The Foundations of Truth-Conditional Semantics[1]

1. The Big Questions

Semantics is the study of meaning. *But what does that mean?*

Well, consider the following, opening passage from 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…

(1) **The Over-Arching Miracle: The Productivity of Language (Semantics)**

The speaker of a language L can understand all the *countably infinite* sentences of L.

<u>How do we do this?</u>
Since brains/lifetimes are finite, this knowledge must be represented in our brains as some kind of **combinatoric system**, one that comprises:

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

b. A finite set of 'rules' for deriving the meanings of complex expressions from the meanings of the primitives.
(*Compositionality*)

What's especially interesting is that this knowledge we all have is *tacit*, *unconscious*…
…***therefore, if we want to find out what this 'combinatory system' is, we need to do SCIENCE!***

(2) **The Over-Arching Research Question**

What is the combinatoric system (algorithm) that our cognitive systems employ to derive/compute the meanings of complex expressions from the meanings of their component parts?

[1] These notes are based on material in Heim & Kratzer (1998: 1-12, 13-26), as well as Chierchia & McConnell-Ginet (2000: 1-33, 53-73, 99-104)

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

a. How does a human being acquire this combinatoric system? How much is already specified by the biology of the organism?

b. How does this combinatoric system vary across languages? Do languages differ in how they 'compute meanings', and if so, in what ways?

Fact: These questions don't receive nearly as much attention within semantics that they do in syntax and phonology (but this is improving…)

---

2. The Meaning of 'Meaning'

(4) **An Immediate Problem for Our Project**
What is a 'meaning'?

- We want to develop a 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 "Obama" and the '**meaning**' of "smokes" such that 'combining them together' gives us the '**meaning**' of "Obama smokes"?

Over-Arching Problem:

- The word "meaning" is a vague, pre-theoretic term from every-day discourse.
- Thus, it may not be an appropriate term for a precise, scientific study of human language. (*cf.* biology doesn't actually employ terms like 'alive' or 'dead' or 'life-form')

(5) **New, Preliminary Objective**

Let's replace our everyday, pre-theoretic concept of 'meaning' with something more precise, so that we might work towards a formal system that can manipulate 'meanings'…

(6) **Side-Benefit: Better Description of 'Meanings'**

A side-benefit of developing (and working within) a formalized semantics is that it forces a richer and more precise *documentation* of linguistic meanings.

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

- By employing a formal semantics, the researcher is forced to ask and answer deeper semantic questions about the language, resulting in greater depth of descriptive coverage…

- Without a formal semantics, certain deeper semantic questions about a given language *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'?...

(7) '**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?
Well, we know a lot of things!

(a) Social Appropriateness of the Statement:
What social contexts the statement is appropriate in.
"That's wonderful." vs. "That kicks ass!"

(b) 'Emotional Content' of the Statement:
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 Statement***
What information about the world the sentence 'conveys'.

Rightly or wrongly, (c) 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 develop our concept of the 'information' that is '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.

(8) **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.

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

- 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.

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

- Sentence S **presupposes** *p* = S is true *or false* 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.

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

- *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.

Those in the know will know that the above discussion is rather crude…*however it suffices to bring us to our main point below…*

<u>The Main Point:</u>

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 are derived from the meanings of its parts.

(ii) How the *presuppositions* of a sentence are derived from the meanings of its parts

(iii) How the *implicatures* of a sentence are derived (in part) from the meanings of its parts.

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

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 at times*…)

(12) **Our Goal (Restated)**

Develop a theory of the combinatory system that derives the *assertions* of a complex (declarative) sentence from the 'meanings' of its component expressions…

<u>Side-Note:</u> 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 (12), let's introduce a new bit of terminology.

(13) **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 (12) can again be restated to the following:

(14) **Our Goal (Restated Again)**

Develop a theory of the combinatory system that derives the *truth conditions* of a sentence from the 'meanings' of its component parts.

A Quick Review of How We Got Here:

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

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…

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 (7), (10) and (11)….

Our restated goal in (14) carries a certain presupposition regarding the nature of our 'language faculty', which it is worth pausing to reflect on:

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

Our goal in (14) 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 how human communication proceeds:

(16) **A Model of How Communication Operates**

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 this model of human communication, a central step in comprehension is the computation of the truth-conditions of the speaker's utterance:

- If listeners have 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 (16) 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***

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

a. The Object Language:
The language we are *describing* (not 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 (14):

(18) **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:*

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

a. In one sense, the *meaning* of the NP "the president of the United States" has just recently changed. It's gone from *meaning* George Bush (with all of his attendant connotations) to *meaning* Barack Obama.
('denotation', 'reference')

b. In another sense, the *meaning* of the NP has stayed the same. 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".

(20) **Extension vs. Intension**

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

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

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

(21) **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 at the sentence's 'conceptual/propositional content'.

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

(22) **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' (Section 4.3)

(23) **The General Picture**

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

b. <u>Extension of "X"</u>
The thing which, at a given time/situation "X" 'picks out'

c. <u>Illustrative Paradigm:</u>

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

ii. Extension of "the president" = Barack Obama

iii. Intension of "the president smokes" = "t.p.s" is T iff the president smokes

iv. Extension of "the president 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

(24) **Some New Notation**
[[X]] = the extension of "X"

(25) **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. <u>A Picture:</u>
[[Obama]] + [[smokes]] = [[Obama smokes]] = True

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

a. In order to make the picture in (25d) work, the extension of the predicate "smokes" must be such that 'combining' it with the extension of "Obama" yields the extension of "Obama smokes".

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

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

c. *Thus:* [[smokes]] + Obama = T
The extension of "smokes" 'combines' with Obama to yield T

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

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

f. <u>The Core Idea:</u>
[[smokes]] =
A function from entities to truth values, which yields T iff that entity smokes.

> f : { x : x is an entity } → { T, F }
> for every y ∈ { x : x is an entity }, f(y) = T *iff* y smokes

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

a. [[Obama]] = Obama

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

We've almost got the picture in (25d) worked out… all we need is a rule for 'combining' the extensions of "Obama" and "smokes" to yield the extension of "Obama smokes"

(28) **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 (28), we now have a system that derives the extension of the sentence "Obama smokes" from the extension of its component pieces! *Here's an illustrative derivation:*

(29) **Computing the Extension of 'Obama smokes'**

a. Syntactic Assumption (to be revised shortly):
The structure of the sentence *Obama smokes* is as follows:

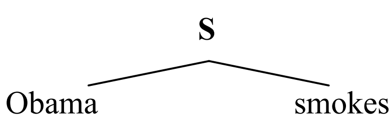

b. Semantic Derivation:

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

(ii) [[smokes]] ([[Obama]]) = (by (27a))

(iii) [[smokes]] (Obama) = (by (27b))

(iv) f (Obama) = (by (27b) and the facts of the world)

(v) True

So, the system in (27) and (28) 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 (25d)…

…Ok, but so what?...

3.3 From Extensions to Truth Conditions

Question: How does the system in (27) and (28) – which derives *extensions* – advance our goal of developing a formal system that derives *truth-conditions?*

(30) **Our Desired Semantic System**

A precise, formal system that for every sentence S of the language, 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 (30)!!!

(31) **Deriving the Truth Conditions of a Sentence via the Function "[[.]]"**

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

a. "Obama smokes" is T *iff* (by Syntax)

b. " **S** [Obama smokes] " is T *iff* (by definition of our notation)

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

c. [[smokes]]([[Obama]]) = T *iff* (by (27))

d. f (Obama) = T *iff* (by definition of 'f' in (27b))

e. Obama 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!!!**

(32) **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.

(33) **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 languages. 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 the value of "[[ ]]" for their component parts.

So our task is to adjust (a) and (b) until our theory predicts *exactly* the correct truth conditions for every sentence of the language!

(34) **A Summary 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 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

In our preceding discussion, we assumed that the structure of “Obama smokes” was that in (35a). However, a more realistic picture of its meaning might be that in (35b).

(35) **The Structure of “Obama smokes”**

a. Simplified Picture (No Non-Branching Nodes)

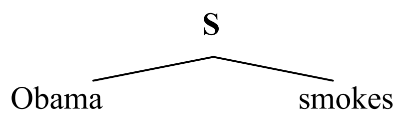

b. More Realistic Picture (With Non-Branching Nodes)

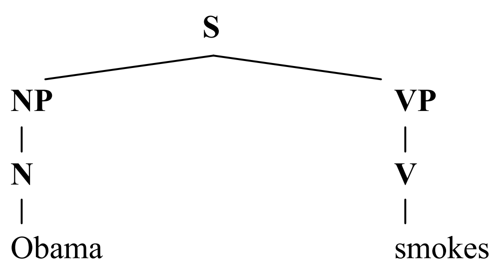

In order to formally compute the truth-conditions of structures like (35b), we will need to introduce some explicit rules for dealing with non-branching nodes like NP, N, VP and V.

(36) **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 “Obama” and “smokes”, rather than continually referring to their specific lexical entries (e.g. (27a,b))

(37) **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 (36) as “NN” and (37) as “TN”

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:

(38) **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. Based on (a) and (b), develop an entry for L which would – in combination with the entries for the other words in the sentences (b) – 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 much of their intuitive 'meaning'

  (the entry in (27b) 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 two properties that often first strike people as 'odd' or 'confusing':

(39) **Two Central (But Initially Confusing) Assumptions of Our Semantic Theory**

a. Sentences of natural language have *extensions* ('denotations', 'reference'), and the *extension* of such sentences is a *truth value*.

b. Our 'semantic interpretation function' "[[.]]" maps the structures of natural language to their *extensions* (not, say, to the conceptual content (intensions) that determine those extensions).

We've given some general conceptual motivation to these claims, by showing that they *do* contribute to our goal of having a formal system that derives truth-conditions….

However, there is also something of an empirical argument for these two assumptions.

(40) **Empirical Argument for the Assumptions in (39)**

There are a variety of natural language expressions (*e.g.* the 'logical connectives') which:

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

If the 'meaning' of a lexical item must take as argument the 'meanings' of its complements/specifiers, such 'logical connectives' *independently* show that:

- At some level, the 'meaning' of a sentence is its truth value.

This (still initially surprising) result makes the most sense within the system we develop above.

(41) **A Logical Connective ('Truth Functional Operator') in Natural Language**

[[it's not the case that]] = f: { T, F } → { T, F }
for all y ∈ { T, F }, f(y) = T *iff* y = F

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

With the lexical entry in (41), our system can derive the following truth-conditional statement:

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

4.4 Sets and Their Characteristic Functions

In the system we just developed, the extension of the verb "smokes" is a curious type of *function*:

a function from entities to truth-values

(42) **The Extension of *Smokes* as a Function from Entities to Truth-Values**

a. [[smokes]] = f : { x : x is an entity } → { T, F }
for every y ∈ { x : x is an entity }, f(y) = T *iff* y smokes

b. [[smokes]] = { <Obama,T>, <Joe,F>, <Seth,F>, <Lou,T>…}

By extension, we can treat *all* intransitive verbs (*runs, dances, laughs*) as functions of this type.

(43) **Generalization Regarding Intransitive Verbs, Part 1**
The extension of an intransitive verb can be modeled as a function from entities to truth-values.

However, it is not unusual to also see, in some systems, the extension of an intransitive verb characterized as a *set of entities.*

(44) **The Extension of *Smokes* as a Set of Entities**
a. [[smokes]] = { x : x smokes }
b. [[smokes]] = { Obama, Lou … }

(44) **Generalization Regarding Intransitive Verbs, Part 2**
The extension of an intransitive verb can be modeled as a set of entities.

There is a very good reason why (43) and (44) both hold:
*In a certain sense, sets of entities and functions from entities to truth-values **are the same thing!***

- For every set of entities, there is a unique function from entities to T-values
- For every function from entities to T-values, there is a unique set of entities

(45) **Transforming a Set of Entities into a Function from Entities to T-Values**
Let *A* be a set of entities. Then 'the characteristic function of *A*' is that function f_A such that for any entity x, f_A(x) = True *iff* x ∈ *A*.

(46) **Transforming a Function from Entities to T-Values into a Set of Entities**
Let *f* be a function from entities to truth-values. Then 'the set characterized by *f*' is that set A_f such that for any entitiy x, x ∈ A_f iff *f*(x) = True.

Given this special correspondence between sets and their characteristic functions, we will often times equivocate between them – sometimes talking about the extension of a V as if it were a set, and sometimes talking about it as if it were a function…

This will be more relevant when we start talking about quantifiers, but it's good to see it now.

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.

(47) **Some (Very Basic) Structures We Will Cover in the Next Few Weeks**

a. Transitive Verbs
b. Non-Verbal Predicates
c. Adjectives
d. Relative Clauses (and Other Movement Structures)
e. Definite Descriptions
f. Pronouns
g. Quantifiers

- Once we've built up enough machinery to correctly interpret each of these structures, we'll be able to branch out into some subject matter that is a bit more 'empirically interesting':
  - The semantics of indefinites
  - The semantics of plural NPs