Propositional Logic

LX 502 - Semantics
September 19, 2008

1. Semantics and Propositions

Natural language is used to communicate information about the world, typically between a speaker and an addressee. This information is conveyed using linguistic expressions. What makes natural language a suitable tool for communication is that linguistic expressions are not just grammatical forms: they also have content. Competent speakers of a language know the forms and the contents of that language, and how the two are systematically related. In other words, you can know the forms of a language but you are not a competent speaker unless you can systematically associate each form with its appropriate content.

Syntax, phonology, and morphology are concerned with the form of linguistic expressions. In particular, Syntax is the study of the grammatical arrangement of words—the form of phrases but not their content. From Chomsky we get the sentence in (1), which is grammatical but meaningless.

(1) Colorless green ideas sleep furiously

Semantics is the study of the content of linguistic expressions. It is defined as the study of the meaning expressed by elements in a language and combinations thereof. Elements that contribute to the content of a linguistic expression can be syntactic (e.g., verbs, nouns, adjectives), morphological (e.g., tense, number, person), and phonological (e.g., intonation, focus). We're going to restrict our attention to the syntactic elements for now. So for us, semantics will be the study of the meaning expressed by words and combinations thereof. Our goals are listed in (2).

(2)
i. To determine the conditions under which a sentence is true of the world (and those under which it is false)
ii. To develop a mechanism that enables us to derive the truth-conditions of a sentence from the meaning of the elements that make up the sentence

Pragmatics is the study of the meaning of a linguistic expression when it is used in a context. The linguistic expression retains its semantic content but can convey additional meanings, inferred from information in the context, rules governing conversation, and other extra-linguistic factors. The distinction between semantics and pragmatics is often viewed as literal versus non-literal meaning, respectively. Pragmatic meaning relies on semantic content, but it obtains from the actual use of a linguistic expression in a specific situation. The difference, then, between semantics and pragmatics can be viewed as a distinction between linguistic competence and performance. Pragmatics is often defined as the study of non-semantic meaning, which includes a wealth of meanings.

There is some terminology we need to adopt to simplify the following discussion. I reserve the terms *sentence*, *proposition*, and *utterance*. A *sentence* is a linguistic expression that is abstract and internalized. Sentences are not part of the outside world. They are part of a native speaker's knowledge of his language.

For us, a sentence is an arrangement of words in a syntactic structure. For example, the complementizer phrase [$_{CP}$ *I'm happy*] is a sentence of English and is a part of every native speaker's knowledge of English. A *proposition* is a logical expression, not a linguistic expression; it is a statement describing a (proposed) state of affairs in the world, that may be true of the world or false. A native speaker's competence includes matching the linguistic expression (a sentence) with the appropriate content (a proposition). An *utterance* is the concrete use of a linguistic expression in a context. The context includes information about the speaker of the utterance, the addressee(s), the location, the time, etc. Uttering a sentence in a context conveys its content, the proposition directly related to the sentence, i.e., its literal meaning. The utterance also conveys additional contextually-implicated meanings. These are also propositions but they are not related to the sentence and constitute its non-literal meaning.

(3)	i.	Sentence	A well-formed arrangement of words into a complementizer phrase
	ii.	Proposition	The description of a state of affairs in the world
	iii.	Utterance	A sentence used in a context

The object language is English. Sentences in English are part of the object language. In order to talk about these sentences and their meanings, I'm going to introduce a metalanguage called Propositional Logic.[1] This metalanguage is artificial but it is still a language, so it too has a syntax and a semantics. The idea is that we translate English expressions (sentences in the object language) into expressions of Propositional Logic (propositions in the metalanguage). The first step in the translation process is to determine what a well-formed proposition looks like. The second step is to specify a way to interpret what propositions say about the world that captures our intuitions about the content of the original English sentences.

Propositional Logic enables us to systematically determine the content of a compound sentence if we know the content of the simple sentences from which it is formed. For example, we can obtain the meaning of (4a) if we know the meaning of (4b) and (4c).

(4) a. John is happy and Mary is sad
  b. John is happy
  c. Mary is sad

Suppose (4b) correctly describes a state of affairs in the world. In other words, it is deemed true of the world. Suppose, however, that (4c) does not. It is does not correctly describe a state of affairs in the world and is therefore deemed false. What, then, can we say about (4a) and its description of the world. In light of what we know about (4b) and (4c), then (4a) must be false as well. Conversely, whenever (4b) and (4c) are true of the world, then it is necessarily the case that (4a) is true of the world. It is impossible for (4b) and (4c) to correctly describe a state of affairs in the world and for (4a) not to. This may not seem like much at present but it gives us a way of determining the meaning of a compound sentence from the meaning of its parts. We use Propositional Logic to express this relationship systematically and unambiguously.

This kind of contingency illustrates a fundamental kind of reasoning we're equipped with, called logical consequence or entailment, that is of great interest in semantics. I'll start by looking at that.

---

[1] Propositional logic goes by other names in the literature, such as sentence logic, statement logic, propositional calculus, or simply PL.

2. Arguments

Logic is the study of *arguments*. We all share a basic understanding of this term. An argument is a series of reasons supporting a conclusion. The following argument supports the conclusion that John is wearing a coat.

(5) If it is snowing, it is cold
If it is cold, John is wearing a coat
It is snowing
———————————————
Therefore, John is wearing a coat

I have separated the reasons in this argument (above the line) from the conclusion (below the line). There is a strong sense that the argument in (5) constitutes a valid piece of reasoning. If the reasons hold, then the conclusion is guaranteed. There are arguments that are less transparent than the one in (5). The argument in (6), for example, is a little more difficult to see but it too is valid.

(6) If it is raining, it is not cold
If is not raining, John is not wearing a coat
It is cold
———————————————
Therefore, John is not wearing a coat

There are also arguments that fall short. Arguments for which the conclusion is not guaranteed. Such arguments do not constitute valid pieces of reasoning. The argument illustrated in (7) is not valid.

(7) If it is snowing, it is cold or it is wet
If it is cold, John is wearing a coat
It is snowing
———————————————
Therefore, John is wearing a coat

The conclusion in (7) seems premature. The fact that it is snowing could imply that it is wet but not cold. When it is cold, John is wearing a coat. But we don't know what he wears when it is not cold. He could be naked. Such a conclusion would be consistent with the reasons given in (7). To show that an argument is not valid, it is customary to give a counterexample in which the same reasons hold, but a different conclusion follows. So imagine a situation in which it is snowing and wet but not cold and John is wearing nothing.

The notion of an argument is made rigorous in Logic. An **argument** is defined as a set of propositions consisting of (i) a set of premises and (ii) a conclusion. A **valid argument** is one in which, if the premises are all true, then the conclusion must be true. I.e., if the premises hold, the conclusion is guaranteed. Otherwise, the argument is not valid. Let's revisit the example in (8).

(8) If it is snowing, it is cold — Premise 1
If it is cold, John is wearing a coat — Premise 2
It is snowing — Premise 3
———————————————
Therefore, John is wearing a coat — Conclusion

The argument in (8) is a valid because whenever the propositions expressed by the premises are true, the proposition expressed by the conclusion is necessarily true. Be careful here. The fact that this is a valid argument is still only an intuition right now. We don't have a systematic way to show validity. Propositional Logic will provide us with the tools to do this.

The examples in (9) - (11) are valid arguments.

(9)	If Mary is dancing, John is dancing	Premise
	Mary is dancing	Premise
	Therefore, John is dancing	Conclusion

(10)	Jacques ate and drank	Premise
	Therefore, Jacques ate	Conclusion

Pay close attention to the argument in (11). The premises are not facts about the actual world. This illustrates an essential property of valid reasoning. A valid argument is a relation between propositions, not a reflection of actual facts: whenever the premises are true, the conclusion is necessarily true. So even though the first premise in (11) is not verifiable in the actual world and the second is false in the actual world, it does not alter the relationship between the premises and the conclusion.

(11)	If the moon is made of cheese, then Mickey is happy	Premise
	The moon is made of cheese	Premise
	Therefore, Mickey is happy	Conclusion

The arguments in (12) and (13) are not valid. For now, we have to rely on our intuitions to see this. Propositional Logic will enable us to show this.

(12)	If Los Angeles is in Canada, then Los Angeles is in North America	Premise
	Los Angeles is in North America	Premise
	Therefore, Los Angeles is in Canada	Conclusion

(13)	There are clouds in the sky or it is not raining	Premise
	There are clouds in the sky	Premise
	Therefore, it is raining	Conclusion

Finally, the notion of validity or a valid argument goes by many names. Sometimes it is called *logical consequence* or *entailment*. The term *logical inference* (or *deductive inference*), on the other hand, designates the mode of reasoning used to get from the premises to the conclusion. The type of inferences that we have been looking at in this section are logical/deductive ones. They rely on the logical properties of words like *and*, *or*, *if…then*, *not*, *if and only if*, etc. There are other kinds of inference. The textbook chapter I provide, for instance, briefly describes what is known as inductive reasoning, i.e., reasoning based on probable facts or experience.

3. Propositional Logic (PL)

3.1. Introduction

Propositional logic is the logical language of propositions. We are going to use PL as our metalanguage to describe English (the object language)—in particular, the meaning of English sentences. We are going to use PL because it is unambiguous and fully determined. As a language, PL has both a syntax and a semantics. Its syntax determines the proper form of propositions, just like the syntax of English determines the grammatical form of sentences in English. The semantics of PL determines the content of these propositions.

The idea is this. We want to translate English sentences into unambiguous propositions of PL. Then we want to interpret the propositions. This sets up a correspondence. The meaning of a sentence in English corresponds to the interpretation of the proposition it expresses.

But first we have to contend with the translation. We need a translation schema that associates simple declarative sentences (which cannot be decomposed into smaller sentences) to atomic propositions (the basic blocks of PL). We also need our translation schema to systematically and unambiguously associate compound sentences in English with compound propositions. To do this translation effectively, however, I need to define the syntax of PL: the component that determines the proper form of propositions in PL.

3.2. Syntax

The goal of this section is to lay out the syntax of Propositional Logic. The syntax consists of a set of basic items (i.e., the lexicon) and a set of construction rules (i.e., syntactic rules) which determine the well-formed expressions in our logical language. In a natural language, like English, a grammatical sentence is a string of words in the right order and with the right structure. Our logical language works in exactly the same way.

(14) **Syntax of Propositional Logic**

A well-formed formula (wff) is a string of lexical items in (A) that satisfies the syntactic rules of the language in (B). A proposition is defined as a wff of PL.

(A) The Lexicon

Basic Terms: atomic propositions represented by the symbols: p, q, r, s, etc…
Connectives: $\neg$ negation, $\wedge$ conjunction, $\vee$ disjunction, $\rightarrow$ implication, $\leftrightarrow$ equivalence
Parentheses: (,)

(B) The Syntactic Rules

i. Any atomic proposition is itself a well-formed formula (wff)
ii. If φ is a wff, then $\neg\varphi$ is a wff
iii. If φ and ψ are wff, then $(\varphi \wedge \psi)$, $(\varphi \vee \psi)$, $(\varphi \rightarrow \psi)$, $(\varphi \leftrightarrow \psi)$ are wffs
iv. Nothing else is a wff

The basic terms of PL are called *atomic propositions* (also *atomic sentences*). They correspond to simple declarative sentences in the object language, e.g., *it is raining*, *Mary is intelligent*, *Bob slept*, etc. These atomic propositions are the building blocks of our logical language. You can think of them as LEGO blocks that come in only one size, constituting a unit of information, i.e., a proposed fact about the world. The set of atomic propositions is taken to be infinite—this is really for convenience and has no consequences for us. Atomic propositions are represented in PL with lower case Roman letters, e.g., *p, q, r, s, etc*. The atomic propositions of PL are the basic terms of the lexicon. Along with these basic terms, the lexicon also includes a set of logical connectives: ¬ (negation), ∧ (conjunction), ∨ (disjunction), → (implication), ↔ (equivalence). The purpose of these connectives is to combine propositions together to form compound propositions. To keep the combinations orderly, we also include parentheses in the lexicon.

The syntactic rules do several things. Rule (i) determines that the atomic propositions are all wffs. Rules (ii) and (iii) determine how compound wffs are formed from simpler ones using the connectives and parentheses. These rules are also **recursive**. Any wff determined by these rules can be used in these rules. For example, rule (i) says the atomic proposition p is a wff. From rule (ii), we get that ¬p is a wff. We can reapply rule (ii) to ¬p and we get that ¬¬p is a wff. Similarly, ¬¬¬p is a wff. Etc. Recursion allows the syntax to determine the set of propositions with two simple rules. Rule (iv) makes Propositional Logic a **closed** system. Only those formulas satisfying rules (i)-(iii) are considered wffs and nothing else; they completely determine the propositions of PL.

(15) Examples of **well-formed formulas**

a. p
b. ¬¬p
c. (p ∧ p)
d. ¬(p ∨ q)
e. (¬(p ∨ q) ∧ p)
f. ¬((p ∨ q) ∧ p)
g. ¬(p ↔ (r ∨ s))
h. (p ↔ ¬(r ∨ s))
i. ((p ∧ q) ∨ (s ∧ r))
j. ((((p → q) → ¬r) ↔ s) ∨ (t ∧ u))

(16) Examples of formulas that are **not well-formed**

a. pqr
b. (p
c. p¬
d. ∨q
e. (¬p ↔ r ∨ s)
f. → ∨ ∧
g. pq →
h. p ∧ p
i. (p) ∧ p
j. →∧pq ∨pq

3.3. Translating between English and Propositional Logic

Translating from English to Propositional Logic takes practice. To do it efficiently is an art. Thankfully, you can do it systematically. Once you learn the semantics of PL, translation will become easier.

(17) Guidelines to Translation

i. If the sentence you are translating is a simple declarative sentence, then you're done. Remember, in PL, these are just represented as atomic propositions, e.g., p.

ii. Assuming the sentence you have to translate is a compound sentence, the first step is to identify the simple, declarative sentences in it. (We will ignore embedded sentences. So for example, the sentence I know Bob is happy is treated as a simple declarative sentence, even though it contains the embedded sentence Bob is happy.)

iii. Once you've identified the simple, declarative sentences, translate these into atomic propositions. You can write each one like this: "p = …" or "Let p = …"

iv. Next you need to identify those words that string the simple sentences together in English, e.g., and, but, or, if…then, if and only if, just in case, it is not the case that, unless, only if, when, etc. (The handout from September 16 should help with this.)

v. Once you've identify these words, translate them into the appropriate connectives in PL. (Again, the handout from September 16 should help with this.)

vi. Lastly, you need to determine the order the atomic propositions combine in with the connectives. When you combine the atomic propositions, remember that the syntactic rules of PL must be satisfied—so you're not working in a vacuum. These rules can help guide you.

vii. Remember to use (but not abuse) the parentheses. They make a difference. Pay attention to Rule (ii) to figure out where they go.

Atomic propositions correspond to simple, declarative sentences. So these are the easiest to translate. To do this, we need only equate an atomic proposition to the sentence we want it to correspond to.

(18) Translating Simple Declarative Sentences

a. Let p = It is raining

b. Let q = Mary is sick

c. Let t = Bob stayed up late last night

d. Let r = Paris is the capital of France

e. Let s = John is a loud-mouth

Suppose that what we understand informally as negation ($\neg$) corresponds to the use of "not" and related terms in natural language. Keeping in mind the translations in (19), we can translate the following compound sentences into PL

(19) Translating Negation

a. It isn't raining
   $\neg p$

b. It is not the case that Mary isn't sick
   $\neg\neg q$

c. Paris is not the capital of France
   $\neg r$

d. John is in no way a loud-mouth
   $\neg s$

e. Bob did not stay up late last night
   $\neg t$

Now suppose what we understand informally as conjunction ($\wedge$) corresponds to the use of "and" and related terms in natural language, like "but". Using the translations in (18) and (19), we can translate the following compound sentences into PL.

(20) Translating Conjunction

a. It is raining and Mary is sick
   $(p \wedge q)$

b. Bob stayed up late last night and John is a loud-mouth
   $(t \wedge s)$

c. Paris isn't the capital of France and It isn't raining
   $(\neg r \wedge \neg p)$

d. John is a loud-mouth but Mary isn't sick
   $(s \wedge \neg q)$

e. It is not the case that it is raining and Mary is sick

   translation 1: It is not the case that both it is raining and Mary is sick
   $\neg(p \wedge q)$

   translation 2: Mary is sick and it is not the case that it is raining
   $(\neg p \wedge q)$

The compound sentence in (20e) is ambiguous. We capture the ambiguity by giving both translations of the sentence, i.e., two distinct propositions. In any given context or with the proper intonation, the sentence has only one meaning and expresses one of these propositions. But the linguistic expression is ambiguous between the two.

Suppose what we understand informally as disjunction ($\vee$) corresponds to the use of "or" and related terms in natural language. Using the translations in (18)-(20), we can translate the compound sentences below into PL.

(21) Translating Disjunction

a. It is raining or Mary is sick
   $(p \vee q)$

b. Paris is the capital of France and it is raining or John is a loud-mouth
   $((r \wedge p) \vee s)$
   $(r \wedge (p \vee s))$

c. Mary is sick or Mary isn't sick
   $(q \vee \neg q)$

d. John is a loud-mouth or Mary is sick or it is raining
   $((s \vee q) \vee p)$
   $(s \vee (q \vee p))$

e. It is not the case that Mary is sick or Bob stayed up late last night
   $\neg(q \vee t)$
   $(\neg q \vee t)$

Whenever a compound sentence includes conjunction and disjunction, ambiguity is quite possible so be on guard. Once again, the sentences are translated into distinct propositions, one for each meaning. The context or intonation can disambiguate which proposition actually corresponds to the sentence whenever it is uttered.

Suppose what we understand informally as implication ($\rightarrow$) corresponds to the use of "if … then …" in natural language and related terms like "when". Using the translations in (18)-(21), we can translate the following compound sentences into PL. Implication is not a straightforward connective the first time around, so don't panic if you don't get right away.

(22) Translating Implication

a. If it is raining, then Mary is sick
   $(p \rightarrow q)$

b. It is raining, when John is a loud-mouth
   $(s \rightarrow p)$

c. Mary is sick and it is raining implies that Bob stayed up late last night
   $((q \wedge p) \rightarrow t)$

d. It is not the case that if it is raining then John isn't a loud-mouth
   $\neg(p \rightarrow \neg s)$

Suppose what we know informally as equivalence ($\leftrightarrow$) corresponds to the use of "if and only if" in natural language and related terms, such as "just in case" and "if … , then … , and vice versa". Using the translations in (18)-(22), we can translate the following compound sentences into our propositional logic. The equivalence is also not a straightforward connective the first time around, so once again my advice is don't panic if you don't get it right away.

(23) Translating Equivalence

a. It is raining if and only if Mary is sick
   $(p \leftrightarrow q)$

b. If Mary is sick then it is raining, and vice versa
   $((p \rightarrow q) \wedge (q \rightarrow p))$
   $(p \leftrightarrow q)$

c. It is raining is equivalent to John is a loud-mouth
   $(p \leftrightarrow s)$

d. It is raining is not equivalent to John is a loud-mouth
   $\neg(p \leftrightarrow s)$

3.4 Truth and Truth-Values

As our theory of meaning, we will adopt the *Correspondence Theory of Truth* (Tarski 1944)[2]. This theory sets up a correspondence between (i) the logical meaning of a sentence and (ii) the outside world. We characterize the outside world in terms of states of affairs. Intuitively, a state of affairs is a relationship between objects in the outside world, and is synonymous for our purposes terms like situation, event, act, state. Propositions are said to describe states of affairs. For example, the sentence *there is a book on John's dinner table* expresses a proposition and that proposition either correctly describes a state of affairs in the outside world or it does not. If this proposition correctly describes a state of affairs in the outside world, I say it is true. Otherwise, I say it is false. I can verify this by going to John's dinner table and checking that there is a book on it. To represent that a proposition is true, we assign it the truth value 1 (true); otherwise, we assign it the truth-value 0 (false). There are no other truth-values. So a proposition is either true (with a value of 1) or false (with a value of 0).

(24) Let p = John is happy

The atomic proposition p either corresponds to the outside world or it does not.

p has value 1: *p correctly describes a state of affairs in which there is an person named John and he has the property of being happy.*

p has value 0: *p does not correctly describe a state of affairs in which there is an person named John and he has the property of being happy.*

We're not actually interested in verifying the truthfulness of a proposition. We're interested in the conditions under which the proposition (and the sentence that expressed it) is true. In other words, we want to know what the outside world has to be like for the proposition to be true; and what it has to be like for the proposition to be false. It is these conditions which constitute the logical meaning of a sentence for us. (I'll try to persuade you of this further down.)

---

[2] Non-declarative statements need a different semantic treatment, though not altogether different. These include but are not limited to questions, imperatives, and exclamatives.

(25) Let p = John is happy

p is true iff it correctly describes a state of affairs in which there is an person named John and he has the property of being happy; otherwise p is false.

The semantics of PL provides a formal way to express truth-values and truth-conditions. The idea is this. The truth-conditions for atomic propositions are straightforward. They either correctly describe the outside world or not. That is, an atomic proposition is true if and only if it correctly describes a state of affairs in the world. However, we need a systematic way to determine the meaning of a compound proposition (its truth-conditions) from the meaning of the atomic propositions that make it up (their truth-conditions). Before we get to the semantics of PL, there are some additional terms that need to be discussed.

There are some sentences that express propositions that are always true, always false, or dependent on the context. Propositions that are always true, no matter what the conditions of the outside world are like, are called a tautologies (26). Propositions that are always false, no matter what the conditions of the outside world are like, are called a contradictions (27). Propositions that may be true or false, depending on the conditions of the outside world, are called contingencies (28).

(26) a. Every man is a man
b. A bachelor is unmarried
c. It is raining or it is not raining

(27) a. Every man is a not man
b. A bachelor is married
c. It is raining and it is not raining

(28) a. Bob is sleeping
b. Doug is Canadian
c. Bob ate all the bacon and Doug drank all the beer

With regards to truth and meaning, let me convince you that truth-conditions are really what we're after and not truth-values. For example, consider my friend Bob. Sometimes Bob is sleeping and other times he is not. We could go over to Bob's house right now and see if he is actually sleeping. Regardless of what we find at any given time, the meaning of the sentence in (28a) is going to be the same. It is independent of the actual truth of the sentence. Even if we cannot verify the actual truth of this sentence, it is still meaningful. Furthermore, whatever truth value we find now may change later, but there is a strong sense that the meaning of the sentence in (28a) does not change. I conclude that there is a difference between truth and meaning.

Semanticists are not really interested in actual truth—in whether Bob is actually sleeping. Instead they are interested in truth-conditions. These are conditions that the outside world must satisfy (what it has to be like) for the proposition expressed by a sentence to be true. Here's a sophisticated formulation of truth-conditions. They are "the conditions that need to be met for there to be an appropriate correspondence between the proposition expressed by a sentence and the state of affairs in the world." (de Swart 1998). From this, I conclude that meaning is equivalent to truth-conditions under the correspondence theory of truth.

3.5. Compositionality

It's time to take stock of our theory. Each proposition can take a truth value of either 1 (true) or 0 (false). For atomic propositions, this truth value corresponds to the state of affairs in the outside world. In other words, an atomic proposition p is true if and only if it correctly describes a state of affairs in the outside world. This does not tell us how the value of compound propositions is obtained. We'll have to figure out a systematic way to obtain these. There is something that will help: the Principal of Compositionality

(29) Principle of Compositionality

The meaning of the whole is a meaning of the parts and the way they put together.
I.e., the truth value of a compound proposition is determined from the values of its syntactic parts and the syntactic structure of the wff (i.e., its connectives and their arrangements).

But how does this all work?

(30) Example

The complex proposition $(\neg p \land (q \rightarrow s))$ has a syntactic structure, fully specified by the syntactic rules of propositional logic. These are the same rules that tell us $(\neg p \land (q \rightarrow s))$ is a wff in PL.

a. The atomic propositions p, q, s are wffs (by rule i)
b. Since p is a wff, then $\neg p$ is a wff (by rule ii)
c. Since q and s are wffs, then $(q \rightarrow s)$ is a wff (by rule iii)
d. Since $\neg p$ is a wff and $(q \rightarrow s)$ is a wff,
then $(\neg p \land (q \rightarrow s))$ (from b. and by rule iii)

Working backwards, the syntactic rules in (30d) tell us that the proposition $(\neg p \land (q \rightarrow s))$ has three components: the proposition $\neg p$, the connective $\land$, and the proposition $(q \rightarrow s)$. We can represent this as the structure in (31).

(31) $(\neg p \land (q \rightarrow s))$

$\neg p$ $\quad\land\quad$ $(q \rightarrow s)$

The rules in (30b) and (30c) tell us that the proposition $\neg p$ has is made up of the connective $\neg$ and the atomic proposition p; and the proposition $(q \rightarrow s)$ is made up of the atomic proposition q, connective $\rightarrow$, and the atomic proposition p. We can represent this as the structures in (32).

(32) a. $\neg p$

$\neg$ $\quad$ p

b. $(q \rightarrow s)$

q $\quad\rightarrow\quad$ s

Putting it all together we get the structure in (33).

(33)

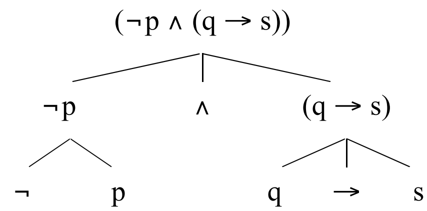

The Principal of Compositionality tells us that the meaning of $(\neg p \wedge (q \rightarrow s))$ is determined from the meaning of its parts $\{p, q, s, \neg, \wedge, \rightarrow\}$ and the way that they are put together. So far we know the following two things: (i) we know the way the parts go together in the syntax and (ii) we know that the truth values of the atomic propositions depend on the outside world. However, there are three things that we do not know.

(34)
- i. The semantics of the connective $\neg$: we don't know how it effects the meaning of p
- ii. The semantics of the connective $\rightarrow$: we don't know how it effects the meaning of p and r
- iii. The semantics of the connective $\wedge$: we don't know how it effects the meaning of $\neg p$ and $(q \rightarrow s)$

Without knowing the semantics of the connectives we have no way to determine the truth-conditions of a compound proposition. The next sections, then, is dedicated to the semantics of the connections. This boils down to figuring out how the connectives affect truth-values.

3.6. Connectives and Truth Tables

The connectives combine propositions into compound propositions. These propositions all have truth-values in the semantics. So it follows that in the semantics, connectives combine the truth-values of each proposition and return a single truth-value for the compound proposition. For example, if φ is a propositions in PL, then $\neg\varphi$ is a proposition in PL. Moreover, the propositions φ and $\neg\varphi$ both have truth-values in PL. The truth-value of compound proposition $\neg\varphi$ is systematically determined from the truth-value of φ and the connective $\neg$. The connective $\neg$ associates each truth-value to another. This association is not ad-hoc. It corresponds to our intuitions about negation in natural language. We write the associations in a table, called a **Truth Table**. The truth-table in (35) represents the semantics of the connective $\neg$.

(35) Truth Table for Negation ($\neg$)

φ	$\neg\varphi$
0	1
1	0

The truth table in (35) provides the truth value of $\neg\varphi$ for each value of φ. The truth values of φ are presented in the first column. And for each truth value of φ, the corresponding value of $\neg\varphi$ is presented in the same row, under the column $\neg\varphi$. In other words, if φ has a truth value of 0, then $\neg\varphi$ has a truth value of 1; and if φ has a truth value of 1, then $\neg\varphi$ has a truth value of 0.

These values were not chosen arbitrarily. The relationship between φ and $\neg\varphi$ is modeled after the logic behind negation. If the sentence in (36a) is true, we expect that the one in (36b) is false. And vice versa.

(36) a. It is raining p
b. It is not raining $\neg p$

The same kind of table (though a little more complex) represents the meaning of each connective.

(37) Truth Table for Conjunction ($\wedge$)

φ	ψ	$\varphi \wedge \psi$
0	0	0
0	1	0
1	0	0
1	1	1

Again, these values are not arbitrary. The relationship between φ, ψ and $(\varphi \wedge \psi)$ is modeled after the logic behind sentential conjunction. This is illustrated in (38). The sentence in (38a) is true just in case both the sentence in (38b) is true and the sentence in (38c) is true.

(38) a. John is happy and Mary is sad $(p \wedge q)$
b. John is happy p
b. Mary is sad q

(39) Truth Table for Disjunction ($\vee$)

φ	ψ	$\varphi \vee \psi$
0	0	0
0	1	1
1	0	1
1	1	1

The relationship between φ, ψ and $(\varphi \vee \psi)$ is modeled after the logic sentential or. This is illustrated in (40). The sentence in (40a) is true just in case the sentence in (38b) is true or the sentence in (38c) is true.

(40) a. John is happy or Mary is sad $(p \vee q)$
b. John is happy p
b. Mary is sad q

(41) Truth Table for Implication ($\rightarrow$)

φ	ψ	$\varphi \rightarrow \psi$
0	0	1
0	1	1
1	0	0
1	1	1

The relationship between φ, ψ and ($\varphi \rightarrow \psi$) is modeled after "if…then" conditional statements. This is illustrated in (42). The sentence in (42a) is true unless both the sentences in (42b) is true and the sentence in (42c) is false. So, a conditional statement is false when the condition (φ) is satisfied but the consequent (ψ) is not. E.g., if it rains and it is not cold.

(42) a. If it is raining, it is cold ($p \rightarrow q$)
b. It is raining p
b. It is cold q

(43) Truth Table for Equivalence ($\leftrightarrow$)

φ	ψ	$\varphi \leftrightarrow \psi$
0	0	1
0	1	0
1	0	0
1	1	1

The relationship between φ, ψ and ($\varphi \leftrightarrow \psi$) is modeled after "if and only if" equivalence statements. This is illustrated in (44). The sentence in (44a) is true if the sentences in (44b) and (44c) have the same value. In this example, Bob's happiness coincides exactly with Mary's dancing. An equivalence is false when one sentence (φ) differs in its truth-value from the other (ψ). E.g., Bob is happy and Mary is not dancing.

(44) a. Bob is happy if and only if Mary is dancing ($p \leftrightarrow q$)
b. Bob is happy p
b. Mary is dancing q

We now have enough to determine the meaning of ($\neg p \wedge (q \rightarrow s)$)

(45) Truth Table

p	q	s	$\neg p$	$(q \rightarrow s)$	$(\neg p \wedge (q \rightarrow s))$
0	0	0	1	1	1
0	0	1	1	1	1
0	1	0	1	0	0
0	1	1	1	1	1
1	0	0	0	1	0
1	0	1	0	1	0
1	1	0	0	0	0
1	1	1	0	1	0

Let me explain this complicated table. The truth values under the atomic propositions p, q, s represent all the possible combinations of truth-values between them. There are 8 combinations. (These correspond to 8 potential states of affairs in the outside world.)

The truth values for $\neg p$ are obtained using the truth table for negation in (35). In each row, when the value for p is 0, the corresponding value for $\neg p$ is 1. When the value for p is 1, the value for $\neg p$ is 0.Similarly, the truth-values for $(q \rightarrow s)$ are obtained using the truth table for implication in (41). When the truth-value of q is 1 and the truth-value of s is 0, the corresponding value of $(q \rightarrow s)$ is 0. Otherwise, the corresponding truth-value is 1. The truth-values for $(\neg p \wedge (q \rightarrow s))$ are obtained using the truth table for conjunction in (37), where $\varphi = \neg p$ and $\psi = (q \rightarrow s)$. In each row, when the truth value of $\neg p$ is 1 and the value of $(q \rightarrow s)$ is 1, then the corresponding value of $(\neg p \wedge (q \rightarrow s))$ is 1.

Why, then, is this a representation of the meaning of $(\neg p \wedge (q \rightarrow s))$? This truth table gives us the conditions under which it is true and those under which it is false. In particular, $(\neg p \wedge (q \rightarrow s))$ is true when the following conditions hold.

(46) the proposition $(\neg p \wedge (q \rightarrow s))$ is true if and only if

p is false, q is false, and s is false; or
p is false, q is false, and s is true; or
p is false, q is true, and s is true
otherwise it is false

Suppose we relate the proposition $(\neg p \wedge (q \rightarrow s))$ back to a sentence in English, like the sentence in (47). The truth table in (45) tells us that the sentence in (47) is true if and only if it correctly describes a state of affairs in which (i) John isn't going out, Mary isn't going out, and Sue isn't going; (ii) John isn't going out, Mary isn't going out, and Sue is going; or (iii) John isn't going out, Mary is going out, and Sue is going.

(47) John isn't going out and if Mary is going out, then Sue is

Before I continue, I need to make two remarks. First, equivalence is a very important connective because it allows us to test whether two propositions coincide in their truth-values. But this is just the same as verifying that two propositions have the same truth-conditions. Since truth-conditions represent meaning for us, this then is a means to test whether two propositions are equivalent in meaning. If they are, we say that the propositions are truth-conditionally equivalent or truth-conditional identical. In other to show equivalence, the simplest way in PL is to set up a truth table that includes columns for the propositions φ, ψ and $(\varphi \leftrightarrow \psi)$. The two propositions φ and ψ are equivalent if the value of $(\varphi \leftrightarrow \psi)$ is always 1.

Finally, a word about truth-tables and connectives. Truth tables are the standard way of representing the semantics of connectives in PL because they are easy to read and apply. Unfortunately, they obscure the nature of connectives as binary relations on the set (or functions, more specifically). You can think of connective $\neg$ as the function that maps 0 to 1 and 1 to 0. You can think of the connective $\wedge$ as the function that maps the ordered pair $\langle 0,0\rangle$ to 0, $\langle 0,1\rangle$ to 0, $\langle 1,0\rangle$ to 0, and finally $\langle 1,1\rangle$ to 1. (Each row in the truth table is the input of this function and the output is a truth-value.)

4.2. Semantics

The goal of this section is to lay out a semantics of Propositional Logic, one that allows us to derives the truth-conditions of compound propositions from the truth-conditions of atomic propositions. The semantics of PL consists of three components. We needs a formal representation of the outside world, called a model. We need a function V that assigns to each proposition a truth-value, but it can't be just any function. It needs to capture the contribution of the connectives and reduce the value of a compound proposition to the value of its parts. Finally we need a set of rules that govern this behaviour of V.

(48) **Semantics of Propositional Logic**

(A) Models

Because we do not know every aspect of the world, we avoid evaluating truth with respect to the whole world; instead we evaluate the truth of a sentence with respect to a model of the world (M): a small, fully-specified, logical representation of a piece of the world.

(B) Valuation Function

The valuation function V_M assigns to each proposition a truth value and it obeys the semantic rules in (C). The valuation function assigns to an atomic proposition the value 1 (true) if and only if the atomic proposition corresponds to the model M.

(C) Semantic Rules of Composition [These rules embody the truth-conditions above]

i. If p is an atomic proposition, then $V_M(p) = 1$ if and only if p correctly describes M
ii. $V_M(\neg\varphi) = 1$ iff $V_M(\varphi) = 0$
iii. $V_M(\varphi \wedge \psi) = 1$ iff $V_M(\varphi) = 1$ and $V_M(\psi) = 1$
iv. $V_M(\varphi \vee \psi) = 1$ iff $V_M(\varphi) = 1$ or $V_M(\psi) = 1$
v. $V_M(\varphi \rightarrow \psi) = 1$ iff $V_M(\varphi) = 0$ or $V_M(\psi) = 1$
vi. $V_M(\varphi \leftrightarrow \psi) = 1$ iff $V_M(\varphi) = V_M(\psi)$

There is an additional component that is not very interesting in PL but will be very important in Predicate Logic. It is called the interpretation function (also the denotation function). In PL, the interpretation function is just equivalent to the valuation function V_M. Nothing more. But I will use it to get you used to it.

(49) Interpretation Function

For any expression α, we write the interpretation of α with respect to the model M as $[\![\alpha]\!]^M$, where $[\![\alpha]\!]^M = V_M(\alpha)$

We now have enough to determine the interpretation (or denotation) of $(\neg p \wedge (q \rightarrow s))$

(50) Example: The semantics of the proposition $(\neg p \wedge (q \rightarrow s))$

$[\![(\neg p \wedge (q \rightarrow s))]\!]^M = 1$	**iff**	$V_M((\neg p \wedge (q \rightarrow s))) = 1$ (To simplify this condition apply rule iii)
	iff	$V_M(\neg p) = 1$ and $V_M((q \rightarrow s)) = 1$ (To simplify this condition apply rule ii)
	iff	$V_M(p) = 0$ and $V_M((q \rightarrow s)) = 1$ (To simplify this condition apply rule v)
	iff	$V_M(p) = 0$ and [$V_M(q) = 0$ or $V_M(s) = 1$] (This is the final condition)

Evaluating the truth of the proposition $(\neg p \wedge (q \rightarrow s))$ is done by specifying a model. In (51), I give three models. Each model is consistent with a different set-up in the world, i.e., a different set of states of affairs. The model in (51a) determines a world in which the proposition $(\neg p \wedge (q \rightarrow s))$ is true. The model in (51c) determines a world in which the proposition $(\neg p \wedge (q \rightarrow s))$ is false.

(51) a. Let the model M_1 be consistent with the truth value assignments: $V(p) = 0$, $V(q) = 0$, $V(s) = 0$

Based on the semantics in (50) and these values, then $[\![(\neg p \wedge (q \rightarrow s))]\!]^{M_1} = 1$,
i.e., the proposition $(\neg p \wedge (q \rightarrow s))$ *is true in the model* M_1

b. Let the model M_2 be consistent with the truth value assignments: $V(p) = 0$, $V(q) = 0$, $V(s) = 1$

Based on the semantics in (50) and these values, then $[\![(\neg p \wedge (q \rightarrow s))]\!]^{M_2} = 1$,
i.e., the proposition $(\neg p \wedge (q \rightarrow s))$ *is true in the model* M_2

c. Let the model M_3 be consistent with the truth value assignments: $V(p) = 1$, $V(q) = 1$, $V(s) = 1$

Based on the semantics in (50) and these values, then $[\![(\neg p \wedge (q \rightarrow s))]\!]^{M_3} = 0$,
i.e., the proposition $(\neg p \wedge (q \rightarrow s))$ *is false in the model* M_3