Orin Percus, University of Nantes / LLING EA3827
Attitude reports at the syntax-semantics interface
EALing, ENS, 2007

A starting point

Where we are going (today).

The purpose of this first class is to sketch a theory of how we associate syntactic structures with information (an adaptation of Heim and Kratzer 1998).

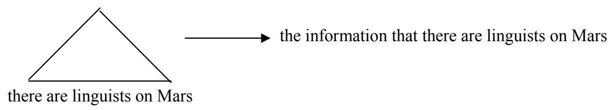

Syntactic structures? Not sentences? An argument:

(1) Luigi left the museum with the painting by Leonardo.

This sentence has two different syntactic structures -- roughly (!) the following:

(2) S1 : … S2 : …

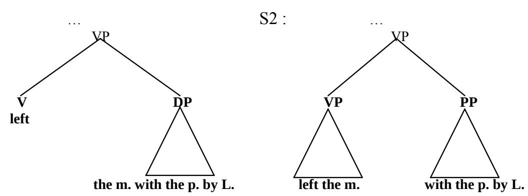

It also has two different meanings:

(3) M1: the museum is the one with the painting M2: Luigi has the painting

We can show that the fact that the sentence has meaning M1 is related to the fact that it has structure S1, and the fact that it has meaning M2 is related to the fact that it has structure S2. So syntactic structure is at least one factor that determines meaning.

Recommended exercise: Show this (specifying all the reasoning). (You might consider conjunctions like *Luigi left the museum and headed towards the gallery with the painting by L.*)

Given that the same words, combining in different ways, yield different interpretations, it makes sense to ask simply how merely the *arrangement* of words in a structure affects the interpretation that the structure has.

The theory I will present specifically addresses this issue. At the same time, I will be introducing a way of describing information, which will affect the way the theory is formulated.

We will get there gradually.

Possible worlds.

<u>Their use to describe the partial information that we have about our world</u>

A way of viewing the fact that there are many ways in which the world might be for all we know: There are many possible worlds among our "candidates for the actual world."

When we learn something new, we modify our set of candidates -- typically narrowing them down, if our new information is compatible with what we already believed at that point.

E.g. You didn't know that my handout would contain just these words.
So, before seeing it, your candidate set included worlds where my handout contained just these words as well as words where my handout contained different words.
But now that you have seen it, you have excluded from your candidate set worlds where my handout had different words.

> W = { w_1, w_2, w_3, w_4, ...} the set of possible worlds
> (The actual world is an element of W. We don't know which.)
>
> In world w_{123}, there exists a handout with these words, but not in w_{248}.
> Before, w_{123} and w_{248} both figured among your "candidates for the actual world."
> Now, w_{248} no longer does.

<u>Their use to describe the way sentences get used to convey information</u>

(4) Nobody is allergic to cheese.

One way of looking at the information content that we associate with (the syntactic structure of) this sentence: it says that a world is somewhere in the set of worlds in which nobody is allergic to cheese.
If what we are talking about is the actual world, and we don't know which world that is, then when we hear the sentence we can at least narrow it down to some world in the set of worlds in which nobody is etc.

For now, we will imagine the following.

On the basis of the syntactic structure of *Nobody is allergic to cheese*, we determine a function like:

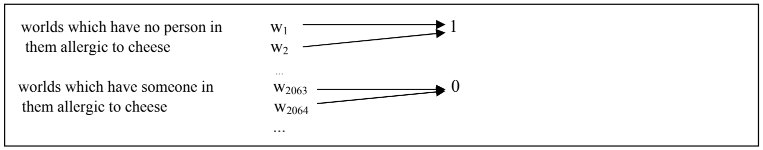

The syntactic structure gets used to convey that this function yields 1 for the actual world.

> Notation: f = λ**u: ϕ. ψ** for (i) dom(f) = {u: ϕ}
> (ii) For all u in dom(f), if ψ then f(u) = 1 and otherwise f(u) = 0
>
> So I drew here:
> λu: u is a possible world. { x: x is a person in u } $\cap$ {x: x is allergic to cheese in u } = $\varnothing$
>
> More abbreviated : **λw. no person in w is allergic to cheese in w**

An initial toy theory.

We arrive at such functions piecewise.

We can see every speaker, at a moment in time, as following a procedure that begins by associating objects with the words in a syntactic structure. Then, on the basis of these values, it progressively calculates values for the larger constituents, until it reaches the top.

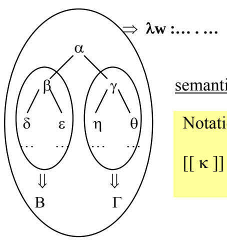

semantic values

Notation :

[[κ]] = z The value that the procedure in question yields for constituent κ is z.

Initial claim: the way it does this is by "functional application." (Always! For everyone! This is UG!)

That is, to arrive at a semantic value for a complex constituent:
The semantic value of one of the subconstituents must be a function.
The semantic value of the other must be an object in the domain of the function.
And we obtain the semantic value of the complex constituent applying the function to the object in its domain.

Functional application (FA)

[[[β γ]]] = [[β]] ([[γ]]) or [[γ]] ([[β]]).

Here is how this approach could potentially lead us to analyze sentences like *Nobody hibernates* or *Orin hibernates*, assuming ultrasimplified syntactic structures.

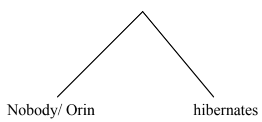

Note: In what follows I will assume that it makes sense to say that the same object may be present in different worlds, and may have different properties in different worlds. Given this, one should now read "x hibernates in w" below as "object x has the property of hibernating in world w (but might not in a different world)" and one should read "there is no person y in w such that ..." as "there is no object y such that, in world w, y has the property of personhood, and such that ..." If you don't like this, don't worry: it's mostly to simplify the exposition, it won't play much of a role, and anyway I am more interested here in the way the mechanics works than in the precise function that gets produced in the end.

Notation:
Below I have used the "λ" notation in a new way. In this way, after the ".", we find a description of the function's value. For example, using this notation, I could write "λu: u is a girl. u's mother" to describe a function whose domain is the set of girls, and which, for every u in its domain, yields u's mother as a value.

f = **λu: ϕ. α** for (i) dom(f) = {u: ϕ}
(ii) For all u in dom(f), f(u) = <u>α</u>

When below I write "λx. λw. x hibernates in w," I am making use of the "λ" in two different ways. First, the new way: I am describing a function that, for every object x in its domain , yields as a value the function <u>λw. x hibernates in w</u>. Then, the old way: in writing "λw. x hibernates in w," I am describing a function that, for every world w in its domain, yields 1 if the individual in question hibernates in w and 0 otherwise.

For someone like me:

[[hibernates]] = λx. λw. x hibernates in w

(I will use **λx. ...** , **λy. ...** , **λz. ...** to abbreviate **λu: u is an "individual". ...**
i.e. something that is *not* a function.)

[[Orin]] = o

(Where o is a certain object who is present in some worlds.)

[[nobody]] = λP. λw. There is no person y in w such that P(y)(w) = 1.

(Where **λP. ...** abbreviates: **λu: u is a function like [[hibernates]]. ... ,** i.e. u is the kind of function which we could notate as λx. λw. ...)

Then the semantic value of the full structure is obtained by FA:

[[Orin hibernates]]
= [[hibernates]]([[Orin]])
= [[hibernates]](o)
= λx. λw. x hibernates in w (o)
= <u>λw. o hibernates in w</u>

[[nobody hibernates]]
= [[nobody]]([[hibernates]])
= λP. λw. There is no person y in w such that P(y)(w) = 1
([[hibernates]])
= λw. There is no person y in w such that
[[hibernates]](y)(w) = 1
= λw. There is no person y in w such that
λx. λv[1]. x hibernates in v (y)(w) = 1
= λw. There is no person y in w such that
λv. y hibernates in v (w) = 1
= <u>λw. There is no person y in w such that y hibernates in w</u>

[1] I have changed the letter here to make the calculation more readable.

Comment: FA is only one of many rules of semantic composition that one could in principle imagine. For example, here is an alternative that one could consider.

FA "passing up world arguments"

$[[\ [\beta\ \ \gamma]\ \ ]] = \lambda w.\ [[\beta]](w)\ (\ [[\gamma]](w))$ [2] or $\lambda w.\ [[\gamma]](w)\ (\ [[\beta]](w))$.

Another (more common) way of writing the same rule:

$[[\ \ [\beta\ \gamma]\ ]]^{w} = [[\beta]]^{w}\ ([[\ \gamma]]^{w})$ or $[[\gamma]]^{w}\ ([[\ \beta\ ]]^{w})$

The way in which this approach could potentially lead us to analyze the same sentences, in order to obtain the same results:

$[[\ \text{Orin}\ ]] = \lambda w.\ o$

$[[\ \text{hibernates}\ ]] = \lambda \underline{w}.\ \lambda \underline{x}.\ x$ hibernates in w

(Note the change in the order of the arguments: the world argument, which gets "passed up," becomes the last one.)

$[[\ \text{nobody}\ ]] = \lambda w.\ \lambda \underline{F}$: <u>F is a function from individuals to {0,1}</u>. There is no person y in w such that <u>F(y) = 1</u>.

[[Orin hibernates]]
= λw. [[hibernates]](w) ([[Orin]](w))
= λw. [[hibernates]](w) (<u>λv. o</u> (w))
= λw. [[hibernates]](w) (o)
= λw. <u>λv. λx. x hibernates in v</u> (w)(o)
= λw. <u>λx. x hibernates in w</u> (o)
= λw. o hibernates in w[3]

[[nobody hibernates]]
= λw. [[nobody]](w) ([[hibernates]](w))
= λw. <u>λv. λF. There is no person y in v such that F(y) = 1</u> (w) ([[hibernates]](w))
= λw. <u>λF. There is no person y in w such that F(y) = 1</u> ([[hibernates]](w))
= λw. <u>There is no person y in w such that [[hibernates]](w)(y) = 1</u>
= λw. There is no person y in w such that <u>λv. λx. x hibernates in v</u> (w)(y) = 1
= λw. There is no person y in w such that <u>λx. x hibernates in w</u> (y) = 1
= λw. There is no person y in w such that y hibernates in w

$[[\ \text{Orin}\ ]]^{w} = o$

$[[\ \text{hibernates}\ ]]^{w} = \lambda x.\ x$ hibernates in w

$[[\ \text{nobody}\ ]]^{w} = \lambda F$. There is no person y in w such that F(y) = 1.

$[[\text{Orin hibernates}]]^{w}$
$= [[\ \text{hibernates}]]^{w}\ (\ [[\text{Orin}]]^{w}\ )$
$= [[\ \text{hibernates}]]^{w}\ (o)$
= <u>λx. x hibernates in w</u> (o)
= 1 if o hibernates in w, and 0 otherwise

$[[\text{nobody hibernates}]]^{w}$
$= [[\text{nobody}]]^{w}\ (\ [[\text{hibernates}]]^{w}\ )$
= <u>λF. There is no person y in w such that F(y) = 1</u> ($[[\text{hibernates}]]^{w}$)
= 1 if there is no person y in w such that $[[\text{hibernates}]]^{w}(y) = 1$, and 0 otherwise
= 1 if there is no person y in w such that <u>λx. x hibernates in w</u> (y) = 1, and 0 otherwise
= 1 if there is no person y in w such that y hibernates in w, and 0 otherwise

---

[2] To be more precise, this expression should read: λw:<u>w∈dom([[β]]) and w∈dom([[γ]]) and [[γ]](w) ∈dom([[β]](w))</u>…
[3] Note that in this last step I have switched to the original use of the "λ" notation, on which what we find after the "." is a condition for the value 1.

Revision 1: complications in the interpretive procedure (variable binding).

First step: The procedure systematically gives us *functions from assignments* (cf. the alternative we just considered, where we had functions from possible worlds), and one way of computing the semantic values of complex constituents is by functional application *passing up assignments*.

Sample semantic values for words:

[[Orin]] = **λg.** o
[[hibernates]] = **λg**. λx. λw. x hibernates in w
[[nobody]] = **λg.** λP. λw. There is no person y in w such that P(y)(w) = 1.
[[hates]] = **λg**. λx. λy. λw. y hates x in w

$[[\ \text{Orin}\ ]]^{\mathbf{g}} = o$
$[[\ \text{hibernates}\ ]]^{\mathbf{g}} = \lambda x.\ \lambda w.\ x$ hibernates in w
$[[\ \text{nobody}\ ]]^{\mathbf{g}} = \lambda P.\ \lambda w.$ There is no person y in w such that P(y)(w) = 1.
$[[\ \text{hates}\ ]]^{\mathbf{g}} = \lambda x.\ \lambda y.\ \lambda w.$ y hates x in w

FA "passing up assignments"

[[[β γ]]] = λg: g is in dom([[β]]) and dom([[γ]]) and [[γ]](g) is in dom([[β]](g))[4]. [[β]](g) ([[γ]](g))
or λg: g is in dom([[β]]) and dom([[γ]]) and [[β]](g) is in dom([[γ]](g)). [[γ]](g) ([[β]](g))

$[[\ \ [\beta\ \gamma]\ ]]^{g} = [[\beta]]^{g}\ ([[\ \gamma]]^{g})$ or $[[\gamma]]^{g}\ ([[\ \beta]]^{g})$

Before we said that a structure S gets used to convey that [[S]] yields 1 for the actual world.

Now we need to make a minor revision, since semantic values for sentences are no longer simply functions from worlds to {0,1}.

Let's say: a structure S gets used to convey that $[[S]]^{\varnothing}$ yields 1 for the actual world, where ∅ is the null assignment.[5, 6]

.

Second step: There are indices in the syntax. We find them inside traces of movement, inside pronouns and elsewhere. Following Heim and Kratzer 1998, we assume that, when an item moves, it leaves just below its new position another instance of the same index that its trace includes.

(5) Nobody hates himself.

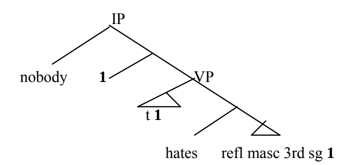

---

[4] I wrote in here the domain condition that I included in a footnote on the previous page.
[5] I am now going to start using the more readable notation with superscripts more often. This is, of course, just another way of saying that [[S]](∅) yields 1 for the actual world.
[6] If we want to maintain an approach of this kind, we will have to say, contrary to most research, that the semantic values of sentences are always constant functions from assignments -- that is, there are no free variables. I will assume that here.

Third step : The interpretive procedure works in such a way that indices like those on traces and pronouns function like variables ...

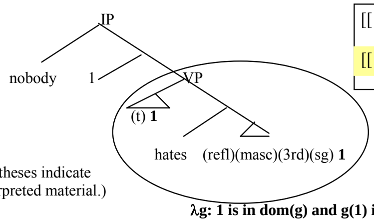

$[[1]] = \lambda g$: 1 is in dom(g). g(1)

$[[1]]^g = g(1)$

(Parentheses indicate uninterpreted material.)

λg: 1 is in dom(g) and g(1) is an individual. λw. g(1) hates g(1) in w

$[[VP]]^g = \lambda w$. g(1) hates g(1) in w

... and there is another (optional) composition rule that has the effect that indices like those below moved items function like abstractors (λ-operators). (I will call these indices <u>binders</u>.)

Abstraction (Abs).

Given a constituent [i β], where i is an index,

[[i β]] = λg: for some u, g[i->u] is in dom([[β]]). λu: g[i->u] is in dom([[β]]). [[β]](g[i->u]).

(Where g[i->u] stands for an assignment that maps i to u and is otherwise just like g.)

$[[\ i\ \beta\ ]]^g = \lambda u$: $[[\beta]]^{g[i->u]}$ is defined. $[[\beta]]^{g[i->u]}$.

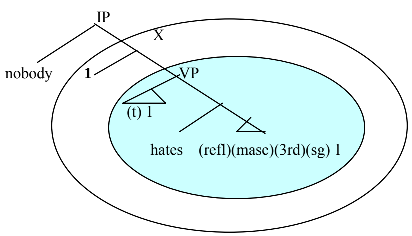

λg. λu. u is an individual. λw. u hates u in w
(a.k.a. λg. λx. λw. x hates x in w)

$[[X]]^g = \lambda x. \lambda w$. x hates x in w

Revision 2: an information state is more than a set of candidates for the actual world.[7]

The absent-minded gentleman[8]. John has gotten into the unfortunate habit of locking himself out of his house at night. When this happens, he climbs into the trunk of his car to sleep. He sleeps well there: there's no light at all, it's pitch black. But since he needs to be at work at 8, he sets three alarms on his cell phone to progressively wake him up, one for 5:30, one for 6:00, one for 6:30. The pattern is always the same. He is always drowsy after the first two alarms, and goes back to sleep – he reaches over with his hand and turns them off – but then the third alarm wakes him up for good. Since he is drowsy, typically he doesn't remember immediately whether an alarm has rung before or not, and it only dawns on him gradually. When the second alarm starts ringing, for example, he asks himself whether this is the first, second or third time. It is only just before he turns it off that he remembers, yes, the alarm rang before. Anyway, this has developed into a real habit, and John, when he climbs into the car to sleep, can count on the fact that this will be the series of events.

John's candidates for the actual world when he wakes up after the second alarm:

Before he realizes: worlds in which the alarm rings three times, he turns the alarm off the first two times and falls asleep again, and he wakes up for good the third time. At some point before he falls asleep the second time, he realizes "Oh, the alarm rang once before." (After all, he counts on this happening.)

After he realizes: the same!

What happened, intuitively: After he realized, he located himself temporally. He knew that "now" was a moment after 6:00.

Let's imagine that there is a time line such that, in all possible worlds, all the events occur at or between two points on that time line. Then we could give a more complete characterization of information states as a set of world-time *pairs*. A world-time pair is a member of a person's "set of candidate pairs" at a given moment when the world is one of his candidates for the actual world and the time is one at which he potentially locates himself in that world at that moment.

John's candidate world-time *pairs* when he wakes up after the second alarm:

Before he realizes: <w_1, the time at which, in w_1, the clocks in France read "5:31">
<w_1, the time at which, in w_1, the clocks in France read "6:01">
...

After he realizes[9]: ~~<w_1, the time at which, in w_1, the clocks in France read "5:31">~~
<w_1, the time at which, in w_1, the clocks in France read "6:01">
...

[7] This point is due to Lewis 1979. (David Lewis 1979, "Attitudes de dicto and de se," *The Philosophical Review*.)
[8] Based on Hass-Spohn's (1995) embellishment of a scenario from Stalnaker 1981. (Ulrike Hass-Spohn 1995, *Hidden Indexicality and Subjective Meaning*, web-accessible translation of 1994 Tübingen PhD dissertation. Robert Stalnaker 1981, "Indexical Belief," *Synthese*.)
[9] This isn't quite accurate since there is a time lag between his not realizing and his realizing, but you get the idea.

The (relatively rational) paranoiac.[10] Bill thinks that there is there is a cruel neurological experiment in progress. The effect of this experiment is that there are two people whose brains are giving them the same signals, who have exactly the same sensory experiences, memories, and so forth. One of them is actually living his normal life. But his brain is being monitored carefully, and a malicious neurosurgeon who is receiving and analyzing the signals is manipulating the brain of an operating table patient in such a way that the poor patient is going through exactly the same thing. Of course, the fact that there is such an experiment is bad enough, but to make matters worse Bill thinks that he is one of those two people. He just, naturally, doesn't know which one he is.

The (less rational) paranoiac. Bill' is just like Bill but he is convinced that he is not the guy on the operating room table. What *he* is seeing is "real."

What is the difference between Bill and Bill' ?

Do they have different sets of candidates for the actual world? No.

Do they have different sets of candidate world-time pairs? No.

The difference has to do with which individuals in their candidate worlds they consider as potential candidates for *themselves*. For Bill, the guy on the operating table is a candidate for himself. For Bill', he isn't.

We could give a more complete characterization of information states as a set of world-time-individual *triples*. A triple is a member of a person's "set of candidate triples" when the world is one of his candidates for the actual world, the time is one at which he potentially locates himself in that world at that moment, and the individual is someone in that world who could for all he knows be him.

Bill's candidate triples:

<w_1, t_1, the person in w_1 who is "actually" seeing … >
<w_1, t_1, the person in w_1 who is on the operating room table>
...

Bill' 's candidate triples:

<w_1, t_1, the person in w_1 who is "actually" seeing … >
~~<w_1, t_1, the person in w_1 who is on the operating room table>~~
...

Terminology: World-time-individual triples are sometimes called <u>contexts</u>.

[10] See the Lewis paper cited for more classic and equally entertaining examples that are meant to demonstrate the same thing.

Revision 3: the semantic values of sentences are more complicated.

We started out by saying that the semantic values of sentences* are functions from worlds to {0,1} and that we use them to convey that the function yields 1 for the actual world.

This fits together with the following approach to truth judgments: a speaker judges a sentence that he has uttered to be true when its semantic value yields 1 for each of his candidates for the actual world. (And false when its semantic value yields 0 for each of his candidates for the actual world.)

Things can't be that simple, though.

(6) It is after 6:00 pm.

Before he realizes, John would not be inclined to find this sentence, as uttered by himself, to be true. After, he would.

(7) I am not on an operating room table.

Bill would not be inclined to find this sentence to be true. Bill' would.

A better guess: we use sentences to convey something not merely about the actual world, but rather about triples of the actual world, the time of utterance, and the speaker. Whether a speaker judges a sentence to be true or not depends not merely on his candidates for the actual world, but rather on his candidate triples.

A first attempt:

The syntactic structures of sentences give us functions that need a world, a time, and an individual in order to arrive at 0 or 1.

(6') [[It is after 6:00 pm]]^g
= λx. λt. λw. in the place where x is located at t in w, the time indicated on the clocks is after 6:00pm

(7') [[I am not on an operating room table]]^g
= λx. λt. λw. x is not on an operating room table at t in w

A speaker judges true a sentence that he has uttered when for all w,t,x such that <w,t,x> is a candidate triple of his, [[S]]$^{\varnothing}$(x)(t)(w) = 1, where S is the syntactic structure that he "has in mind for it".

(...false...when for all w,t,x... [[S]]$^{\varnothing}$(x)(t)(w) = 0)

A contrast to consider:

(8) a. I am happy.
    b. I am happy now.

(9) a. It is never the case that I am happy.
    b. ?? It is never the case that I am happy now.

Since (8a) and (8b) seem to convey exactly the same thing -- we would judge them true in exactly the same situations -- this approach would lead us to the conclusion that they have the same semantic value (namely (8')):

(8') $[[\text{ (8a) }]]^{g} = [[\text{ (8b) }]]^{g} = \lambda x.\ \lambda t.\ \lambda w.$ x is happy at t in w

But then we would expect them to make the same contribution to sentences in which they are embedded (assuming crucially that embedded clauses that are pronounced like (8a) and (8b) have the same syntactic structures as unembedded clauses pronounced like (8a) and (8b)). And evidently they don't.

Many conclude from this that another revision is needed.

Revision 3': semantic values are more complicated still.

Semantic values are even more complicated than we have been imagining.

On the one hand, words that previously took only world arguments now frequently take time arguments as well -- so, if we stopped here, the semantic values of sentences would be objects that, once we provide an assignment, need times and worlds to produce 1 or 0.

On the other hand, now semantic values systematically take *context* (i.e. triple) arguments in addition to all the others we have posited.

$[[\text{ 1 }]] = \lambda g$: 1 is in dom(g) . $\boldsymbol{\lambda c.}$ g(1)	$[[\text{ 1 }]]^{g,c} = g(1)$
$[[\text{ hibernates }]] = \lambda g.\ \boldsymbol{\lambda c.}\ \lambda x.\ \lambda t.\ \lambda w.$ x hibernates at t in w	$[[\text{ hibernates }]]^{g,c} = \lambda x.\ \lambda t.\ \lambda w.$ x hibernates at t in w
$[[\text{ nobody }]] = \lambda g.\ \boldsymbol{\lambda c.}\ \lambda P.\ \lambda t.\ \lambda w.$ There is no person y at t in w such that P(y)(t)(w) = 1.	$[[\text{ nobody }]]^{g,c} = \lambda P.\ \lambda t.\ \lambda w.$ There is no person y at t in w such that P(y)(t)(w) = 1.
$[[\text{ hates }]] = \lambda g.\ \boldsymbol{\lambda c}.\ \lambda x.\ \lambda y.\ \lambda t.\ \lambda w.$ y hates x at t in w	$[[\text{ hates }]]^{g,c} = \lambda x.\ \lambda y.\ \lambda t.\ \lambda w.$ y hates x in w

The semantic composition rules are more complicated than we have been imagining only in that they systematically pass the context argument up -- along with the assignment argument in the case of functional application.

What we arrive at for full sentences are objects that are more complicated than we have been imagining in one sense, and less in another sense. More complicated in that they want us to provide a *context* as well as an assignment. Less complicated in that, when we do this, they give us something that characterizes world-time pairs, and not triples: something that gives us 0 or 1 when we provide a time and a world.

FA "passing up assignments and contexts"
$[[\ [\beta\ \gamma]\ ]]^{g,c} = [[\beta]]^{g,c}([[\ \gamma]]^{g,c})$ or $[[\gamma]]^{g,c}([[\ \beta]]^{g,c})$

Abs "passing up contexts":
$[[\ i\ \beta\ ]]^{g,c} = \lambda u$: $[[\beta]]^{g[i \to u],c}$ is defined. $[[\beta]]^{g[i \to u],c}$.

$[[IP]]^{g,c} = \lambda t.\ \lambda w.$ There is no person y at t in w such that y hibernates at t in w

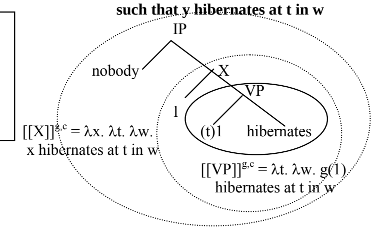

The role of the context argument is specifically to determine the value of terms like *I* and *now* (<u>indexicals</u>) The result that sentences with these terms have semantic values that wind up characterizing different world-time pairs when they are given different context arguments.

$[[\text{ I }]]^{g,c} = c^3$

$[[\text{ now }]]^{g,c} = c^2$

Notation: I will write "c^3" to mean the third member of the triple c. Writing "$[[\text{ I }]]^{g,c} = c^3$" is like writing "$[[\text{ I }]]^{g,\langle w,t,x\rangle} = x$," where $c = \langle w,t,x\rangle$.

These resources will be sufficient to get us out of our problem: the semantic values of *I am happy* and *I am happy now* will now be different.

$[[\text{ I am happy }]]^{g,c} = \lambda t.\ \lambda w.\ c^3$ is happy at t in w
$[[\text{ I am happy now }]]^{g,c} = \lambda t.\ \lambda w.\ c^3$ is happy at c^2 in w

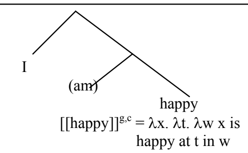

(Don't worry about the structure for *I am happy now*.)

On this approach, however, the role of semantic values will be more complicated when it comes to characterizing the use of a sentence. On this approach, we use a sentence with structure S to say that we get 1 when we provide [[S]] with:

(i) a context argument consisting of the actual world, the time of utterance, and the speaker;
(ii) an assignment argument that is the null assignment;
(iii) a time argument that is the time of utterance;
(iv) a world argument that is the actual world.

And when it comes to characterizing truth judgments, we would say that a speaker judges true a sentence that he has uttered when for all c such that c is a candidate triple of his, $[[S]]^{\varnothing,c}(c^2)(c^1) = 1$, where S is the syntactic structure that he "has in mind for it".

(8) a. I am happy.
b. I am happy now.

(9) a. It is never the case that I am happy.
b. ?? It is never the case that I am happy now.

Let's pretend that *it is never the case that* is a single word, which combines directly with (8a) or (8b). If its semantic value is something like

(10) $[[\text{...never...}]]^{g,c} = \lambda p.\ \lambda t.\ \lambda w.$ For no time t' is it the case that p(t')(w) = 1.

Then

(11) $[[(9a)]]^{g,c} = [[\text{... never ...}]]^{g,c} ([[\text{I am happy}]]^{g,c})$ (FA)
= <u>λp. λt. λw. For no time t' is it the case that p(t')(w) = 1</u> ($[[\text{I am happy}]]^{g,c}$)
= λt. λw. For no time t' is it the case that $[[\text{I am happy}]]^{g,c}$(t')(w) = 1
= λt. λw. For no time t' is it the case that <u>λs. λv. c^3 is happy at s in v</u> (t')(w) = 1
= λt. λw. For no time t' is it the case that <u>λv. c^3 is happy at t' in v</u> (w) = 1
= λt. λw. For no time t' is it the case that c^3 is happy at t' in w

(12) $[[(9b)]]^{g,c} = [[\text{... never ...}]]^{g,c} ([[\text{I am happy now}]]^{g,c})$ (FA)
= <u>λp. λt. λw. For no time t' is it the case that p(t')(w) = 1</u> ($[[\text{I am happy now}]]^{g,c}$)
= λt. λw. For no time t' is it the case that $[[\text{I am happy now}]]^{g,c}$(t')(w) = 1
= λt. λw. For no time t' is it the case that <u>λs. λv. c^3 is happy at $\mathbf{c^2}$ in v</u> (t')(w) = 1
= λt. λw. For no time t' is it the case that <u>λv. c^3 is happy at $\mathbf{c^2}$ in v</u> (w) = 1
= λt. λw. For no time t' is it the case that c^3 is happy at $\mathbf{c^2}$ in w

Some exercises to try and questions to think about, related to this final approach.

Q1. What would be an appropriate semantic value for the word *you* ?

Q2. We proposed what it means for a speaker to judge true a sentence that he himself has uttered. What does it mean for a speaker to judge true a sentence that another person has uttered?

Q3. In a certain sense, the sentence *I am happy*, uttered by me, and the sentence *You are happy*, uttered by you to me, are synonymous. How can the relevant notion of synonymy be explicated given the notions that this final approach makes use of?

Q4 (to think about). Throughout, I have been suggesting that there is a direct relation between our truth judgment for a sentence, and a single semantic value that we associate with that sentence. This could be questioned. Take a situation in which I know that John is one of the people in the room, but I don't know which, and all the people in the room are blond. I would certainly judge the sentence *John is blond* to be true. One way of thinking about this could be as follows: I don't associate *John* with a single semantic value and therefore don't associate the sentence with a single semantic value. At the same time, I consider a variety of ways of associating *John* with a semantic value, and each of these yields a semantic value for the sentence that would lead to a "true" judgment. This way of thinking suggests a different approach to (6) and (7). What changes could it cause in the final theory?