Basics of intensionality

Semantics II

January 29 & February 1, 2018

Up from extensions

The extensional state of play

The possible types of linguistic meanings:

$$\tau ::= \underbrace{e \mid t \mid \tau \to \tau}_{\text{Backus-Naur Form}}$$

Basic ways for meanings to compose (FA, BA, PM):

$$[\![A\ B]\!]^g = \begin{cases} [\![A]\!]^g\,[\![B]\!]^g & \text{or} \\ [\![B]\!]^g\,[\![A]\!]^g & \text{or} \\ \lambda x.\,[\![A]\!]^g\,x \wedge [\![B]\!]^g\,x & \text{whichever's defined} \end{cases}$$

Predicate abstraction for variable binding:

$$[\![n\ A]\!]^g = \lambda x.\,[\![A]\!]^{g[n \to x]}$$

That is the essence of Heim & Kratzer (1998).

Whence communication?

Language is (often) used for *communication*: encoding and transmitting knowledge about the ways the world is or is not.

It is not so obvious how the extensional system is compatible with communication. For expressions to have interpretations at all, we need to fix a model M. But then M determines all the relevant facts (whether it's raining, who saw whom, etc.).

Non-synonymy

Let's suppose a model with the following features:

- It's cloudy.
- It's raining.

Then relative to this model, we'll have the following:

$$[\![\text{it's cloudy}]\!]^g = [\![\text{it's raining}]\!]^g$$

But it's quite obvious these expressions do not mean the same thing.

Non-synonymy

Let's suppose a model with the following features:

- It's cloudy.
- It's raining.

Then relative to this model, we'll have the following:

$$[\![\text{it's cloudy}]\!]^g = [\![\text{it's raining}]\!]^g$$

But it's quite obvious these expressions do not mean the same thing.

The obvious thing to do here is to avoid identifying meanings *per se* with semantic value relative to a model: *it's cloudy* and *it's raining* mean different things because there are models M where $[\![\text{it's cloudy}]\!]^{M,g} \neq [\![\text{it's raining}]\!]^{M,g}$.

Non-truth-functionality

Now, let's suppose a model with the following features:

- It's cloudy.
- It's raining.
- I thinks that it's cloudy, but I don't think that it's raining.

In view of these facts, *we would hope to* have the following:

$$[\![\text{Simon thinks it's cloudy}]\!]^g \neq [\![\text{Simon thinks it's raining}]\!]^g$$

Could we?

Non-truth-functionality

Now, let's suppose a model with the following features:

- It's cloudy.
- It's raining.
- I thinks that it's cloudy, but I don't think that it's raining.

In view of these facts, *we would hope to* have the following:

$$[\![\text{Simon thinks it's cloudy}]\!]^g \neq [\![\text{Simon thinks it's raining}]\!]^g$$

Could we? We couldn't. Compositionality is hard-wired into the system:

$$\text{If } [\![\Gamma]\!]^g = [\![\Delta]\!]^g\text{, then } [\![\ldots\Gamma\ldots]\!]^g = [\![\ldots\Delta\ldots]\!]^g.$$

What are some other cases where this kind of substitution fails?

Circumstantial “variability”

We know that the extensions (i.e., truth values) of expressions like *it’s raining* or *it’s cloudy* **can shift**, depending on the circumstances.

We’ve also seen that this shifty-ness is relevant for judgments of synonymy, and that it can have interpretive consequences in embedded contexts (e.g., under *believes*).

Our task today will be to extend our semantics to do justice to this variability (his will also help us say something useful about communication).

Thinking back to pronouns

Pronouns are chameleons: their values shift with the context of utterance — and sometimes multiple times within a single sentence!

1. John saw **her**.
2. Every philosopher$_i$ thinks **they**$_i$'re a genius.

We needed some way to interpret pronouns. The standard device is an *assignment function*, essentially something whose only job is to value pronouns.

Once we have assignments, we can use them to for interpreting certain kinds of things, and (more or less) ignore them in other cases:

$$[\![\text{she}_n]\!]^g = g_n \qquad [\![\text{saw}]\!]^g = \textbf{saw} \qquad \ldots$$

Circumstantial "variability", redux

The standard solution to the problems we've seen is *remarkably* similar in character to the way that variables and variable binding are handled.

Essentially, we will treat expressions like *it's raining* as kinds of variables. Rather than varying with an assignment function, though, they'll vary with something like *the circumstances* under which a sentence is evaluated.

- *It's raining* is True in certain circumstances and False in others.

It is convenient (though not necessary) to identify these circumstances with **possible worlds**. A possible world an extremely complete piece of data, something that renders a True-or-False verdict for any (True-or-False) sentence imaginable.

- Further, because possible worlds must be *possible*, a world can't have it that *it's raining* and *it's not raining* are both True.

Circumstantial "variability" as world variability

Just like pronouns' interpretations are relativized to the assignment, the interpretation of some expressions will now be relativized to a possible world:

$$[\![\text{it's raining}]\!]^{w,g} = \text{True} \iff \text{it's raining in } w$$

For convenience, we'll generally abbreviate the right-hand side:

$$[\![\text{it's raining}]\!]^{w,g} = \textbf{rain}_w$$

We can imagine a world *u* where it's raining and a world *v* where it isn't. Relative to these worlds, *it's raining* will have different extensions:

$$[\![\text{it's raining}]\!]^{u,g} = \textbf{rain}_u = \text{True} \qquad [\![\text{it's raining}]\!]^{v,g} = \textbf{rain}_v = \text{False}$$

World-variability and lack thereof

What kinds of expressions have meanings that depend on a possible world?

- Do nouns and adjectives?

World-variability and lack thereof

What kinds of expressions have meanings that depend on a possible world?

- Do nouns and adjectives? Yes! (Think: sleepy professors.)
- Do intransitive verbs?

World-variability and lack thereof

What kinds of expressions have meanings that depend on a possible world?

- Do nouns and adjectives? Yes! (Think: sleepy professors.)
- Do intransitive verbs? Yes! (Different people snore in different worlds.)
- Do transitive verbs?

World-variability and lack thereof

What kinds of expressions have meanings that depend on a possible world?

- Do nouns and adjectives? Yes! (Think: sleepy professors.)
- Do intransitive verbs? Yes! (Different people snore in different worlds.)
- Do transitive verbs? Of course! (I'm teaching you right now, but it's not *so* hard to imagine the shoe being on the other foot.)
- How about connectives like *and*?

World-variability and lack thereof

What kinds of expressions have meanings that depend on a possible world?

- Do nouns and adjectives? Yes! (Think: sleepy professors.)
- Do intransitive verbs? Yes! (Different people snore in different worlds.)
- Do transitive verbs? Of course! (I'm teaching you right now, but it's not *so* hard to imagine the shoe being on the other foot.)
- How about connectives like *and*? Seems like no: no matter what the world is like, *p and q* requires both *p* and *q* to be True.
- How about determiner meanings like *the* and *every*?

World-variability and lack thereof

What kinds of expressions have meanings that depend on a possible world?

- Do nouns and adjectives? Yes! (Think: sleepy professors.)
- Do intransitive verbs? Yes! (Different people snore in different worlds.)
- Do transitive verbs? Of course! (I'm teaching you right now, but it's not *so* hard to imagine the shoe being on the other foot.)
- How about connectives like *and*? Seems like no: no matter what the world is like, *p and q* requires both *p* and *q* to be True.
- How about determiner meanings like *the* and *every*? Again, seems like no (for the same reason as the connectives).
- How about proper names?

World-variability and lack thereof

What kinds of expressions have meanings that depend on a possible world?

- Do nouns and adjectives? Yes! (Think: sleepy professors.)
- Do intransitive verbs? Yes! (Different people snore in different worlds.)
- Do transitive verbs? Of course! (I'm teaching you right now, but it's not *so* hard to imagine the shoe being on the other foot.)
- How about connectives like *and*? Seems like no: no matter what the world is like, *p and q* requires both *p* and *q* to be True.
- How about determiner meanings like *the* and *every*? Again, seems like no (for the same reason as the connectives).
- How about proper names? This is somewhat contentious. Many people take proper names to have world-invariant meanings (following Kripke). There are dissenters though. For convenience, we'll be Kripkean today.

Expressions whose meanings don't vary with possible worlds are called *rigid*.

World-variability and lack thereof, in the semantics

It's easy to model world-(in)variance in our semantics. Some expressions will have meanings that depend on the possible world:

$$[\![\text{sleepy}]\!]^{w,g} = \textbf{sleepy}_w \qquad [\![\text{prof}]\!]^{w,g} = \textbf{prof}_w \qquad [\![\text{teaches}]\!]^{w,g} = \textbf{teaches}_w$$

And other expressions won't make use of the possible world at all:

$$[\![\text{and}]\!]^{w,g} = \lambda q.\lambda p.\, p \wedge q \qquad [\![\text{every}]\!]^{w,g} = \lambda f.\lambda g.\forall x : f\,x \Rightarrow g\,x \qquad [\![\text{Al}]\!]^{w,g} = \textbf{a}$$

It bears re-emphasizing that this is precisely analogous to the situation with pronouns, though world-variability seems to be far more pervasive!

The intensional state of play, Take 1

The possible types of linguistic meanings:

$$\tau ::= \underbrace{e \mid t \mid \tau \to \tau}_{\text{Unchanged from before}}$$

Basic ways for meanings to compose (FA, BA, PM):

$$[\![A\ B]\!]^{w,g} = \begin{cases} [\![A]\!]^{w,g}\,[\![B]\!]^{w,g} & \text{or} \\ [\![B]\!]^{w,g}\,[\![A]\!]^{w,g} & \text{or} \\ \lambda x.\,[\![A]\!]^{w,g}\,x \wedge [\![B]\!]^{w,g}\,x & \text{whichever's defined} \end{cases}$$

Predicate abstraction for variable binding:

$$[\![n\ A]\!]^{w,g} = \lambda x.\,[\![A]\!]^{w,g[n \to x]}$$

Variability projects up

Given this setup, it's straightforward to calculate meanings for complex phrases:

$$[\![\text{every good ling loves semantics}]\!]^{w,g} = \forall x.\,(\mathbf{good}_w\,x \wedge \mathbf{ling}_w\,x) \Rightarrow \mathbf{loves}_w\,\mathbf{s}\,x$$

The result is True at worlds where the good linguists all love semantics, and false at worlds where any of the good linguists don't. This is as desired.

Notice that meanings can depend on both the world and the assignment:

$$[\![\text{she}_1 \text{ knows Petrous}]\!]^{w,g} = \mathbf{knows}_w\,\mathbf{p}\,g_1$$

A complete derivation

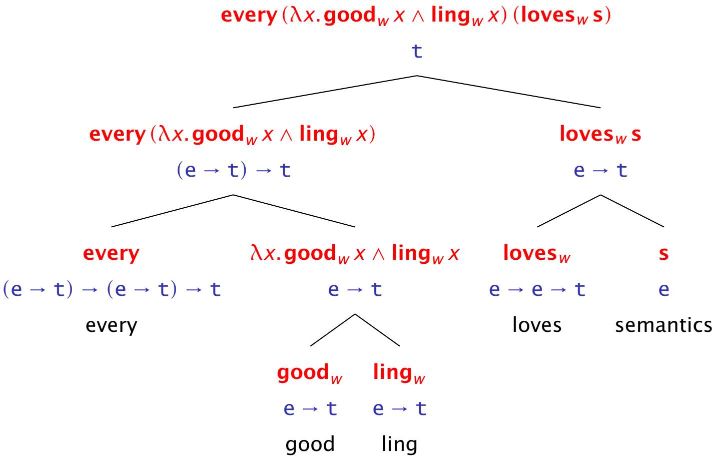

(Since **every** abbreviates $\lambda f.\lambda g.\forall x : f\,x \Rightarrow g\,x$, this β-reduces to the term on the last slide.)

Stepping back

We've folded in circumstantial variability, by extending our models with possible worlds, and extending $[\![\cdot]\!]$ to compose meanings relative to a possible world.

But have we solved our problems?

Stepping back

We've folded in circumstantial variability, by extending our models with possible worlds, and extending $[\![\cdot]\!]$ to compose meanings relative to a possible world.

But have we solved our problems? Not yet, but we now have the resources to do so.

The meaning of an expression

We can now define *the meaning of an expression* as its **intension**:

The **meaning** of A at an assignment g is $\lambda w.[\![A]\!]^{w,g}$.

von Fintel & Heim (2011) notate this as $[\![A]\!]^{g}_{¢}$. Sometimes you'll simply see $[\![A]\!]^{g}$.

To have functions from worlds to extensions as possible denotations, we must tweak the type theory a little bit (s is the type of worlds):

$$\tau ::= \mathrm{e} \mid \mathrm{t} \mid \underbrace{\mathrm{s} \to \tau}_{\text{Adding intensions}} \mid \tau \to \tau$$

Food for thought: why not simply say $\tau ::= \mathrm{e} \mid \mathrm{t} \mid \mathrm{s} \mid \tau \to \tau$?

Kinds of intensions

The meaning of a sentence is a function from possible worlds to True or False:

$$\begin{aligned}[\![\text{she}_1 \text{ teaches Al}]\!]^g_{¢} &= \lambda w.\, [\![\text{she}_1 \text{ teaches Al}]\!]^{w,g} \\ &= \lambda w.\, \mathbf{teaches}_w\, \mathbf{a}\, g_1\end{aligned}$$

This is (isomorphic to) a set of worlds, also known as a *proposition*:

$$\{w \mid \mathbf{teaches}_w\, \mathbf{a}\, g_1\}$$

The meaning of definite descriptions is a function from worlds to individuals. This is known as an *individual concept*:

$$[\![\text{the prof}]\!]^g_{¢} = \lambda w.\, \iota x.\, \mathbf{prof}_w\, x$$

Exercise: calculate $[\![\text{the professor who slept snored}]\!]^g_{¢}$.

A calculation

$[\![\text{the prof who 1 } t_1 \text{ slept snored}]\!]^{g}_{¢} =$

A calculation

$$
\begin{aligned}
[\![\text{the prof who 1 } \text{t}_1 \text{ slept snored}]\!]^g_{¢} &= \lambda w.\,[\![\text{the prof who 1 } \text{t}_1 \text{ slept snored}]\!]^{w,g} \qquad ¢ \\
&=
\end{aligned}
$$

A calculation

$$
\begin{aligned}
[\![\text{the prof who 1 t}_1\text{ slept snored}]\!]^{g}_{¢} &= \lambda w.\,[\![\text{the prof who 1 t}_1\text{ slept snored}]\!]^{w,g} && ¢\\
&= \lambda w.\,\textbf{snored}_w\,[\![\text{the prof who 1 t}_1\text{ slept}]\!]^{w,g} && \text{BA}\\
&=
\end{aligned}
$$

A calculation

$$\begin{aligned}
[\![\text{the prof who 1 } \text{t}_1 \text{ slept snored}]\!]^g_{¢} &= \lambda w.\, [\![\text{the prof who 1 } \text{t}_1 \text{ slept snored}]\!]^{w,g} && ¢ \\
&= \lambda w.\, \textbf{snored}_w\, [\![\text{the prof who 1 } \text{t}_1 \text{ slept}]\!]^{w,g} && \text{BA} \\
&= \lambda w.\, \textbf{snored}_w\, (\textbf{the}\, [\![\text{prof who 1 } \text{t}_1 \text{ slept}]\!]^{w,g}) && \text{FA} \\
&=
\end{aligned}$$

A calculation

$$
\begin{aligned}
[\![\text{the prof who 1 t}_1\text{ slept snored}]\!]^{g}_{¢} &= \lambda w.\,[\![\text{the prof who 1 t}_1\text{ slept snored}]\!]^{w,g} && ¢ \\
&= \lambda w.\,\textbf{snored}_w\,[\![\text{the prof who 1 t}_1\text{ slept}]\!]^{w,g} && \text{BA} \\
&= \lambda w.\,\textbf{snored}_w\,(\textbf{the}\,[\![\text{prof who 1 t}_1\text{ slept}]\!]^{w,g}) && \text{FA} \\
&= \lambda w.\,\textbf{snored}_w\,(\textbf{the}\,(\lambda x.\,\textbf{prof}_w\,x \wedge [\![\text{1 t}_1\text{ slept}]\!]^{w,g}\,x)) && \text{PM} \\
&=
\end{aligned}
$$

A calculation

$$
\begin{aligned}
[\![\text{the prof who 1 t}_1\text{ slept snored}]\!]^g_¢ &= \lambda w.\, [\![\text{the prof who 1 t}_1\text{ slept snored}]\!]^{w,g} && ¢ \\
&= \lambda w.\, \mathbf{snored}_w\, [\![\text{the prof who 1 t}_1\text{ slept}]\!]^{w,g} && \text{BA} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, [\![\text{prof who 1 t}_1\text{ slept}]\!]^{w,g}) && \text{FA} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, (\lambda x.\, \mathbf{prof}_w\, x \wedge [\![\text{1 t}_1\text{ slept}]\!]^{w,g}\, x)) && \text{PM} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, (\lambda x.\, \mathbf{prof}_w\, x \wedge (\lambda y.\, \mathbf{slept}_w\, y)\, x)) && \text{PA, BA} \\
&=
\end{aligned}
$$

A calculation

$$
\begin{aligned}
[\![\text{the prof who 1 t}_1\text{ slept snored}]\!]^g_{¢} &= \lambda w.\,[\![\text{the prof who 1 t}_1\text{ slept snored}]\!]^{w,g} && ¢\\
&= \lambda w.\,\mathbf{snored}_w\,[\![\text{the prof who 1 t}_1\text{ slept}]\!]^{w,g} && \text{BA}\\
&= \lambda w.\,\mathbf{snored}_w\,(\mathbf{the}\,[\![\text{prof who 1 t}_1\text{ slept}]\!]^{w,g}) && \text{FA}\\
&= \lambda w.\,\mathbf{snored}_w\,(\mathbf{the}\,(\lambda x.\,\mathbf{prof}_w\,x \wedge [\![\text{1 t}_1\text{ slept}]\!]^{w,g}\,x)) && \text{PM}\\
&= \lambda w.\,\mathbf{snored}_w\,(\mathbf{the}\,(\lambda x.\,\mathbf{prof}_w\,x \wedge (\lambda y.\,\mathbf{slept}_w\,y)\,x)) && \text{PA, BA}\\
&= \lambda w.\,\mathbf{snored}_w\,(\mathbf{the}\,(\lambda x.\,\mathbf{prof}_w\,x \wedge \mathbf{slept}_w\,x)) && \beta\\
&=
\end{aligned}
$$

A calculation

$$
\begin{aligned}
[\![\text{the prof who 1 } \mathrm{t}_1 \text{ slept snored}]\!]^{g}_{¢} &= \lambda w.\, [\![\text{the prof who 1 } \mathrm{t}_1 \text{ slept snored}]\!]^{w,g} && ¢ \\
&= \lambda w.\, \mathbf{snored}_w\, [\![\text{the prof who 1 } \mathrm{t}_1 \text{ slept}]\!]^{w,g} && \text{BA} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, [\![\text{prof who 1 } \mathrm{t}_1 \text{ slept}]\!]^{w,g}) && \text{FA} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, (\lambda x.\, \mathbf{prof}_w\, x \wedge [\![\text{1 } \mathrm{t}_1 \text{ slept}]\!]^{w,g}\, x)) && \text{PM} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, (\lambda x.\, \mathbf{prof}_w\, x \wedge (\lambda y.\, \mathbf{slept}_w\, y)\, x)) && \text{PA, BA} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, (\lambda x.\, \mathbf{prof}_w\, x \wedge \mathbf{slept}_w\, x)) && \beta \\
&= \lambda w.\, \mathbf{snored}_w\, (\iota x.\, \mathbf{prof}_w\, x \wedge \mathbf{slept}_w\, x) && \textbf{the}, \beta \\
&\simeq
\end{aligned}
$$

A calculation

$$
\begin{aligned}
[\![\text{the prof who 1 } \mathrm{t}_1 \text{ slept snored}]\!]^{g}_{¢} &= \lambda w.\,[\![\text{the prof who 1 } \mathrm{t}_1 \text{ slept snored}]\!]^{w,g} && ¢ \\
&= \lambda w.\,\mathbf{snored}_w\,[\![\text{the prof who 1 } \mathrm{t}_1 \text{ slept}]\!]^{w,g} && \text{BA} \\
&= \lambda w.\,\mathbf{snored}_w\,(\mathbf{the}\,[\![\text{prof who 1 } \mathrm{t}_1 \text{ slept}]\!]^{w,g}) && \text{FA} \\
&= \lambda w.\,\mathbf{snored}_w\,(\mathbf{the}\,(\lambda x.\,\mathbf{prof}_w\,x \wedge [\![\text{1 } \mathrm{t}_1 \text{ slept}]\!]^{w,g}\,x)) && \text{PM} \\
&= \lambda w.\,\mathbf{snored}_w\,(\mathbf{the}\,(\lambda x.\,\mathbf{prof}_w\,x \wedge (\lambda y.\,\mathbf{slept}_w\,y)\,x)) && \text{PA, BA} \\
&= \lambda w.\,\mathbf{snored}_w\,(\mathbf{the}\,(\lambda x.\,\mathbf{prof}_w\,x \wedge \mathbf{slept}_w\,x)) && \beta \\
&= \lambda w.\,\mathbf{snored}_w\,(\iota x.\,\mathbf{prof}_w\,x \wedge \mathbf{slept}_w\,x) && \mathbf{the}, \beta \\
&\simeq \{w \mid \mathbf{snored}_w\,(\iota x.\,\mathbf{prof}_w\,x \wedge \mathbf{slept}_w\,x)\} && \simeq
\end{aligned}
$$

A calculation

$$
\begin{aligned}
[\![\text{the prof who 1 } t_1 \text{ slept snored}]\!]^{g}_{¢} &= \lambda w.\, [\![\text{the prof who 1 } t_1 \text{ slept snored}]\!]^{w,g} && ¢ \\
&= \lambda w.\, \mathbf{snored}_w\, [\![\text{the prof who 1 } t_1 \text{ slept}]\!]^{w,g} && \text{BA} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, [\![\text{prof who 1 } t_1 \text{ slept}]\!]^{w,g}) && \text{FA} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, (\lambda x.\, \mathbf{prof}_w\, x \wedge [\![\text{1 } t_1 \text{ slept}]\!]^{w,g}\, x)) && \text{PM} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, (\lambda x.\, \mathbf{prof}_w\, x \wedge (\lambda y.\, \mathbf{slept}_w\, y)\, x)) && \text{PA, BA} \\
&= \lambda w.\, \mathbf{snored}_w\, (\mathbf{the}\, (\lambda x.\, \mathbf{prof}_w\, x \wedge \mathbf{slept}_w\, x)) && \beta \\
&= \lambda w.\, \mathbf{snored}_w\, (\iota x.\, \mathbf{prof}_w\, x \wedge \mathbf{slept}_w\, x) && \mathbf{the}, \beta \\
&\simeq \{w \mid \mathbf{snored}_w\, (\iota x.\, \mathbf{prof}_w\, x \wedge \mathbf{slept}_w\, x)\} && \simeq
\end{aligned}
$$

Small book-keeping note: the second-to-last line is a *partial* function from worlds to truth values. The last line is a set of worlds where that partial function is true (ergo, also defined).

Attitudes and modals

Back to attitudes

In principle, we want to be able to derive the result below, even at worlds w where it's cloudy and it's raining. Can we? Not exactly, but we are *so close*.

$$[\![\text{Simon thinks it's cloudy}]\!]^{w,g} \neq [\![\text{Simon thinks it's raining}]\!]^{w,g}$$

The idea: even though, at w, it may be true that $[\![\text{it's cloudy}]\!]^{w,g} = [\![\text{it's raining}]\!]^{w,g}$, it will in general also be true that $[\![\text{it's cloudy}]\!]^{g}_{¢} \neq [\![\text{it's raining}]\!]^{g}_{¢}$ (since, in the kinds of intensional models we care about, the sentences express different propositions).

The intuition is that *thinks* composes with the *intension* of the embedded clause:

$$\mathbf{thinks}_w :: (\mathtt{s} \to \mathtt{t}) \to \mathtt{e} \to \mathtt{t}$$

The final state of play

The possible types of linguistic meanings:

$$\tau ::= e \mid t \mid \underbrace{s \to \tau}_{\text{Adding intensions}} \mid \tau \to \tau$$

Basic ways for meanings to compose (extended to allow application to an ***intension****):*

$$[\![A\ B]\!]^{w,g} = \begin{cases} [\![A]\!]^{w,g}\,[\![B]\!]^{w,g} \quad \text{or} \quad [\![B]\!]^{w,g}\,[\![A]\!]^{w,g} \quad \text{or} & FA, BA \\ [\![A]\!]^{w,g}\,[\![B]\!]^{g}_{¢} \quad \text{or} \quad [\![B]\!]^{w,g}\,[\![A]\!]^{g}_{¢} \quad \text{or} & IFA, IBA \\ \lambda x.\,[\![A]\!]^{w,g}\,x \wedge [\![B]\!]^{w,g}\,x & PM\text{, whichever's defined} \end{cases}$$

Predicate abstraction for variable binding (unchanged):

$$[\![n\ A]\!]^{w,g} = \lambda x.\,[\![A]\!]^{w,g[n \to x]}$$

The choice of rule at any point is stull fully deterministic (i.e., strongly type-driven).

Deriving an attitude ascription

$[\![\text{Simon thinks it's raining}]\!]^{w,g} =$

Deriving an attitude ascription

$$\begin{aligned}[\![\text{Simon thinks it's raining}]\!]^{w,g} &= [\![\text{thinks it's raining}]\!]^{w,g}\,\mathbf{s} && \text{BA} \\ &= \end{aligned}$$

Deriving an attitude ascription

$$
\begin{aligned}
[\![\text{Simon thinks it's raining}]\!]^{w,g} &= [\![\text{thinks it's raining}]\!]^{w,g}\,\mathbf{s} && \text{BA}\\
&= \mathbf{thinks}_w\,[\![\text{it's raining}]\!]^{g}_{¢}\,\mathbf{s} && \text{IFA}\\
&=
\end{aligned}
$$

Deriving an attitude ascription

$$
\begin{aligned}
[\![\text{Simon thinks it's raining}]\!]^{w,g} &= [\![\text{thinks it's raining}]\!]^{w,g}\,\mathbf{s} && \text{BA}\\
&= \mathbf{thinks}_w\,[\![\text{it's raining}]\!]^{g}_{¢}\,\mathbf{s} && \text{IFA}\\
&= \mathbf{thinks}_w\,(\lambda v.\,[\![\text{it's raining}]\!]^{v,g})\,\mathbf{s} && ¢\\
&=
\end{aligned}
$$

Deriving an attitude ascription

$$
\begin{aligned}
[\![\text{Simon thinks it's raining}]\!]^{w,g} &= [\![\text{thinks it's raining}]\!]^{w,g}\,\mathbf{s} && \text{BA}\\
&= \mathbf{thinks}_w\,[\![\text{it's raining}]\!]^{g}_{¢}\,\mathbf{s} && \text{IFA}\\
&= \mathbf{thinks}_w\,(\lambda v.\,[\![\text{it's raining}]\!]^{v,g})\,\mathbf{s} && ¢\\
&= \mathbf{thinks}_w\,(\lambda v.\,\mathbf{rain}_v)\,\mathbf{s} && \text{Lex}
\end{aligned}
$$

The result requires me to stand in the thinking-at-w relation to the *proposition* that it's raining, which is in general different from the proposition that it's cloudy!

Note that we can also, of course, calculate the intension of *Simon thinks it's raining*:

$$\lambda w.\,\mathbf{thinks}_w\,(\lambda v.\,\mathbf{rain}_v)\,\mathbf{s}$$

Which will be useful for sentences like *Bob thinks Simon thinks it's raining*.

Thinking-at-w?

We have as yet an unanalyzed piece of meaning: $\mathbf{thinks}_w$.

We know its type has to be $(\mathtt{s} \to \mathtt{t}) \to \mathtt{e} \to \mathtt{t}$. But it would be useful to say a bit more about what kind of meaning this represents.

Thinking-at-w?

We have as yet an unanalyzed piece of meaning: $\textbf{thinks}_w$.

We know its type has to be $(\mathtt{s} \to \mathtt{t}) \to \mathtt{e} \to \mathtt{t}$. But it would be useful to say a bit more about what kind of meaning this represents.

Let us assume (reasonably) that agents entertain a variety of candidates for the actual world. Then we can characterize an agent's belief-state in a world as follows:

$$\textbf{Bel}_{x,w} = \{v \mid x \text{ considers } v \text{ a candidate for the actual world in } w\}$$

Thinking-at-w?

We have as yet an unanalyzed piece of meaning: $\textbf{thinks}_w$.

We know its type has to be $(\mathtt{s} \to \mathtt{t}) \to \mathtt{e} \to \mathtt{t}$. But it would be useful to say a bit more about what kind of meaning this represents.

Let us assume (reasonably) that agents entertain a variety of candidates for the actual world. Then we can characterize an agent's belief-state in a world as follows:

$$\textbf{Bel}_{x,w} = \{v \mid x \text{ considers } v \text{ a candidate for the actual world in } w\}$$

The last step: $\textbf{thinks}_w$ *universally quantifies over* its subject's belief-worlds:

$$\textbf{thinks}_w = \lambda p.\lambda x.\forall v \in \textbf{Bel}_{x,w} : p\,v$$

Putting the pieces together

Our compositionally derived meaning from a couple slides back:

$$\textbf{thinks}_w\,(\lambda v.\,\textbf{rain}_v)\,\textbf{s}$$

Substituting in the definition for $\textbf{thinks}_w$ and doing a couple β-reductions:

$$\forall v \in \textbf{Bel}_{\textbf{s},w} : \textbf{rain}_v$$

In other words, I only entertain worlds in which it's raining. Seems reasonable.

Multiple formulations

We can decompose **Bel** a bit further. Evidently, it determines, based on an agent x and an initial evaluation world w, a set of worlds that x considers candidates for the actual world in w:

$$\mathbf{Bel}_x :: \mathsf{s} \to \{\mathsf{s}\}$$

This is isomorphic to a relation between worlds:

$$\mathbf{Bel}_x :: \mathsf{s} \to \mathsf{s} \to \mathsf{t}$$

These kinds of relations between worlds are so common, and important, that they have a special name: **accessibility relations**.

$$\text{Notational conventions: } w\, R\, u \iff R\, w\, u$$
$$w\, R\, u\, S\, v \iff R\, w\, u \wedge S\, u\, v$$

Properties of accessibility relations

Researchers have found it fruitful to think about the formal properties that accessibility relations might have. Are they reflexive? Transitive? Symmetric?

Reflexive	$w\ R\ w$
Transitive	$w\ R\ u\ R\ v \Longrightarrow w\ R\ v$
Symmetric	$w\ R\ u \Longrightarrow u\ R\ w$

We won't dwell on these, but it's worth thinking about which $\mathbf{Bel}_x$ should satisfy:

- Reflexivity: if you believe p, it's true. This can't be right.
- Transitivitity: if you believe p, then you believe you believe p. Perhaps!
- Symmetry: if p is true, you believe it's possible. That seems pretty fishy.

Modals

Modals can be handled very similarly to attitude verbs. The principal difference is that modals are *context-sensitive*:

$$\textbf{may}_w = \underbrace{\lambda R.\lambda p.\, \exists v : w\, R\, v \wedge p\, v}_{(\mathrm{s}\to\mathrm{s}\to\mathrm{t})\to(\mathrm{s}\to\mathrm{t})\to\mathrm{t}}$$

In general *R* will be filled by a variable in the logical form:

- [may R] [it be raining]

This variable gets its value from the assignment (i.e., from the context). This allows different kinds of worlds to be quantified over, depending on what's relevant:

- You may eat a cookie.
- John may be in the garden.

von Fintel, Kai & Irene Heim. 2011. Intensional semantics, Spring 2011 edition.

Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Oxford: Blackwell.