Semantics: Propositional Logic

LING 106, April 1, 2009

1. CLASSICAL LOGIC

Consider the following two arguments:

(1) Everything in this class so far has been hard.
∴ Semantics will probably also be hard.

(2) All humans are mortal.
Socrates is a human.
∴ Socrates is mortal.

(Things above the line are *premises*; the thing below the line, marked with the "therefore" symbol, is the *conclusion*.)

Neither one of these arguments is inherently better than the other, but they demonstrate two different kinds of thinking: *deduction* and *induction*, respectively.

1.1. Induction

The first argument illustrates INDUCTION: generalization from premises. Unfortunately, some induction is better than other induction...

(3) We've seen thousands of swans; they've all been white.
∴ All swans are white.

(4) The nine Penn students in this class can draw a FSA.
x∴ In general, Penn students can draw a FSA.

(5) Of the nine students here, one is an anthropology major.
x∴ In general, 1/9 of all Penn students are anthropology majors.

(Using "x" to mark an invalid conclusion.) Generally speaking, whether induction is valid depends on the particular content.

[Historical footnote: (3) was the long-standing example of inductive reasoning, until Europeans reached Australia and found *Cygnus atratus*, a black species of swan. Which just goes to show you.]

1.2. Deduction

DEDUCTION draws conclusions without generalizing beyond the premises. One form is given in (2). Other deductive arguments with the same form:

(2) All humans are mortal.
Socrates is a human.
∴ Socrates is mortal.

(6) All squorfles are brachtish.
The Archon of Ptharn is a squorfle.
∴ The Archon of Ptharn is brachtish.

(7) All Penn students are from Ireland.
George W. Bush is a Penn student.
∴ George W. Bush is from Ireland.

We don't have to know what the words mean. And the premises can both be false (in which case the conclusion may be false as well, though it's possible to get a true conclusion from false premises). But the argument will still be valid by virtue of its form.

Though you do have to be careful about blind substitution:

(8) I'm looking at a picture of Mark Twain.
Mark Twain is the author of *Tom Sawyer*.
∴ I'm looking at a picture of the author of *Tom Sawyer*.

(9) I'm looking at a picture of someone.
Someone is the author of *Tom Sawyer*.
x ∴ I'm looking at a picture of the author of *Tom Sawyer*.

(10) Jimmy Olsen thinks Superman can fly.
Clark Kent is Superman.
x ∴ Jimmy Olsen thinks Clark Kent can fly.

We'll worry about these later, if at all.

1.3. Why look at logic

Classical logic is a way of classifying arguments as valid or invalid based solely on their forms, regardless of the details of the content.

Compare this to what we did in syntax:

- Basic elements aren't *particular words*, but syntactic *categories*.
  e.g., **john threw the ball** and **mary likes a cat** are syntactically identical

- Categories combine to create structures.
  article + *noun* = *noun-phrase* $[_{NP}$ **the**$_{art}$ **ball**$_{n}]$
  verb + *noun-phrase* = *verb-phrase* $[_{VP}$ **threw**$_{v}$ $[_{NP}$ **the**$_{art}$ **ball**$_{n}]]$
  noun-phrase + *verb-phrase* = *sentence* $[_{S}$ **john**$_{np}$ $[_{VP}$ **threw**$_{v}$ $[_{NP}$ **the**$_{art}$ **ball**$_{n}]]]$

In logic:

- Basic elements aren't particular sentences, but sentence types
- Sentence types combine to create arguments

So formal logic can serve as a foundation for semantics.

2. PROPOSITIONAL LOGIC

A proposition is a unit expressing a fact...basically, a sentence in the indicative mood. For example:

- Snow is white.
- Socrates is mortal.
- John will give me the ball.

But not:

- Give me the ball.
- Can I have the ball?

Basic fact about propositions: a proposition may be either true or false. Some true propositions:

- Mark Twain is the author of *Tom Sawyer*.
- Philadelphia is in Pennsylvania.

Some false propositions:

- George W. Bush is from Ireland.
- Natural language can be modeled with FSAs.

Some more propositions—are they true or false? (Answer: yes.)

- Jessica's middle name is "Lynn".
- The population of the United States is greater than 306,129,207.

2.1. Propositional Logic: A Definition

What we want here, as with syntax, is a way to represent propositions at a more abstract level. To that end, philosophers devised a language to represent propositions, which we'll call *PL*.

- Σ = {**p**, **q**, **r**, ..., ∧, ∨, →, **~**, **[**, **]**}
- **p**, **q**, **r**, ..., are propositions.
- **[** and **]** are brackets.
- Spaces can be used freely.

We then have strings over Σ, e.g. **p**, ∧, **pqp ~** ∨**p[[[q**, **[p** ∨ **q]** → **~[q** ∧ **r]**. But only some of these are strings of propositional logic. We'll define the set *PL* recursively as follows:

- The propositions **p**, **q**, **r**, ... ∈ *PL*.
- If $x \in PL$, then **[~***x***]** ∈ *PL*.
- If $x, y \in PL$, then **[***x* ∧ *y***]** ∈ *PL*.
- If $x, y \in PL$, then **[***x* ∨ *y***]** ∈ *PL*.
- If $x, y \in PL$, then **[***x* → *y***]** ∈ *PL*.
- Nothing else is in *PL*.

Question: are the following in *PL*?

- **p**
- **[[p** ∨ **p]** ∨ **[q** ∨ **q]]**
- **[p** → **~[q** ∨ **r]]**
- **[p** ∨ **[p** ∨ **[p** ∨ **p]]**
- **[~[p** → **r]]**

Finally, we need a semantics for the language:

- Each proposition **p**, **q**, **r**, ... may be true or false.
- **[~***x***]** is true iff *x* is false.
- **[***x* ∧ *y***]** is true iff *x* and *y* are both true.
- **[***x* ∨ *y***]** is false iff *x* and *y* are both false.
- **[***x* → *y***]** is false iff *x* is true and *y* is false.

2.2. Proofs in propositional logic

When we talked about formal reasoning at the start of the semester, we considered two types of reasoning: "within the system", which involved following a series of rules of proof, and "outside the system", which involved reasoning at a higher level.

The former is possible. For example, to prove that **[p** → **p]** is always true, given three axioms and two rules of inference:

$[[\mathbf{s} \rightarrow [\mathbf{p} \rightarrow \mathbf{q}]] \rightarrow [[\mathbf{s} \rightarrow \mathbf{p}] \rightarrow [\mathbf{s} \rightarrow \mathbf{q}]]]$	1. Axiom #2
$[[\mathbf{s} \rightarrow [\mathbf{r} \rightarrow \mathbf{q}]] \rightarrow [[\mathbf{s} \rightarrow \mathbf{r}] \rightarrow [\mathbf{s} \rightarrow \mathbf{q}]]]$	2. Rule of substitution (**r** for **p**)
$[[\mathbf{s} \rightarrow [\mathbf{r} \rightarrow \mathbf{p}]] \rightarrow [[\mathbf{s} \rightarrow \mathbf{r}] \rightarrow [\mathbf{s} \rightarrow \mathbf{p}]]]$	3. Rule of substitution (**p** for **q**)
$[[\mathbf{p} \rightarrow [\mathbf{r} \rightarrow \mathbf{p}]] \rightarrow [[\mathbf{p} \rightarrow \mathbf{r}] \rightarrow [\mathbf{p} \rightarrow \mathbf{p}]]]$	4. Rule of substitution (**p** for **s**)
$[[\mathbf{p} \rightarrow [\mathbf{q} \rightarrow \mathbf{p}]] \rightarrow [[\mathbf{p} \rightarrow \mathbf{q}] \rightarrow [\mathbf{p} \rightarrow \mathbf{p}]]]$	5. Rule of substitution (**q** for **r**)
$[\mathbf{p} \rightarrow [\mathbf{q} \rightarrow \mathbf{p}]]$	6. Axiom #1
$[[\mathbf{p} \rightarrow \mathbf{q}] \rightarrow [\mathbf{p} \rightarrow \mathbf{p}]]$	7. Rule of modus ponens (6 + 5)
$[[\mathbf{p} \rightarrow [\mathbf{q} \rightarrow \mathbf{p}]] \rightarrow [\mathbf{p} \rightarrow \mathbf{p}]]$	8. Rule of substitution ($[\mathbf{q} \rightarrow \mathbf{p}]$ for **q**)
$[\mathbf{p} \rightarrow \mathbf{p}]$	9. Rule of modus ponens (6 + 8)

—Alonzo Church, *Introduction to Mathematical Logic*, p. 81

If that kind of thing looks exciting, there's surely a course out there that covers it. We'll skip ahead to the point where it's been proven that this is equivalent to using a TRUTH TABLE: a visual display of all the options for things to be true or false, which makes it a way to reason outside the system.

[~x**]** is true iff x is false ≡ "it is not the case that x"

p	[~p]
T	F
F	T

[$x \wedge y$**]** is true iff x and y are both true. ≡ "x and y"

p	q	[p ∧ q]
T	T	T
T	F	F
F	T	F
F	F	F

[$x \vee y$**]** is false iff x and y are both false. ≡ "x or y"

p	q	[p ∨ q]
T	T	T
T	F	T
F	T	T
F	F	F

[$x \rightarrow y$**]** is false iff x is true and y is false. ≡ "If x, then y"

p	q	[p → q]
T	T	T
T	F	F
F	T	T
F	F	T

It may not be immediately clear why the fourth of these corresponds to "if...then". The first two lines should be pretty clear—suppose **p** = "It's raining" and **q** = "John is carrying an umbrella." Then "If it is raining, then John is carrying an umbrella", which is **[p → q]**...

- In cases where it's raining and John is carrying an umbrella, the sentence is true.
- In cases where it's raining and John is not carrying an umbrella, the sentence is false.

- And then, in cases where it's not raining and John is carrying an umbrella, the sentence is true, according to that table. Why?

  Because if I say the sentence, I'm saying that **if it's raining**, so-and-so is true. If it's not even raining, how can I be lying? (Compare other sentences that start with something untrue in an *if* clause: "If your first name starts with a Q, you're passing this class." Question: does the truth of that depend on there being *someone else* whose name starts with a Q?)

With truth tables, we can take sentences and determine whether they're necessarily true; and, similarly, take arguments and determine whether, solely based on their forms, they're valid. So, letting **p** = "It's raining", **q** = "John is carrying an umbrella", and **r** = "It's snowing"....

If it's raining, then it's raining. = **[p → p]**

Always true!

p	p	[p → p]
T	T	T
F	F	T

Either it's raining, or it's not raining. = **[p ∨ [~p]]**

Always true!

p	[~p]	[p ∨ [~p]]
T	F	T
F	T	T

If it's raining, John is carrying an umbrella. **[p → q]**
John is not carrying an umbrella. **[~q]**
∴ It is not raining. **[~p]**

If premises are true, conclusion is true!

p	q	[p → q]	[~q]	[~p]
T	T	T	F	F
T	F	F	T	F
F	T	T	F	T
F	F	T	T	T

Either it is raining or it is snowing. **[p ∨ r]**
It is not snowing. **[~r]**
∴ It is raining. **p**

If premises are true, conclusion is true!

p	r	[p ∨ r]	[~r]	p
T	T	T	F	T
T	F	T	T	T
F	T	T	F	F
F	F	F	T	F

Definition: a TAUTOLOGY is a sentence which, solely based on its form, is necessarily true.
Definition: a VALID argument is one in which, solely based on its form, whenever its premises are true, its conclusion is necessarily true.

2.3. The limits of propositional logic

The good news for semantics is that propositional logic does a good job of modeling connected propositions, and giving a meaning for the connected form based on the meanings of the pieces. (Well, at least insofar as that meaning will be "true" or "false".) Sentences like "If you've taken LING 106, you can draw FSAs" and "The exam will be hard, or the exam will not be hard" are possible to evaluate.

Unfortunately...

2.3.1. The real meanings of if *and* and

The claim above: “If *p*, *q*” = **[p → q]**, which is true unless **p** is true and **q** is false.

- This is great for “If you fail the exam, you’ll fail the class.” (True in cases where: you fail the exam and you fail the class; you pass the exam and you pass the class; you pass the exam and you fail the class.)

- “If Abe Lincoln was a Russian spy, then I’m a professor of linguistics.”
  In fact, I am a professor of linguistics...is this still true?

- “If you’re thirsty, there’s beer in the fridge.”
  If you’re not thirsty and there’s no beer in the fridge, is this true?

- “If you see a unicorn, you should shoot it.”
  If you never see a unicorn, is this true?

The claim above: “*p* and *q*” = **[p ∧ q]**, which is true if **p** is true and **q** is true; and
“*p* or *q* = **[p ∨ q]**, which is true if **p** is true or **q** is true.

- “I started lecturing, turned on my laptop, and came into the room.”
  Really? That’s true?

- “I wrote a letter to my grandmother yesterday, and six men can fit in the back seat of a Ford.”
  ...what? Does it even make sense to call this “true” or “false”?

- Either it’s April, or George W. Bush is an Irish Penn student.
  ...again, this technically true, but...?

2.3.2. Doing anything else with the semantics

None of this helps provide meanings to the propositions themselves, or to arguments that involve propositions other than the ones with “if”, “and”, and so on:

	All humans are mortal.	**p**
	<u>Socrates is a human.</u>	**q**
∴	Socrates is mortal.	**r**

Instead, we need to know that **human** = </hjumən/, *noun*, ‘sapient bipedal mammal’>, that **mortal** = </mortəl/, *adjective*, ‘subject to death’>, and from there, build the meaning of the proposition **all humans are mortal** from the meaning of all (whatever that is), the meaning ‘sapient bipedal human’, and the meaning ‘subject to death’, just as we built the sentence from *noun phrase* (*article* + *noun*) and *verb phrase* (*verb* + *adjective*).

3. The Goal of Semantics

We want to know the "meaning" of sentences. What do we mean by "meaning"? One useful way to look at this: *truth conditions*.

- That is: under what circumstances is a sentence true? If you know that, you know its meaning.
    - For instance: **the circle is in the square**—you'd recognize whether that's true or false for any given condition.
- Of course, those circumstances need to come from the meanings of the parts of the sentence. Propositional logic was step in that direction:
    - If it is raining, John is carrying an umbrella. **[p → q]**

  True in what circumstances? Those circumstances in which it's raining and John is carrying an umbrella, or in which it's not raining.
- But we need to go deeper than propositional logic—in what circumstances is it true that "Socrates is a human"?

4. Predicates

Here's the intuition. The sentences

(11) Socrates is a human.
The Archon of Ptharn is a squorfle.
George W. Bush is a Penn student.

are "about" Socrates, the Archon of Ptharn, and George W. Bush, respectively. And Socrates, the Archon of Ptharn, and George W. Bush are all **individuals**. So what are *human / squorfle / Penn student*?

- Answer: "predicates". That is: they're things that don't have truth conditions themselves (in what circumstances is "human" or "Penn student" true?), but they give truth conditions once they have subjects.

We can think of predicates as sets, and *subject + predicate* being "true" if the subject is in the set. (Also read as "the predicate holds of the subject".)

So rather than being about *propositions* (and their connectors)...

(12) If it is raining, John is carrying an umbrella. **[p → q]**
John is not carrying an umbrella. **[~q]**
∴ It is not raining. **[~p]**

...the classical, Aristotelian syllogism is about *predicates*.

(13) All humans are mortal. (humans, mortals)
Socrates is a human. (humans)
∴ Socrates is mortal. (mortals)

4.1. Aside #1: Parts of Speech

Syntacticians use “predicate” to mean “verb phrase”. Here it looks like we’re using “predicate” to mean “noun”. Should we worry? In fact...

- Every child **produces art**. (verb)
- Every child is **artistic**. (adjective)
- Every child is **an artist**. (noun)

Answer: no. From the point of view of semantics, these mean (roughly) the same thing.

4.2. Aside #2: Plurals

We have *all humans are mortal* (plural). Of course, “human” isn’t plural in other sentences, like “Socrates is human” or “Someone is human”. Again, should we worry?

- All humans are mortal. / Every human is mortal. (Synonymous.)
- Some woman is mortal. / Some women are mortal. (Kind of synonymous?)
- All unicycles have wheels. (Just confusing.)

Answer: no. Someone should worry about plurals, especially in that last sentence, but not us.

4.3. Predicate Logic: A Definition

We can define Predicate Logic the way we defined Propositional Logic—as a language with an alphabet, syntax, and semantics.

- In fact, we still need propositional logic to account for connected sentences, so we’ll still have the elements of propositional logic in our lexicon:

  $\{\sim, \vee, \wedge, \rightarrow, [,]\} \subset \Sigma$

  ...as well as, when we get there, the same syntax and semantics for them.

- We'll also need to add individuals and predicates to the system, as well as a way of connecting them:

  $\{\mathbf{A}, \mathbf{B}, ..., \mathbf{Z}, \mathbf{a}, \mathbf{b}, ..., \mathbf{w}\} \subset \Sigma$

  where capital letters represent predicates, lowercase letters represent individuals, and once we lay out the syntax, we'll connect them.

The full definition is as follows.

- $\Sigma = \{\mathbf{\sim}, \vee, \wedge, \rightarrow, \mathbf{[}, \mathbf{]}, \mathbf{A}, ..., \mathbf{Z}, \mathbf{a}, ..., \mathbf{w}, x, y, z, \forall, \exists, \mathbf{=}\}$, where
  - **CAPITAL LETTERS** represent predicates;
  - **lowercase letters** represent individuals.
  - *italicized lowercase letters* represent variables over individuals (and we'll reserve *x*, *y*, *z* for variables).

- *F*, the set of formulas of Predicate Logic, is defined recursively as follows.
  - If *A* is a predicate and *b* is an individual[1], then $Ab \in F$.
  - If *a*, *b* are individuals, then $\mathbf{[a = b]} \in F$.
  - If $m, n \in F$, then $\mathbf{[\sim} m\mathbf{]}$, $\mathbf{[}m \wedge n\mathbf{]}$, $\mathbf{[}m \vee n\mathbf{]}$, $\mathbf{[}m \rightarrow n\mathbf{]} \in F$.
  - If $m \in F$ and *a* is a variable, then $\mathbf{[}\forall a\ m\mathbf{]}$ and $\mathbf{[}\exists a\ m\mathbf{]} \in F$.
  - Nothing else is in *F*.

  For instance, some formulas:

  - **Hs**, **H**x, **[Hs** $\wedge$ **Ms]**, **[H**x $\wedge$ **[s =** x**]]**, **[**$\forall x$ **Hs]**, **[~[**$\forall x$ **H**x**]]**, **[**$\exists x$ **[H**x $\wedge$ **[s =** x**]]]**

  Some non-formulas:
  - **H**, **Hs**, x, **[Hs** $\wedge$ **Ms** $\wedge$ **Ps]**, **[Hs = Ps]**, **[**$\forall$**s Hs]**

And one more definition:

- In the formulas **[**$\forall a\ m$**]** and **[**$\exists a\ m$**]**, any instance of *a* appearing in *m* is BOUND. For example:
  - In **[**$\forall x$ **H**x**]**, *x* is bound.
  - In **[**$\forall x$ **[H**$x \rightarrow$ **M**y**]]**, *x* is bound, and *y* is not.
  - In **[**$\exists y$ **[**$\forall x$ **[H**$x \rightarrow$ **M**y**]]]**, both *x* and *y* are bound.
  - In **[[**$\exists y$ **[**$\forall x$ **[H**$x \rightarrow$ **M**y**]]]** $\wedge$ **H**y**]**, the last *y* is not bound.

  A formula is a PROPOSITION if it contains no unbound variables.

[1] Where "an individual" means either a lowercase letter or a lowercase italicized letter—i.e., a specific individual or a variable.

And finally, the semantics—i.e., the conditions under which a formula is true.

- $[\sim m]$, $[m \wedge n]$, $[m \vee n]$, $[m \rightarrow n]$ are true if [insert facts from propositional logic here].
- If *A* is a predicate and *b* is an individual, then *Ab* is true if the predicate holds of the individual.
  - e.g., if **H** = "is human" and **s** = "Socrates", then **Hs** is true.
  - Whether **H***x* is true will depend on who *x* is.
- $[a = b]$ is true if *a* and *b* are the same individual.
- $[\forall a\ m]$ is true if substituting *any* individual for *a* in *m* makes *m* true.
- $[\exists a\ m]$ is true if substituting *some* individual for *a* in *m* makes *m* true.

4.4. A Very Small World

	Humans (**H**)	Cats (**C**)	Deities (**D**)
Male (**B**)	Socrates (**s**) Plato (**p**)	Mr. Kitty (**k**)	Hades (**h**)
Female (**G**)	Lysistrata (**l**)	Fluffy (**f**)	Athena (**a**) Demeter (**d**)

Gray shaded squares: Mortals (**M**)

Question: In this world, which of the following are true?

- **Hs**
- **Ms**
- **Hk**
- **Mk**
- $[\mathbf{Hs} \wedge \mathbf{Ms}]$
- $[\mathbf{Hk} \wedge \mathbf{Mk}]$
- $[\mathbf{Hk} \vee \mathbf{Mk}]$
- **Mk**
- $[\forall x\ \mathbf{H}x]$
- $[\exists x\ \mathbf{H}x]$
- $[\exists x\ [\mathbf{H}x \wedge \mathbf{G}x]$
- $[\exists x\ [\mathbf{H}x \wedge \mathbf{C}x]$
- $[\exists x\ [\mathbf{H}x \rightarrow \mathbf{C}x]$
- $[\forall x\ [\mathbf{H}x \rightarrow \mathbf{M}x]$

How does this help with natural language? Well...

- *All humans are mortal* = $[\forall x\ [\mathbf{H}x \rightarrow \mathbf{M}x]]$

  Why? Because this is only false when there's some x for which $[\mathbf{H}x \rightarrow \mathbf{M}x]$ is false, i.e. some x for which $\mathbf{H}x$ is true and $\mathbf{M}x$ is false—i.e., when there's someone who's human, but not mortal.

- *Some female is mortal* = $[\exists x\ [\mathbf{G}x \wedge \mathbf{M}x]]$

  Why? Because this is true as long as there's some x for which $[\mathbf{G}x \wedge \mathbf{M}x]$ is true, i.e. some x for which $\mathbf{G}x$ and $\mathbf{M}x$ are both true—i.e., when there's someone who's both female and mortal.

4.5. Digression: back to the form of arguments

Now that we can represent propositions, we may want to determine whether a given argument is valid, just as we did for propositional logic—perhaps by setting up a truth table?

(14)	All humans are mortal.	$[\forall x\ [\mathbf{H}x \rightarrow \mathbf{M}x]]$
	Socrates is a human.	$\mathbf{Hs}$
∴	Socrates is mortal.	$\mathbf{Ms}$

(15)	All humans are mortal.	$[\forall x\ [\mathbf{H}x \rightarrow \mathbf{M}x]]$
	Some female is human.	$[\exists x\ [\mathbf{G}x \wedge \mathbf{H}x]]$
∴	Some female is mortal.	$[\exists x\ [\mathbf{G}x \wedge \mathbf{M}x]]$

The problem is that we'd need to consider all of the cases, as we did for propositions. But rather than two things that are each true or false—suppose we have just two individuals, **s** and **p**. Then we need to consider *sixteen* cases: neither is human and either neither is mortal, **s** is mortal, **p** is mortal, or both are mortal; **s** is human and either neither is mortal, **s** is mortal, **p** is mortal, or both are mortal; and so on.[2]

Instead, we need a few rules of inference, the way we had for shapes:

- From $[\forall x\ m]$, conclude m with any constant replacing each occurrence of x.
  - Example: from $[\forall x\ [\mathbf{H}x \rightarrow \mathbf{M}x]]$, conclude $[\mathbf{Hs} \rightarrow \mathbf{Ms}]$.

- From m, conclude $[\exists x\ m]$, with x replacing each occurrence of any one constant.
  - Example: From $[\mathbf{Hs} \rightarrow \mathbf{Ms}]$, conclude $[\exists x\ [\mathbf{H}x \rightarrow \mathbf{M}x]]$.

[2] In general, there will be $2^{I \cdot P}$, where I is the number of individuals and P is the number of properties. It gets really unwieldy really fast.

For instance, for the syllogisms on the previous page:

(14′)	$[\forall x\ [\mathbf{H}x \rightarrow \mathbf{M}x]]$	(given)
	$[\mathbf{Hs} \rightarrow \mathbf{Ms}]$	(rule of inference)
	$\mathbf{Hs}$	(given)
	$\therefore \mathbf{Ms}$	($[\mathbf{p} \rightarrow \mathbf{q}]$ and $\mathbf{p}$, infer $\mathbf{q}$—*modus ponens*, from prop. logic)

(5′)	1.	$[\exists x\ [\mathbf{G}x \wedge \mathbf{H}x]]$	(given)
	2.	$[\mathbf{Ga} \wedge \mathbf{Ha}]$	(assumption from 1)
	3.	$\mathbf{Ga}$	(inference from 2, prop logic)
	4.	$\mathbf{Ha}$	(inference from 2, prop logic)
	5.	$[\forall x\ [\mathbf{H}x \rightarrow \mathbf{M}x]]$	(given)
	6.	$[\mathbf{Ha} \rightarrow \mathbf{Ma}]$	(inference from 5)
	7.	$\mathbf{Ma}$	(inference from 4+6, *modus ponens*)
	8.	$[\mathbf{Ga} \wedge \mathbf{Ma}]$	(inference from 3+7, prop logic)
	9.	$[\exists x\ [\mathbf{G}x \wedge \mathbf{M}x]]$	(inference from 8)
	$\therefore$	$[\exists x\ [\mathbf{G}x \wedge \mathbf{M}x]]$	(un-assumption)

But once again, that kind of thing can be left to logicians. (i.e., not us.) The question for us as linguists is: how helpful is this in modeling language?

5. PREDICATE LOGIC AS A MODEL FOR LANGUAGE

So far, it's not too bad; we can model sentences like

- Socrates is mortal.
- All humans are mortal.
- Some human is mortal.
- Socrates is mortal and Plato is not mortal.
- If all humans are mortal, then Socrates is mortal.

What else?

5.1. Modeling no

Let's try *no deity is mortal*. In the same way that *all S are P* = $[\forall x\ [\mathbf{S}x \ldots \mathbf{P}x]]$ and *some S are P* = $[\exists x\ [\mathbf{S}x \ldots \mathbf{P}x]]$, we could introduce a new symbol and say *no S are P* = $[\mathbf{N}x\ [\mathbf{S}x \wedge \mathbf{P}x]]$, where $\mathbf{N}$, an upside-down N on analogy with the upside-down A and E, represents "no"; we'd add a rule that says

- $[\mathbf{N}a\ m]$ is true if substituting *no* individual for *a* in *m* makes *m* true.

i.e., no individual *x* makes $[\mathbf{D}x \wedge \mathbf{M}x]$ true, so *No deity is mortal* = $[\mathbf{N}x\ [\mathbf{D}x \wedge \mathbf{M}x]]$ is true.

But instead...

- This could be a statement about all individuals, just like *all humans are mortal*: instead of saying that any individual with the "human" property has the "mortal" property, it says that any individual with the "deity" property lacks the "mortal" property. Since this is like *all*, we don't need the **N**; we can write **[**∀*x* **[D***x* → **[~M***x***]]]**. (Or perhaps we want **[**∀*x* **[M***x* → **[~D***x***]]]**? Is that the same thing?)

- Or maybe this is a statement about which individuals exist, just like *some humans are mortal*: instead of saying that you can find an individual with both the "human" property and the "mortal" property, it says that you can't find any individual with the "deity" property and the "mortal" property. Since this is like *some*, once again we don't need the **N**; we can write **[~[**∃*x* **[D***x* ∧ **M***x***]]]**.

So which one is right? In fact: they're identical!

> PROOF: Let's call **[**∀*x* **[D***x* → **[~M***x***]]]** "Formula A". When is Formula A false? When there's some *x* for which **[D***x* → **[~M***x***]]** is false. That is to say, Formula A is false when there's some *x* such that **D***x* is true, but **[~M***x***]** is false, since that's the only way to make the "→" formula false. So Formula A is false when there's some *x* such that **D***x* and **M***x* are both true, which we can write as **[**∃*x* **[D***x* ∧ **M***x***]]**. If Formula A is false only when that's true, then Formula A is true whenever its negation is true—that is, Formula A is true exactly when **[~[**∃*x* **[D***x* ∧ **M***x***]]]** is true. So they're equivalent. QED

Exercise: Prove that **[**∀*x* **[D***x* → **[~M***x***]]]** is equivalent to **[**∀*x* **[M***x* → **[~D***x***]]]**.
Exercise: Prove that **[~[**∃*x* **[D***x* ∧ **M***x***]]]** is equivalent to **[**∀*x* **[~[D***x* ∧ **M***x***]]]**.

5.2. *Modeling* at least [number]

- To model *At least one cat is mortal*...in fact, this is the same as "some cat is mortal".

- To model *At least two cats are mortal*: we can write this as "some cat is mortal, and some other cat is mortal". Perhaps:

  [[∃*x* **[C***x* ∧ **M***x***]]** ∧ **[**∃*y* **[C***y* ∧ **M***y***]]]**

Except that this turns out to be true in our small world from several pages back, even though there's only one mortal cat! Why? Because: (a) the formula to the left of the ∧ is true: there is some *x* for which **[C***x* ∧ **M***x***]** is true (namely, *x* = **k**); and (b) the formula to the right of the ∧ is true: there is some *y* for which **[C***y* ∧ **M***y***]** is true (namely, *y* = **k**).

What's missing is the "other" from "some other cat". We need the following (note that some brackets have been left out for readability):

∃*x* ∃*y* **[[C***x* ∧ **M***x***]** ∧ **[C***y* ∧ **M***y***]** ∧ **~[***x* **=** *y***]]**

Now, if we take x = **k**, the inside formula is false: there's no y we can pick to make all three parts true. If y = **f**, **[C**y ∧ **M**y**]** is false because **Mf** is false; if y = **s**, **[C**y ∧ **M**y**]** is false because **Cs** is false; if y = **k**, **[~[**x **=** y**]]** is false because **[~[k = k]]** is false—in fact **k** and **k** are the same individual.

- To model *At least three cats are mortal*: we do much the same, with clauses to express that none of x, y, z are the same:

$$\exists x\, \exists y\, \exists z\, [[\mathbf{C}x \wedge \mathbf{M}x] \wedge [\mathbf{C}y \wedge \mathbf{M}y] \wedge [\mathbf{C}z \wedge \mathbf{M}z] \wedge \sim[x = y] \wedge \sim[x = z] \wedge \sim[y = z]]$$

- To model *At least fifty cats are mortal*: we do much the same, except that it starts with fifty ∃n for fifty different variables n; has fifty conjoined **[C**n ∧ **M**n**]** clauses, and ends with 1,225 different **[~[**m **=** n**]]** clauses.

  You wouldn't want to write that out. Then again, you wouldn't want to write out a lot of possible FSAs; it suffices to know *how* to do it, i.e. that it *can* be done.[3]

- To model *Fewer than two cats are mortal*: all we need to do is treat this as "it's not true that at least two cats are mortal", i.e. **[~[**(formula for *at least two cats are mortal*)**]]**

5.3. Even more sentences

- *Socrates is a mortal human*.

  This is equivalent to *Socrates is mortal and Socrates is human*, i.e. **[Ms** ∧ **Hs]**.

- *Socrates, who is human, is mortal*.

  This is also equivalent to *Socrates is mortal and Socrates is human*.

- *Socrates taught Plato*.

  We can't really do that, except by introducing some predicate **Q** that represents "taught Plato"—but that's kind of like just using **p** to represent "Socrates is human". It misses the fact that there's a second individual in there.

  What we *can* do is modify predicate logic to include predicates that take *pairs* instead of *individuals* as arguments. For instance, just as **H**x is true if x is human, we might say that **T**$\langle x,y\rangle$ (or **T**xy, or however we write it) is true if x taught y; then **T**$\langle$**s**,**p**$\rangle$ represents "Socrates taught Plato".

---

[3] Computer scientists might want to write $[\exists x_1, x_2, ..., x_{50}\ [\forall n : 1 \leq n \leq 50\ [[\mathbf{C}x_n \wedge \mathbf{M}x_n] \wedge \forall m : 1 \leq m \leq 50\ [x_m = x_n \leftrightarrow m = n]]]]$. It's a lot more concise. Unfortunately, it's not a formula of predicate logic, at least not unless we introduce, among other things, numbers.

6. Some Problem Cases

So far, so good. But how accurate was all this?

6.1. Existential Import

	Humans (**H**)	Cats (**C**)	Deities (**D**)	Unicorns (**U**)
Male (**B**)	Socrates (**s**) Plato (**p**)	Mr. Kitty (**k**)	Hades (**h**)	
Female (**G**)	Lysistrata (**l**)	Fluffy (**f**)	Athena (**a**) Demeter (**d**)	

Gray shaded squares: Mortals (**M**)

- *Some unicorn is male* is $[\exists x\, [\mathbf{U}x \wedge \mathbf{B}x]]$.

  Technically this comes out false: we can't find any individuals for which $\mathbf{U}x$ and $\mathbf{B}x$ are both true. The sentence is also pretty obviously false, so that's fine.

- *All unicorns are male* is $[\forall x\, [\mathbf{U}x \rightarrow \mathbf{B}x]]$.

  Technically this comes out true: no matter what individual we choose, if $\mathbf{U}x$ is true, then $\mathbf{B}x$ is also true. (After all, $\mathbf{U}x$ is always false.) The sentence is...less obviously true.

  What about *All unicorns have one horn*? *All unicorns have two horns*? (*All Penn students under the age of six live in college housing*?)

This is a classical debate: does "All *S* are *P*" require that there actually are some *S* things? Obviously "Some unicorns are white" can only be true if there are unicorns. But is "All unicorns are white" true even if there are no unicorns? Or is it false? Or is it kind of weirdly neither?

And what about "No *S* are *P*"—is it true/false/neither to say that "No unicorns are white"? (*No unicorns are mortal*? *No unicorns are mammals*?)

The typical classical answer is that things just work out better if we assume that there must be some *S* things. Thus, the classical Square of Opposition:

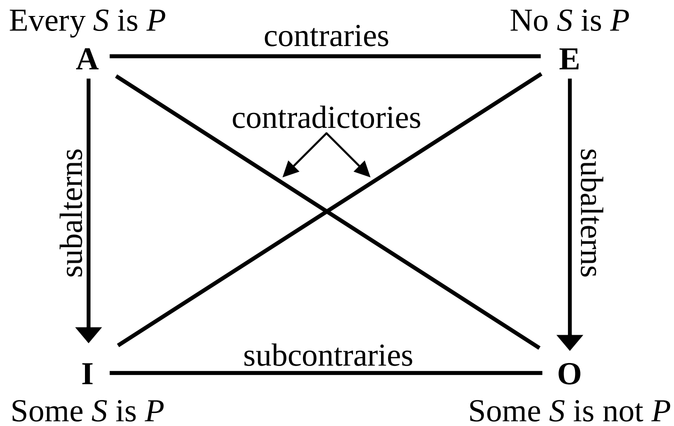

- CONTRADICTORY propositions cannot both be true and cannot both be false.

  e.g., *some cats are immortal* vs. *no cats are immortal* – exactly one of these is true.

- CONTRARY propositions cannot both be true (but can both be false).

  e.g., *all cats are immortal* vs. *no cats are immortal* – these can't both be true

- SUBCONTRARY propositions cannot both be false (but can both be true).

  e.g., *some cats are mortal* vs. *some cats are immortal* – these can't both be false

- SUBALTERNS: if the upper one is true, the lower is true.

  e.g., if *all cats are mortal*, then certain *some cat is mortal*

For this square to hold, *all S are P* must be <u>false</u> when there aren't any *S* things—i.e., the sentence "All unicorns are white" is false if there aren't any unicorns. That's the opposite of the way things go in $[\forall x\ [\mathbf{U}x \rightarrow \mathbf{W}x]]$, which is <u>true</u> if there aren't any unicorns. (Which breaks the Square of Opposition, which is why it was abandoned in the 1800s.)

6.2. Even more Existential Import

We've done *all*, *some*, *at least two*...what about *the*?

- *The LING106 TA is male*.

  Perhaps something like: $\exists x$ **[T**$x \wedge$ **B**x**]**? That's true, since we can find an x, namely **n** (Satoshi), for whom "LING106 TA" is true and "male" is true.

- *The LING106 student is male*.

  In this case, $\exists x$ **[S**$x \wedge$ **B**x**]**, which is true, because we can find an x, for instance **j** (Joseph), for whom "LING106 student" is true and "male" is true.

...and that's kind of unfortunate, because *The LING106 student is male* isn't exactly true. So we need a better solution. This one comes from Bertrand Russell:

- *The LING106 TA is male* means:

$\exists x$ [$\mathbf{T}x \wedge$	there is some person x who's a LING106 TA, and
$\forall y$ [$\mathbf{T}y \rightarrow$ [$y = x$]] $\wedge$	anyone who's a LING106 TA must be person x, and
$\mathbf{B}x$]	person x is male

- This makes *The LING106 TA is male* true (because there's a person, Satoshi, who's a LING106 TA; and anyone who's a LING106 TA is Satoshi; and Satoshi is male.

- And *The LING106 student is male* is false: $\exists x$ [$\mathbf{S}x \wedge \forall y$ [$\mathbf{S}y \rightarrow$ [$y = x$]] $\wedge$ $\mathbf{B}x$] – there's no x that makes the bracketed part true. (For instance, if x = Joseph: it's true that he's a LING106 student and true that he's male, but it's false that anyone who's a LING106 student is Joseph.)

- Also: *The King of France is male* is false: $\exists x$ [$\mathbf{K}x \wedge \forall y$ [$\mathbf{K}y \rightarrow$ [$y = x$]] $\wedge$ $\mathbf{B}x$] – there's no x that makes the bracketed part true, since $\mathbf{K}x$ is false for any x.

That's not nearly as unfortunate...but it's not the most intuitive result. But there's even more:

- We proved that $[\mathbf{p} \vee [\sim\mathbf{p}]]$ is a tautology—for instance, *Either it's raining or it's not raining* must be true.

  And therefore, *Either the King of France is bald, or the King of France is not bald* must also be true.

Once again: not the most intuitive result. But what alternative do we have? $[\mathbf{p} \vee [\sim\mathbf{p}]]$ isn't necessarily true? (What can our logic look like for that not to be true?)

6.3. *More problems:* Most

- All humans are mortal: $[\forall x\ [\mathbf{H}x \rightarrow \mathbf{M}x]]$
- Some human is mortal: $[\exists x\ [\mathbf{H}x \wedge \mathbf{M}x]]$
- No human is mortal: $[\mathbf{N}x\ [\mathbf{H}x \wedge \mathbf{M}x]]$ (which worked, though we didn't need it)

- Most humans are mortal: $[\text{Ɯ}x\ [\mathbf{H}x \ldots \mathbf{M}x]]$? Is that possible?

  That is: can we say "for most *x*, (something about *x* being human and *x* being mortal)"?

Bringing back our Very Small World:

	Humans (**H**)	Cats (**C**)	Deities (**D**)
Male (**B**)	Socrates (**s**) Plato (**p**)	Mr. Kitty (**k**)	Hades (**h**)
Female (**G**)	Lysistrata (**l**)	Fluffy (**f**)	Athena (**a**) Demeter (**d**)

Gray shaded squares: Mortals (**M**)

- *Most humans are male* is true.
  But $[\text{Ɯ}x\ [\mathbf{H}x \wedge \mathbf{B}x]]$ = "for most *x*, *x* is a human and *x* is male" is not!
  (It's false for 6 of 8 individuals: **l**, **k**, **f**, **h**, **a**, **d**)

- *Most humans are female* is false.
  But $[\text{Ɯ}x\ [\mathbf{H}x \rightarrow \mathbf{B}x]]$ = "for most *x*, if *x* is a human then *x* is female" is not!
  (It's true for 6 of 8 individuals: **l**, **k**, **f**, **h**, **a**, **d**)

In general: nothing we do with ∀, ∃, or any other upside-down letters is going to get the right meaning for *most*—or for *many*, *few*, *half*...

And the way our system is set up*, the King of France is bald* is forced to be true or false, when in fact it's...well, it's not really either, but "neither" isn't really an option.

Therefore: we need something a step up from predicate logic.