The Conceptual Foundations of Truth-Conditional Semantics[1]

1. The Big Questions

(1) **Semantics** The study of *meaning* in natural language.

- But, what does it really *mean* to study 'meaning'?
- Well, considering the opening passage to Stephen Pinker's *The Language Instinct*

"As you are reading these words, you are taking part in one of the wonders of the natural world. For you and I belong to a species with a remarkable ability: *we can shape events in each other's brains with exquisite precision…*
…Simply by making noises with our mouths, we can reliably cause precise new combinations of ideas to arise in each other's minds. The ability comes so naturally that we are apt to forget what a miracle it is. Asking you only to surrender your imagination to my words… I can make you think some very specific thoughts…"

(2) **Uncontroversial Idea:** Language is a system for encoding (and expressing) thought

(3) **Immediate Follow-Up Question:** *How*, exactly, does language encode thought?

- Classic 'productivity arguments' show us that it cannot be a system like the one in (4)

(4) **The Simplest System (Which Language is Not)**

A stipulated (memorized) list of symbols paired with their meaning:

/bɛl/ ⇔ 🔔
/hæpi/ ⇔ ☺
/dɛθ/ ⇔ ☠
…

So, if it's not a list like in (4), what kind of system is it?...

(5) **Common Way for Linguists to Frame the Question**
What does one have to know in order to understand the sentences of a language?

To get a handle on this question, consider the sentence in (6)…

[1] These notes are based on material in Heim & Kratzer (1998: 1-12, 13-26), Chierchia & McConnell-Ginet (2000: 1-33, 53-73, 99-104), Larson (1995: 361-368), and Partee (1995: 311-316).

(6) **Sentence of St'át'imcets (Lillooet; Salish, British Columbia)**

Kelkálenas tits'í7a na sqaxa7lhkálha.

To figure out what this sentence means, what do you have to know?

a. The meaning of the words (duh).

b. The structure of the sentence (what's modifying what, etc.)

c. <u>Often Overlooked Ingredient (Usually Implicit in Second Language Instruction):</u>
A set of rules for deriving the meaning of the sentence from (6a) and (6b).

(7) **Fundamental Conclusion about Natural Language Semantics**

The semantic competence of a native speakers of a language L must comprise:

a. A finite number of primitive meaningful units (lexical items, idioms)

b. A finite number of *rules* for deriving the meaning of a complex expression from:
(i) The meanings of the component expressions, *and*
(ii) The syntactic structure of the complex expression.

(8) **The Principle of Compositionality (see Partee 1995)**

The meaning of a complex expression (in natural language) can be effectively computed from (i) **the meaning of its component expressions**, and (ii) **their 'mode of combination'** (i.e., syntax of complex expression)

(9) **Another Path to the Conclusions in (7) and (8)**
The meaning of a complex expression in natural language depends upon both:

a. <u>The meaning of its component expressions</u>
After all, if you change the words in a sentence, you change its meaning:

The dog bit the man. VERSUS *The dog bit the cat.*

b. <u>The syntactic structure of the expression</u>
After all, if you change the syntactic structure, you also change the meaning.

The dog bit the man. VERSUS *The man bit the dog.*

So, part of what's in our brains is a system ***of rules*** *for deriving the meanings of complex expressions…*

(10) **Burning Question:**
What are these semantic rules that our brains use for computing meanings?

(11) **Unfortunate (or Fortunate) Fact:**
Our knowledge of these rules is 'tacit' (unconscious), and so we cannot simply introspect them.

- Rather, to answer the question in (10), we have to do science; we have to **frame and test hypotheses!**

(12) **The Over-Arching Research Question (for Semanticists)**

What is the system of rules that our cognitive systems employ to compute the meanings of complex expressions (in a way consistent with the Principle of Compositionality (8))

(13) **Some Related, More Specific Research Questions**

a. How does a human being acquire this system? How much is already specified by the biology of the organism? (semantic acquisition)

b. How does this system vary across languages? Do languages differ in how they compute meanings, and if so, in what ways? (semantic typology)

---

2. The Meaning of 'Meaning'

(14) **An Immediate Problem for Our Project: What the Heck is a 'Meaning'**

- To answer the question in (12), we want to develop some hypotheses about what the system is like, and then test them…

- So, we'll want to design a hypothetical formal system that will manipulate primitive '**meanings**' to derive the '**meanings**' of more complex sentences.

- But, *how do we formally represent the '**meaning**' of a sentence or its component phrases?*

- What is the '**meaning**' of "Barack" and the '**meaning**' of "smokes" such that 'combining them together' gives us the '**meaning**' of "Barack smokes"?

Meaning($[_{NP}$ Barack]) + Meaning($[_{VP}$ smokes]) = Meaning($[_{S}$ Barack smokes])

<u>Over-Arching Problem:</u>

- The word 'meaning' is a vague, pre-theoretic term from every-day discourse.
  (like 'alive', 'hot', 'rock', 'heavy')

- Thus, it may not be an appropriate term for a precise, scientific study of human language.
  (*e.g.* biology doesn't actually employ terms like 'alive' or 'dead' or 'life-form')
  (*e.g.* physics doesn't actually employ terms like 'hot', 'heavy', 'fast', *etc.*)

(15) **The General Issue**

- Sometimes, our 'everyday words' for things don't actually 'carve nature at her joints'.

  - That is, it's often the case that our everyday words lump together things that turn out to be – upon further inquiry – actually quite different.

  - A well-known example of this in the philosophy of science is 'jade'. It turns out that 'jade' applies to two rather different minerals (jadeite and nephrite).

- As we'll see in a moment, it seems like the everyday concept 'meaning' is like this.

  - 'Meaning' is a rather vague cover-term for a whole host of different properties and phenomena related to language use.

(16) **New, Preliminary Objective**
Let's replace our pre-theoretic concept of 'meaning' with some more precise terminology, which will:

(i) Divide up 'meaning' into more properly identifiable (and manageable) subparts.

(ii) Thereby allow us to work towards a formal system that manipulates 'meanings'…

(17) **Side-Benefit of Technical Vocabulary: Better Description of 'Meanings'**

- In traditional grammars, particularly of 'minority languages', one often finds frustratingly vague descriptions of the meanings of certain elements.

- By employing a formal semantic theory, the researcher is forced to ask and answer deeper (and more precise) questions, resulting in greater depth of empirical coverage

- Moreover, without this technical semantic theory, certain (quite interesting) semantic questions about various languages *cannot even be framed*.

So how can we better pin down the phenomena/properties that we are interested in, those that are typically, loosely categorized under the general umbrella of 'meaning'?...

(18) **'Meaning' is as 'Meaning' Does**
"In order to say what a meaning is, we may first ask what a meaning *does*, and then find something that does that" (David Lewis; "General Semantics")

- What (exactly) do we *know* when we *know* 'the meaning' of a sentence?

(a) Social Appropriateness of the Sentence:
What social contexts the statement is appropriate in.

"That's wonderful." vs. "That kicks ass!"

(b) 'Emotional Content' of the Sentence:
What the statement reveals about the emotional state of the speaker.

"I disagree with Dave's judgment." vs. "Dave is a damn fool!"

(c) ***The Informational Content of the Sentence***
What information about the world the sentence conveys.

Rightly or wrongly, (18c) *has received by-far-and-away the greatest attention over the centuries*
It will also be the aspect of meaning that we will be concerned with in this course. Thus, let's try to further develop this notion of the 'information conveyed' by a sentence

---

2.1 The Different Ways that Information Can Be 'Conveyed'

When we examine this notion of a sentence 'conveying information', we find that it is not so simple either.
There seem to be different *ways* that information can be 'conveyed' by a sentence.

(19) **Example Dialog**

Person 1: How did Dave's physical go?
Person 2: Well, he's stopped smoking.

Person 2's utterance 'conveys' all the following information:
- Dave has stopped smoking.
- Dave *has been smoking*.
- Dave's physical did not go well. (Dave received bad news at his physical.)

Each of these different bits of information is 'conveyed' in a different way by the utterance.

(20) **Assertion**
The information that *Dave has stopped smoking* is **asserted** by the utterance / speaker

- **Assertion** = the information that is 'explicitly added' by the utterance

- Sentence S **asserts that** *p* = S is true if and only if *p*

Test:	"Dave stopped smoking" is true if and only if Dave stopped smoking.

(21) **Presupposition**
The information that *Dave has been smoking* is **presupposed** by the utterance / speaker

- **Presupposition** = the information that is 'taken for granted' by the utterance

- Sentence S **presupposes** *p* = S is true *or false* only if *p*
  S *and negation of S* is true only if p

Test:	"Dave stopped smoking" can only be true if Dave has been smoking. "Dave **didn't** stop smoking" can only be true if Dave has been smoking.

(22) **Implicature**
The info that *D's physical didn't go well* is an **implicature** of the utterance / speaker

- **Implicature** = the information that is not explicitly asserted in the utterance, but which the speaker (clearly) intends the listener to conclude.

- *p* is **an implicature** of S = *p* is 'conveyed' by S, but 'not *p*' is consistent with S

Test:	"Dave stopped smoking, but he did fine on his physical" is logically consistent.

The Main Point:

The information 'conveyed' by a sentence / utterance can be divided into (at least):
(a) The information *asserted* by the sentence / utterance
(b) The information *presupposed* by the sentence / utterance
(c) The *implicatures* of the sentence / utterance.

Consequently, if we want to understand the overall system by which complex sentences can 'convey' information, we will need to understand each of the following
(i) How the *assertions* of a sentence S are derived from the meanings of its parts.
(ii) How the *presuppositions* of S are derived from the meanings of its parts
(iii) How the *implicatures* of S are derived (in part) from the meanings of its parts.

Ultimately, a formal semantic theory will need to do all of (i) – (iii) above

However, we also need to start *somewhere*, and so (in this class) we will start with (i)…
(… *though we will also touch a bit on presuppositions and implicature later on*…)

(23) **Our Goal (Restated)**

Develop a theory of the rule system that derives the *assertions* of a complex (declarative) sentence from (i) the 'meanings' of its component expressions, and (ii) its syntactic structure

Side-Note: What about *non*-declarative sentences, like questions and imperatives?
They seem to also be meaningful, but they don't seem to 'assert' anything!...

Suspend your disbelief!
If you go on in semantics, you will find that a treatment of interrogatives and imperatives can be built using the formal tools we initially develop for the treatment of declaratives…

2.2 The Importance of 'Truth Conditions' to a Theory of Meaning

To build towards our (restated) goal in (23), let's introduce a new bit of terminology.

(24) **Truth Conditions**

The 'truth conditions' of a sentence S are the conditions under which S is true.

Truth-Conditional Statement: 'S is true if and only if *p*'

Some Consequences:

a. The 'truth conditions' of S are another name for the 'assertions' of S
b. Thus, our goal in (23) can again be restated to the following:

(25) **Our Goal (Restated Again)**

Develop a theory of the rule system that derives the *truth conditions* of a sentence from (i) the 'meanings' of its component parts, and (ii) its syntactic structure.

A Quick Review of How We Got Here:

a. We want to know how the 'meaning' of a sentence is computed from the 'meanings' of its parts (and the way they are syntactically combined).
b. This requires us to make more precise what we mean by 'meaning'.
c. This leads us to the notion of *the information that a sentence conveys*
d. This requires us to make more precise what we mean by 'conveying information'
e. This leads us to the notion of *the information 'asserted' by a sentence.*
f. This notion can be recast as *the truth conditions of a sentence*.
g. Thus, we want to know how the *truth conditions* of a sentence can be derived from the 'meanings' of its component parts, etc.

An Important Reminder:
As we've seen, 'truth-conditions' aren't all there is to the general phenomenon of 'meaning'.
At some point, we will have to come back to the other phenomena in (18), (21) and (22)….

Our restated goal in (25) carries a certain presupposition regarding the nature of our 'language faculty', which it is worth pausing to reflect on:

(26) **The 'Psychological Reality' of Truth-Conditions**

Our goal in (25) presupposes that *part of our cognitive capacity as speakers of a language is a system that derives* ***truth conditions****.*

This isn't so far fetched a claim… consider the following (plausible) characterization of the information computed during a typical conversation…

(27) **A Model of Information Computed During Sentence Comprehension**

a. Speaker's Utterance: / ðə haʊs ɪz an fajɹ /

b. Listener's Computations:

(i) *Syntax*: The string /ðə haʊs ɪz an fajɹ/ has the following structure:
[[the house][is[on[fire]]]]

(ii) *Semantics:* [[the house][is[on[fire]]]] is true *iff* the house is on fire

(iii) *Pragmatics:* The speaker is an honest guy, so he believes what he says...
The speaker is smart, so what he believes is true...
So "[[the house][is[on[fire]]]]" must be true…
So, **given its truth conditions,** the house must be on fire…
…oh my god the house is on fire!...

Observation:
In the model of human communication in (27), a central step in comprehension is the computation of the truth-conditions of the speaker's utterance:

- If a listener has reason to believe that a speaker's utterance is *true* (*e.g.* because the speaker is honest and knowledgeable)….

- …then knowing the **truth conditions** of a sentence allows the listener to thereby deduce *facts about the world!*

CONCLUSION:

To the extent to which (27) is an accurate characterization of the kind of information our systems compute when we 'understand' an utterance, ***then part of what our language systems 'do' during sentence comprehension is compute truth-conditions***

(28) **Brief Terminological Aside: 'Object Language' and 'Meta-Language'**

a. The Object Language:
The language we are *describing* (not necessarily using).

We want our theory to associate sentences of the *object* language with their truth conditions.

b. The Meta-Language:
The language we are *using* to characterize the truth conditions of the sentences of our *object language.*

We want our theory to associate sentences of the object language with their truth conditions, *as characterized in the meta-language*.

The object language and meta-language can both be *the same language*….

c. "The house is on fire" is T (true) *iff* the house is on fire.
Object language: English *Meta-Language: English*

…or, they can be two different languages:

d. "Het huis brandt" is T (true) *iff* the house is on fire.
Object language: Dutch *Meta-Language: English*

3. Obtaining a Theory of Truth Conditions from 'Extensions'

The following is a crude sketch of our over-arching goal in (25):

(29) **Crude Sketch of a Compositional Semantic Theory that Yields Truth Conditions**

SEMANTICS('Dave') + SEMANTICS('smoke') =

SEMANTICS('Dave smokes') =

TRUTH-CONDITIONS('Dave smokes')

Some Important Questions:

- How do we construct a theory that obtains this result for us?
- What are the 'meanings' of "Dave" and "smokes" such that combining them together can yield for us the truth conditions of "Dave smokes"?

... more in a moment!

3.1 More about the Meaning of 'Meaning': An Excursus of 'Extensions'

We've already seen that the everyday word 'meaning' is vague and confusing in a number of ways…*here's another:*

(30) **The Meaning of the Phrase 'The President'**

a. In one sense, the *meaning* of the NP "the president of the United States" changed in January 2009. It went from *meaning* George Bush (with all of the attendant connotations) to *meaning* Barack Obama.

('denotation', 'reference')

b. In another sense, the *meaning* of the NP is the same now as it was in 2008. It still *means* 'the person who holds the office of the presidency of the United States'.

('sense', 'concept')

Instead of using the word 'meaning' in this vague and ambiguous fashion, let's introduce two different terms to refer unambiguously to these two different 'senses' of the word "means".

(31) **Extension vs. Intension**

a. The *extension* of an NP is the thing in the world that the NP (currently) refers to.

- The *extension* of "the president" is *Barack Obama.*

b. The *intension* of an NP is the 'general concept' behind the NP, which determines (for a given time/situation) what the extension of the NP is.

- The *intension* of "the president" is *the person who holds the office…*

So, the 'meaning' of an NP can be broken up into its 'extension' and its 'intension'…
…can the meaning of a sentence likewise be broken up in this way?

(32) **Intension of a Sentence = 'Truth Conditions'**

We might take the 'intension' of a sentence to be (something like) its *truth conditions*…

As we've already seen, the truth conditions of a sentence are akin to what we might vaguely describe as the sentence's 'asserted content'.

… But if the 'intension' of a sentence is its truth conditions, what is its 'extension'?...

(33) **Extension of a Sentence = Truth Value**

If we take the 'intension' of a sentence to be its truth conditions, then we should take the 'extension' of a sentence to be its *truth value*.

<u>Why?</u>
Recall that the 'intension' determines for a given time/situation what the *extension* is.
Truth conditions determine for a given time/situation what the *truth value is.*

<u>Side-Note:</u>

The (crucial) idea that the extension of a sentence is its *truth value* is one of those 'weird ideas' that you just have to get used to…

For some further motivation behind this crucial concept:

a. Chierchia & McConnell-Ginet (2000: 62-65)
b. Our later discussion of the 'truth-functional sentence connectives'

(34) **The General Picture**

a. Intension of "X":

A kind of 'concept'/ 'formula' / 'function' which – for any given time/situation – determines what "X" 'picks out' at that time/situation.

b. Extension of "X"

The thing which, at a given time/situation "X" 'picks out'

c. Illustrative Paradigm:

i. Intension of "the president" = *whoever holds the office of the presidency*

ii. Extension of "the president" = Barack Obama

iii. Intension of "Barack smokes" = "Barack smokes" is T iff Barack smokes

iv. Extension of "Barack smokes" = TRUE

What's the point of all of this?

You can actually build a decent theory of how meanings can 'compose' to yield truth-conditions by paying attention to extensions (defined as above)!

3.2 Towards a Compositional Semantics Based on Extensions

(35) **Some New Notation** [[X]] = the extension of "X"
(function from X to its extension)

(36) **The Compositionality of Extensions**

a. We've broken down the 'meaning' of a S/NP into its *extension* and its *intension*.

b. Recall, however, that our semantic system must be such that the 'meaning' of a complex expression should be *derived from* the 'meaning' of its component parts.

c. *Thus, the extension of a complex expression should be derived from the extensions of its component parts!*

d. A Picture: [[Barack]] + [[smokes]] = [[Barack smokes]] = True

(37) **The Extension of a Predicate = Function**

a. In order to make the picture in (36d) work, the extension of the predicate "smokes" must be such that 'combining' it with the extension of "Barack" yields the extension of "Barack smokes".

b. *How can we model this?* Well, we know the following:

(i) [[Barack]] = Barack
(ii) [[Barack smokes]] = T(rue)

c. *Thus:* [[smokes]] + Barack = T

The extension of "smokes" 'combines' with Barack to yield T

d. *So, the extension of "smokes" is like a 'device' that takes Barack as input and yields T as output…*

e. ***So, the extension of "smokes" is like a FUNCTION!***

f. The Core Idea:

[[smokes]] = A function from entities to truth values, which yields T iff that entity smokes.

$$f : \{ x : x \text{ is an entity} \} \rightarrow \{ T, F \}$$
$$\text{for every } y \in \{ x : x \text{ is an entity} \},\ f(y) = T \textit{ iff } y \text{ smokes}$$

(38) **Interim Summary: Our Lexical Entries**

a. [[Barack]] = Barack

b. [[smokes]] = $f : \{ x : x \text{ is an entity} \} \rightarrow \{ T, F \}$
for every $y \in \{ x : x \text{ is an entity} \}$, $f(y) = T$ *iff* y smokes

We've almost got the picture in (36d) worked out… all we need is a rule for 'combining' the extensions of "Barack" and "smokes" to yield the extension of "Barack smokes"

(39) **The Rule of Function Application (Heim & Kratzer 1998: 44)**

If X is a branching node that has two daughters – Y and Z – and if [[Y]] is a function whose domain contains [[Z]], then [[X]] = [[Y]]([[Z]])

With the rule in (39), we now have a system that derives the extension of the sentence "Barack smokes" from (i) the extension of its component pieces, and (ii) the syntax of the sentence.

(40) **Computing the Extension of 'Barack smokes'**

a. <u>Syntactic Assumption (to be revised shortly):</u>
The structure of the sentence *Barack smokes* is as follows:

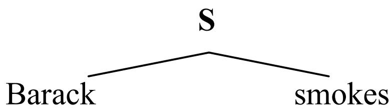

b. <u>Semantic Derivation:</u>

(i) The extension of "Barack smokes" = (by Syntactic Assumption)

(ii) The extension of ' 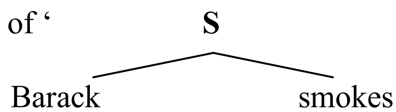 ' = (by notation)

(iii) [[S]] = (by F(unction) A(pplication))

(iv) [[smokes]] ([[Barack]]) = (by (38a))

(v) [[smokes]] (Barack) = (by (38b))

(vi) f (Barack) = (by (38b) and the facts of the world)

(vii) True

So, the system in (38) and (39) can derive the extension of a sentence (its truth value) from the extension of its component parts (given the facts of the world)…

So, we've obtained the system sketched in (36d)…

…Ok, but so what?...

3.3 From Extensions to Truth Conditions

<u>Question:</u> How does the system in (38) and (39) – which derives *extensions* – advance our goal of developing a formal system that derives *truth-conditions?*

(41) **Our Desired Semantic System**

A precise, formal system that for every sentence S of the language (*i.e.*, English), will derive a *correct* statement of the following form: *"S" is True iff X*

Believe it or not, our system for deriving extensions achieves what we want in (41)!!!

Side-Note: Some logical truths to keep in mind

a. *Transitivity of 'iff'* A *iff* B and B *iff* C **entails** A *iff* C

b. *Substituting 'equals' for 'equals'* If x = y, then x = z *iff* y = z

(42) **Deriving the Truth Conditions of a Sentence via Our Theory of Extensions**

To Prove: "Barack smokes" is T iff Barack smokes

a. "Barack smokes" is T *iff* (by Syntactic Assumptions)

b. " **S** " is T *iff* (by definition of our notation)

Barack smokes

c. [[**S**]] = T *iff* (by FA)

c. [[smokes]]([[Barack]]) = T *iff* (by (38))

d. f (Barack) = T *iff* (by definition of 'f' in (38b))

e. Barack smokes.

What just happened:

- In the preceding section, we first showed how a compositional extensional semantics can, *given the facts in the world*, compute the truth value of a sentence.
- *Conversely*, if we take as *hypothesis* that the truth value of a sentence is 'TRUE', our compositional extension semantics can 'work backwards', and compute *how the world must be constituted in order for the sentence to be true!*
- **Thus, we can use our extensional semantic function "[[ ]]" to compute the truth conditions of sentences!!!**

(43) **The Big Upshot**

An 'extensional semantics' – a formal system that maps complex structures to their extensions in the world – can provide us with a theory of how our brains recursively map sentences to their truth conditions.

(44) **Our Over-Arching Project, Redefined Again**

We wish to develop the *right* theory of the function "[[.]]", by examining the truth-conditions of particular sentences of the language. Such a theory will consist of:

a. Primitive statements for the lexical items of the language

b. Rules for deriving the value of "[[ ]]" for larger structures from (i) the value of "[[ ]]" for their component parts and (ii) the syntax of the larger structures

So our task is to adjust (a) and (b) until our theory predicts *exactly* the correct truth conditions for every sentence of English (or whatever language we're working on)!

(45) **A Review of Where We Are and How We Got Here**

a. We want to know how the 'meaning' of a sentence is computed from the 'meaning' of its parts.

b. We found that (for declarative sentences) the term 'meaning' can (largely) be recast as *the truth-conditions of a sentence.*

c. Thus, we wish to know how the *truth conditions* of a sentence can be derived from the 'meanings' of its component parts.

(i) We took a detour and introduced the notion of the *extension* of a given linguistic expression.

(ii) We noted (stipulated) that the extension of a sentence is its *truth-value*.

(iii) We then developed a system that derives the extensions of sentences (their truth values) from the extensions of their component expressions.

(iv) The following are two key features of this system:

- The extensions of some expressions are *functions*

- A rule for 'combining' the extensions of two sister expressions to yield the extension of their mother is FUNCTION APPLICATION

d. **We found that this 'extensional semantics' can be used to derive the truth-conditions of sentences from the extensions of their component parts…**

4. Some Follow-Up Technical Points

4.1 Deriving the Meaning of Non-Branching Nodes

Thus far, we've been assuming that the structure of "Barack smokes" is that in (46a). However, a more realistic picture of its syntax might be that in (46b).

(46) **The Structure of "Barack smokes"**

a. Simplified Picture (No Non-Branching Nodes)

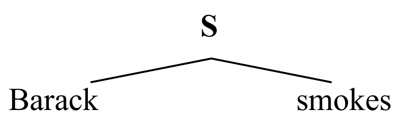

b. More Realistic Picture (With Non-Branching Nodes)

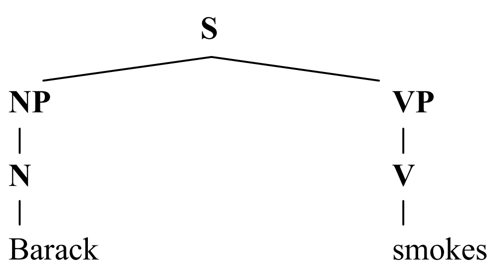

In order to formally compute the truth-conditions of structures like (46b), we will need to introduce some explicit rules for dealing with non-branching nodes like NP, N, VP and V.

(47) **The Rule for Non-Branching Nodes (Heim & Kratzer 1998: 44)**

If X is a non-branching node, and Y is its sole daughter, then [[X]] = [[Y]]

As we will see, it will also be useful to adopt a general rule for terminal nodes such as "Barack" and "smokes", rather than continually referring to their specific lexical entries (e.g. (38a,b))

(48) **The Rule for Terminal Nodes (Heim & Kratzer 1998: 43)**

If X is a terminal node, then [[X]] is specified in the lexicon

In our proofs, we can abbreviate the rule in (47) as "NN" and (48) as "TN"

So, let's get some practice by using these rules in a proof!

(49) **Starting Off the Proof**

a. “ 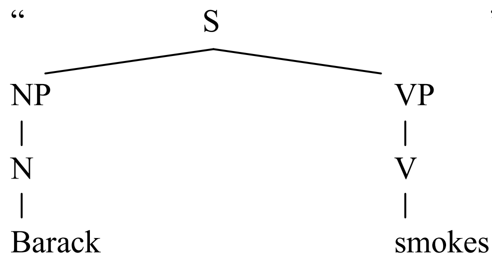

” is T *iff* (by notation)

b. [[S]] = T *iff* … uh oh!...

(50) **PROBLEM**

- We want to use our rule of ‘Function Application’ to compute [[S]]…
- But, we don’t know whether the rule can apply unless we know the values of [[NP]] and [[VP]].
- But, the values of [[NP]] and [[VP]] must now themselves be computed (rather than just easily looking them up in our lexicon)…

(51) **The Answer: ‘Subproofs’**

- A *subproof* is a proof inside a larger proof.
- The subproof proves/derives something that is then used in the larger proof.

On the following page, we can see how the proof above can be completed by using a ***subproof…***

(52) **Illustrative Derivation with Subproofs**

a. “ 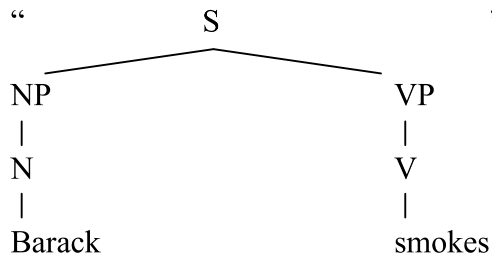

” is T *iff* (by notation)

b. [[S]] = T

c. **<u>Subproof</u>**

(i) [[NP]] = (by NN)

(ii) [[N]] = (by NN)

(iii) [[Barack]] = (by TN)

(iv) Barack

d. **<u>Subproof</u>**

(i) [[VP]] = (by NN)

(ii) [[V]] = (by NN)

(iii) [[smokes]] = (by TN)

(iv) f

e. [[S]] = T *iff* (by FA, **c**, **d**)

f. [[VP]]([[NP]]) = T *iff* (by **c**)

g. [[VP]](Barack) = T *iff* (by **d**)

h. f (Barack) = T *iff* (by def. of *f* in (38b))

i. Barack smokes.

Later on, when calculations start getting complicated, we will let you leave out ‘subproofs’…
For now, however, make sure to include any ‘subproofs’ your proofs rely upon, using the conventions above

4.2 On the Semantics of Lexical Items

In Section 3.2, we developed a theory of the semantic value of the intransitive verb 'smoked', whereby we identified its extension with a particular type of function (from entities to T-values).

- It's instructive to reflect on *how* we came to came to this conclusion:

(53) **Determining the 'Meanings' of a Lexical Item L**

a. Consider the truth-conditions of sentences in which L appears.

b. Consider the (already established) extensions of the other lexical items in these sentences.

c. Consider the rules of semantic composition at our disposal (e.g., FA, NN, TN).

d. Based on (a)-(c), develop an entry for L which would – in combination with the entries for the other words in the sentences (b) and our semantic rules (c) – correctly derive the truth conditions of the sentences it appears in (a).

Thus, to determine the 'meaning' of a lexical item within this general program for semantics, we don't (say) 'introspect' our 'concept' for the word…

(*e.g.* 'smoking' means lighting tobacco on fire and inhaling it, *etc. etc.*)

… rather, we look *only* to how the word contributes systematically to the truth conditions of a sentence.

- As we will see, this kind of an approach works especially well for the semantic analysis of *function words* like 'and', 'the', 'every', *etc.*

- For *content words* like 'smoke' and 'child', however, this approach admittedly ignores (or can ignore) much of their intuitive 'meaning'

  (the entry in (38b) doesn't say anything about what distinguishes 'smoking' from just 'inhaling smoke'.)

- *This rather 'sparse' treatment of the semantics of 'content words' may or may not be a critical weakness of the overall 'truth-conditional' approach to semantics.*
  (Jackendoff 1990; Lakoff 1990)

4.3 On Identifying 'Truth Values' as the Extensions of Sentences

Our semantic system has one core property that often first strikes people as 'odd' or 'confusing':

(54) **The Central (But Initially Confusing) Assumption of Our Semantic Theory**
Sentences of natural language have *extensions* ('denotations', 'reference'), and the *extension* of such sentences is a *truth value*.

We've given some general conceptual motivation to this claim, and have shown that it *does* contribute to our goal of having a formal system that derives truth-conditions….

…However, there is also something of an *<u>empirical argument</u>* *for these two assumptions.*

(55) **Empirical Argument for the Assumption in (54)**

a. <u>Fact:</u>
There are a number of natural language expressions – called the 'logical connectives' – that have the following key properties:

(i) *Syntactically, they take sentences as complements / specifiers / sisters*
(ii) *Semantically, they seem to take <u>truth values</u> as arguments.*

b. <u>Argument:</u>
If the 'meaning' of one lexical item must take as argument the 'meanings' of its complements/specifiers, then such 'logical connectives' *independently* show that **at some level, the 'meaning' of a sentence is its truth value.**

(56) **Example: '*It is not the case that*'**

a. <u>Syntax:</u> Takes sentences as complement.

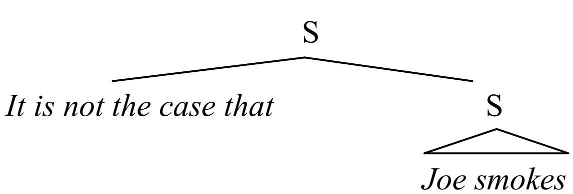

b. <u>Semantics:</u>
[[it's not the case that]] = $h: \{ T, F \} \rightarrow \{ T, F \}$
for all $y \in \{ T, F \}$, $h(y) = T$ *iff* $y = F$

The extension of 'it is not the case that' is a function which takes as argument a truth-value y and yields the value T iff y is F.

(57) **Question**
But, why should we assume the semantics in (56b)?

Answer: With the lexical entry in (56b), our system can derive the following truth-conditional statement:

"Its not the case that Joe smokes" is T *iff* Joe doesn't smoke.

(58) **Illustrative Derivation**

a. "It is not the case that Joe smokes" is T *iff* (by Syntactic Assumptions)

b. " 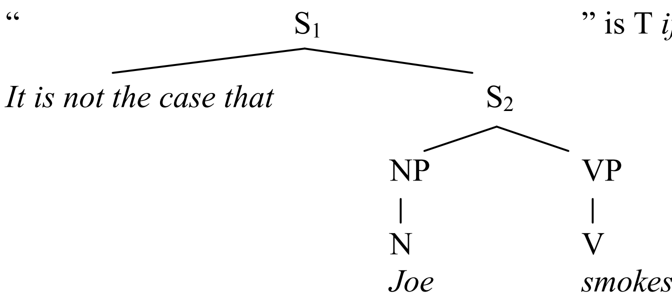 " is T *iff* (by notation)

c. [[S_1]] = T

d. **Subproof**

(i) [[NP]] = (by NN)

(ii) [[N]] = (by NN)

(iii) [[Joe]] = (by TN)

(iv) Joe

e. **Subproof**

(i) [[VP]] = (by NN)

(ii) [[V]] = (by NN)

(iii) [[smokes]] = (by TN)

(iv) f

f. **<u>Subproof</u>**

(i) [[S_2]] = (by FA, d, e)

(ii) [[VP]]([[NP]]) = (by d)

(iii) [[VP]](Joe) = (by e)

(iv) f(Joe)

g. [[S_1]] = T *iff* (by FA, NN, f)

h. [[it is not the case that]]([[S_2]]) = T *iff* (by TN)

i. h([[S_2]]) = T *iff* (by def. of *h* in (56b))

j. [[S_2]] = F *iff* (by f)

k. f(Joe) = F *iff* (by def. of *f* in (38b))

l. Joes doesn't smoke.

(59) **The Main Point**

a. The proof in (58) suggests that the lexical entry in (56b) is on the right track.

b. If the lexical entry in (56b) is correct, then the expression '*it is not the case that*' has as its extension a function that that takes **truth-values** as argument.

c. If the syntax in (56a) is correct, then the extension of '*it is not the case that*' must take as argument the extension of its sentential complement.

d. **Consequently, we must conclude that the extension of the sentential complement of '*it is not the case that*' is a <u>truth-value.</u>**

e. Thus, the extension of a sentence is a truth-value…

5. Where Do We Go From Here?

- Right now, we have a system that can derive truth-conditions for a *very limited* set of English sentences: *those headed by intransitive Vs and whose subjects are proper names.*

- We are going to quickly expand this system, so that it can interpret ever more complex (and therefore 'realistic') structures of English.

(60) **Some (Very Basic) Structures We Will Cover in the Next Few Months**

a. Transitive Verbs

b. Non-Verbal Predicates (Nouns and Verbs as Main Predicates)

c. NP Modification Structures

d. Relative Clauses (and Other Movement Structures)

e. Definite Descriptions

f. Pronouns

g. Quantifiers (and 'quantificational DPs')

h. Sentential complements

Once you've seen how the phenomena in (60) can receive treatment in a formal, truth-conditional semantic framework, you will have gained a relatively rich 'toolbox', which you should be able to apply to other phenomena in other languages…