P only if Q", which are $Q \to P$ and $P \to Q$, respectively. We can draw up a big truth table to see that $P \leftrightarrow Q$ is equivalent to the complex sentence "P if Q, and P only if Q", or, more formally, $(Q \to P) \wedge (P \to Q)$:

P	Q	$P \leftrightarrow Q$	$P \to Q$	$Q \to P$	$(Q \to P) \wedge (P \to Q)$
1	1	1	1	1	1
1	0	0	0	1	0
0	1	0	1	0	0
0	0	1	1	1	1

We'll say more about complex sentences in the next lecture.