Questions: Round 2

Hi dear students --

Here's second round of questions and my answers/discussion. Sorry for being late with them. It's been lots of fun thinking about the questions you put, although they're by and large somewhat trickier than the last set. Good for my mental processes, though. Feel free to send follow-ups, or, if we get time, to bring further questions up in class! As before, questions are in italics, answers are normal. These questions appear in no particular order; in fact, Travis, who's last, sent me his question first (actually, now that I see what order I've got them in, I think they're in order of [most general]→ [most specific]). Apologies, Travis, for taking so long composing this answer! Katie and Erika seem to have arrived at similar states of mind as a result of all this coursework, as once again, (at least one of) their questions are similar in intent. (So too, apparently, for Kris and Gondy, although Gondy then later changed her question to a different one). Probably means something (something good, I hope!) The discussion for most of the questions, though, is in some wise relevant to the discussion for others, so I guess you've all arrived at similar states of mind. Some of the answers are very long this time -- sorry about that!

Enjoy -
hh

Gondy

The whole semantics representation for the meaning/truth condition of a sentence seems to be funded on syntactic trees. And this seems to be the case regardless of the discussion: "building the tree + then meaning" OR "get meaning while you build the tree".

So here are my questions:
1. is it assumed that people actually build these trees?
2. if it is NOT believed that people use it, why are we using it? and what is used instead of trees?
3. if it is believed that people use it, is there evidence for it?

I believe more in mental models than this, for the following reasons:
· *if your tree is incorrect --> understanding of the sentence will be wrong*
· *if your model is incorrect --> your model will probably be incomplete and you will still be able to understand a sentence (in a way)*
· *can everyone build these trees correctly? why doesn't it come naturally? why devote a whole syntax class to it?*

(Ed's note: Later postscript*:*

I can actually simplify my question quite a lot. What do you think of this summary question:

--------- *Why trees?* ----------)

Heidi: I think ultimately what you're getting at is the question whether this is a psychologically realistic model of linguistic performance on-line or not, but I'm going to take it at face value at first and explain what trees are and what they're supposed to do.

Trees are a model of linguistic representations — models of mental representations of sentences . What they are intended to represent, in their most raw form, are constituency facts. It's a fact of most (probably all) languages that certain words in a sentence form units (constituents) with certain other words, and fail to form units with yet other words. So, for instance, "the" and "cat"

form a unit (you can tell it's a unit by a number of tests, e.g., you can replace "the cat" with a single word, "it", and still have a well-formed sentence). However, "the" and "arrive" don't form a unit ("*the arrive"). Then you take these sub-units ("the cat", "arrived") and put them together to form a bigger unit, in this case the sentence ("the cat arrived"). (Again, you can tell a sentence is a unit with a number of tests -- it's not just an intuition that 'these words sort of belong together' or anything as imprecise as that).

So, in fact, trees are simply a convenient notation with which to represent certain syntactic groupings. In the earliest work on generative grammar, sentence structure wasn't represented using the tree notation, but rather, with brackets:

1. (a) $[_{S}\ [_{DP}$The cat] $[_{VP}$arrived]]
   (b) [Sue [[went [to [the garden]] [with Mary]] and [picked [a [[very red] apple] [to eat]]]]]

The bracketing in (1a) for example, is simply a notationally equivalent representation of the tree structure in 2 (modulo the category labels on the lexical items, which we can probably take for granted):

2.

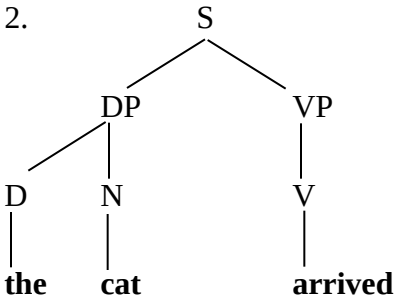

So, in fact, trees are not necessarily a representation of *semantic* facts, but rather a representation of *syntactic* facts. And yes, it is supposed that they are a good model of what people actually do in their heads when they construct or interpret a sentence. Nearly all theories of language include some form of structural representation of constituency facts. This is because what a theory of language is trying to do is to model what it is that people know when they know a language -- and one of the things that people know, apparently, is the ways in which words group together.

Now, the ultimate goal of a theory of the language faculty is how you get from the string of sounds "The cat chased the rat" to the conscious understanding of the proposition [[the cat chased the rat]]. One thing you clearly do is break the string up into words, which have meaning of their own — words are symbols for certain meanings. Then, you also break the string of words up into small groups of words. The null hypothesis about why you would do this is that the groups themselves also have a meaning of their own. And the groups fit together in certain well-defined ways (DPs fit into the same position in the tree relative to Vs, over and over again) to make a whole sentence. So presumably, you do interpret a sentence piecewise, first figuring out what the words mean, then figuring out what the subgroups mean, then figuring out what the whole sentence means. (Obviously you can still do this while you're getting new parts of the string — these processes go on in parallel, not in sequence. You don't have to hear a whole sentence to understand its first half, and to understand that that first half by itself is incomplete).

One thing that isn't clear to me from your question is why "mental models" seem like a plausible replacement for a model of the language faculty that represents constituency facts, probably because I don't know what kind of "mental models" you have in mind. What I think you

must be referring to is an idea that the meaning of a sentence is in some way a "mental picture" of the proposition it denotes. In fact, that's exactly to be the end result of combining the meanings of all the subparts of the sentence: a specific way of putting together all the little "mental pictures" that the individual words denote to get to a big "mental picture" of the whole proposition. Perhaps you're intending to imply that the structural relationships that words enter into don't necessarily have anything to do with the meaning of the whole sentence — but if that's the case, then it's not clear to me why any combinations of the words "the cat chased the rat" don't mean the same thing ("The cat chased the rat" is not equivalent to "the rat chased the cat" for example, or "the chased the cat rat" or "chased rat cat the the" etc. etc., to put it explicitly).

Ok. Now to specifically answer some of your your specific points (at last, eh?)

1. is it assumed that people actually build these trees?

Well, not in the sense that they have a little piece of paper in their heads and they draw lines and such. But yes, they do put the words of a sentence together in the way that the tree structure (or bracketing, or whatever equivalent) represents.

2. if it is NOT believed that people use it, why are we using it? and what is used instead of trees?

I'll interpret this as, what are the neural correlates of syntactic structure? Nobody knows, but they're studying it. Ask again in about 50 years; we'll know a lot more about the actual neural representation of syntactic structure by then.

Ok, and on to your specific arguments:

· *if your tree is incorrect --> understanding of the sentence will be wrong*

This question seems to be, what if you misparse the tree? Here's a potential example:

"Sue hit the girl with the book"

This can mean either "Sue hit the girl who is holding the book" or "Sue hit the girl by using the book".

Now, if you misparse the structure, attaching the PP "with the book" to the VP instead of the NP, you'll end up interpreting the sentence as meaning the second thing, when it was intended to mean the first thing, So in some sense, your understanding of the sentence is wrong. However, you still get it partially right -- you understand it has to do with Sue, and a girl, and an act of hitting (by Sue) and a book fits into it somewhere. So you'll still be able to understand the sentence sort of.

This example, however, is probably too easy. What if, when you're building the tree, you accidentally put "the" by itself as the sister of the verb, and adjoin "girl" to the VP? Then you really will be in interpretive trouble (although you'll still be able to access the meaning of the words -- so you can still see what the sentence is about, again, sort of). But you won't understand that "the girl" is supposed to be the object of the hitting.

The answer is, not only do we all know how to build trees, but tree-building is hardwired into our processing mechanism in such a way that we can't do it as wrongly as that. In generative grammar, we represent this hardwiring as a set of rules for generating tree structures — in their simplest form, they look simply like rewriting rules:
S --> DP VP
VP --> V (DP) (PP)
PP--> P DP ... etc.

· *if your model is incorrect --> your model will probably be incomplete and you will still be able to understand a sentence (in a way)*

but this is still true on an account whereby tree-interpretation generates a mental model....

· *can everyone build these trees correctly? why doesn't it come naturally? why devote a whole syntax class to it?*

Aha! the question everyone wonders sooner or later. But think about it: you can walk, or catch a ball, right? But to make an explicit model of exactly *how* someone walks, or catches a ball, you have to know a hell of a lot of math, physics and engineering — these are not things that come naturally. You're programmed to walk and catch a ball, but you don't have any direct access to the code. Similarly, you're programmed to parse sentences, but you don't have any direct access to the code, so in order to figure out what the program is doing, we have to look at the sentences explicitly, observe how they work, and make a model of what our brain must be doing to produce and understand something in exactly the way that they do.

I hope this provides a at least a partial answer to your question — I hope I interpreted it correctly!! Correct me if I've misparsed anything... Sorry to be so long-winded; but really, a truly full answer to the question of "why trees" *does* take a whole syntax course to provide. :)

Kris

We talk about entities and functions. we have been given ways to interpret their semantic meanings. Can these representations be somehow captured and then used to parse through text (or even voice recordings) to assign meaningful understanding of what was said or written? Can we begin to understand the relationships between different lexical entities using this technique?

I have seen what I would call toy attempts like LFG, but these are woefully incomplete and are not computationally practical. The approach we have in class seems much more powerful and parsimonious. Is this also due to the fact that we are merely looking at a subset or are these rules and definitions generalizable to this larger task?

(This seems sort of like a bigger version of Gondy's question of last week, so forgive me for any duplication of discussion).

The problem of creating meaningful understanding independently of a person's head (i.e. in a computer) is far bigger than the problem we've been addressing in this class. We've been seeing how (assumed) meanings (i.e. concepts) are put together to get bigger meanings, and we've been able to draw some pretty robust conclusions about the sorts of meanings that different words and types of words must have. But we've been taking for granted the concepts themselves (the notion of "cattiness" from last week, e.g.)

Now, it would be fairly straightforward (I think -- I don't really know how hard it would be) to make a program that mapped particular words to their particular lambda-expressions (or some equivalent computer-interpretable notation for functions), which is what the model lexicon that we've been developing in this course does, give it phrase-structure rules for building trees (i.e. hook it up to a parser), and then have it combine the subconstituents of the parsed trees in the appropriate ways to calculate the truth-conditions of a whole sentence, according to our six rules of interpretation. And the computer would be able, given a sentence, to tell you what its purely logical entailments are, in the Aristotelian fashion.

However, what's a lot less straightforward, as I intimated last time, is getting the computer to understand the truth-conditions for the lexical items themselves. How can we give a computer a notion of "cattiness", and a notion of "grayness", so that it can combine those two notions and really understand what it means to say "Felix is a gray cat". Given the model we're working with,

the computer *will* understand that "Felix is gray" and "Felix is a cat" are entailed by "Felix is a gray cat", just from the calculation. But, how can we get it to understand that "Felix is a gray cat" entails that "Felix is a quadruped" and "Felix is a mammal" as well as "Felix is not a dog" and "Felix is not made of plastic" and "Felix is from Earth" and "Felix is fairly drab" and "Felix might eat mice" and "Felix can't do arithmetic"? These are all things that we humans know about the notion of "cattiness", but the computer, armed only with [$\lambda x \quad D_e$. x is a cat], doesn't know.

So, what you're asking, is, I think, can we use the present mechanisms to create a model of our understanding of the world? And the answer is, sort of yes, and sort of no.

We could add to our lexicon a representation of all the entailments between (the truth-conditions of) lexical items that we can think of – perhaps, e.g., something like WordNet. Then the computer still wouldn't have a notion of what it means to be gray, or a cat in the experiential sense, but it would have a network of entailment relationships that connected the symbols for "grayness" and "catness" to other symbols for other concepts in an appropriate way. That would be a sort of a model of our understanding of the world. Unfortunately, at this point we've moved away from our nice parsimonious mechanism (well, we already did when we added the parser, but that's not such a big deal). Another possible approach is to assume that words are all made up of semantic "primitives" (the famous proposal that, e.g., "bachelor" really means UNMARRIED MAN, "cake" might mean something like SWEET FOOD, etc.).These semantic primitives would combine via our parsimonious mechanisms and then we'd have a much smaller WordNet type thing expressing entailment relations between the semantic primitives.

And *then*, to make it really work right, we'd have to include a bunch of heuristics like the ones we use everyday when we're assigning parses to sentences as we hear them. For instance, if I've been in a conversation about Wall Street, I know that the string "bank", when I come across it, most likely refers to a financial institution, and not any other possible meaning it has. If I've been talking about royalty, I can parse the string /prIntz/ as "prince" rather than "prints". If someone says, "You're soooooooo smart", (you can infer the intonational contour), I know they really mean, You're not smart at all.

So I imagine that given our current level of understanding, it'd be possible to take the necessary steps to build a pretty good program. It's a HUGE job, evidently. At least, it seems clear to me what its necessary parts would be. I'm almost certainly really naive about all this. Any thoughts?

Erika

#1 (Ed: somewhat like 2nd part of Gondy's, above, in that it's about the psychological reality of the model)

We've been dealing with sentences and their truth-values in a very mathematical way. This is nice because it's consistent and follows a strict set of rules. However, as a human speaker, I'm not sure that I have this kind of consistency. The short version (sans example) of my question is: how true-to-life is this mathematical, logical interpretation of human language? Or, how far can we take this interpretation, given human fallibilities in the area of logic?

It's clear to me that there are many simple (and not-so-simple) sentences that can be accurately interpreted in terms of their truth-values. But humans are fallible. What about the case where I say "I saw the cat"? It seems clear what the truth-value of this sentence is, but what if what I saw was a large Tucson cockroach and I thought it was a cat? It seems that the sentence I uttered is technically false, and yet isn't exactly false. I did, after all, get my meaning across.

*Even worse: what about the 'believe' sentences? In terms of truth conditions, 'people only drank vodka at the party' entails 'no gin was consumed at the party'. So if I say 'I believe that people only drank vodka at the party', that should be logically equivalent to saying 'I believe that no gin was consumed at the party'. But this doesn't *feel* logically equivalent. I could have been mistaken,*

since vodka and gin look so similar, and neglected to consider the possibility that gin might have even been present. But if I believe that the world is such that 'people only drank vodka at the party' is true, I should also believe that the world is such that 'no gin was consumed at the party' is true, since they're the same world. This is a bad example; I had hoped to find one where my human fallibilities would mean that I fail to notice that two statements are logically equivalent, and so I could believe one but not the other. I hope the idea got across anyway.

#2 (ed: see Katie #1 below)

When you talk about words having several different types (like the determiners under the type shifting approach, and like 'and', for instance), does this mean that the words have several different 'definitions'? That is, under the psycholinguistic perspective, would they be different lexical entries?

Heidi: Your #1 question is pretty clear, no worries. And (I think!) it has a two fairly straightforward answers.

The question touches on the deep distinction between Sense and Reference, again, and also on the notion that the meanings of sentences are True and False (and Undefined), rather than some proposition. Sense and Reference first:

Recall the early discussion about the difference between "Hesperus is Hesperus" and "Hesperus is Phosphorus". Although "Phosphorus" and "Hesperus" both refer to the same thing, and hence both sentences are trivially true, Frege noticed (clever fellow) that the second one could be informative, but the first one could not. So, therefore, it's possible for someone to entertain the notion that "Hesperus is Hesperus" (indeed, if they entertain the notion "Hesperus" at all, they must), without entertaining the notion that "Hesperus is Phosphorus". Yet, the real-world referent of the two names is identical, so you'd think that one should logically entail the other.

Frege's answer, as I discussed earlier, is to say that while the reference of the two expressions is the same, their sense is different. Put another way, it's possible for someone to believe that the first is true and the second is false, because they think that the truth-conditions for understanding "Hesperus" and those for understanding "Phosphorus" are distinct. Put yet another way, the statement "Hesperus is Hesperus" and "Hesperus is Phosphorus" ARE synonymous *in the real world*, that is, in the real world, one entails the other. But in the mind of the speaker, they aren't synonymous, and hence can be uttered informationally. (This is a bit specious: the only place that such statements can exist is in the mind of the speaker, but I'm trying to point out that the real world and a speaker's representation of the real world can be mismatched).

So, in the case of "People only drank vodka at the party", if I utter that proposition truthfully, it is true *in the real world* that that entails that no gin was consumed at the party. However, it doesn't follow that I asserted "No gin was consumed at the party". It does logically follow -- so if I hear someone say the former, I can pause and think and then say, hey! That means no gin was consumed! (What about that bottle I lent to Jim that he said he was bringing to that party??) But the point is that the pausing and thinking step is necessary: we don't immediately, upon hearing the sentence about only vodka, at the same time as we comprehend that, comprehend an infinite host of other propositions which are entailed by it. This is, again, because what we understand when we understand the sentence is its Sense, i.e., its truth-conditions, and we understand the speaker to be asserting that the state of affairs entailed by the Sense to hold. *Once we understand what that state of affairs is*, we can then make inferences about other Senses that are true of that state of affairs -- but we don't immediately understand Sense B from Sense A without an intermediate step of pausing to consider what the state of affairs indicated by Sense A is.

These point hold of the believe-context, too. Frege's distinction between Sense and Reference enables an account of the fact that we can believe two mutually contradictory statements at the same time, as long as we don't know that the actual reference of their Sense is mutually contradictory. So, for instance, let's say you've noticed, over the past few days, that there's always a man in sight. On some days, it's a man in a brown coat, and on others , it's a man in a gray coat. For whatever reason, the brown-coated man has aroused your suspicions. In that case, you can say, "I

believe the man in the brown coat is a spy" (or even "The man in the brown coat is a spy", which amounts to the same thing) and also, "I believe that the man in the gray coat is not a spy", and nothing about your beliefs is inconsistent — until you discover, much to your surprise, that the man in the brown coat and the man in the gray coat are the same man. (This is sort of the opposite of your case with the cockroach and the cat: you've conveyed your belief that you saw the cat, even though you really saw the cockroach). This inconsistency is possible because the sense of the expressions differs, although their reference doesn't.

The difference between Sense and Reference is what makes sentences like "My name is Heidi" and "The earth orbits the sun" feel different, even though they both denote the same thing, i.e. True. The notion that the meaning of a sentence is a truth-value is like the notion that the meaning of a name is a person (or thing, in the case of Hesperus and Phosphorus). The starkness of the difference between Sense and Reference is much easier to see in the case of sentences than names, because there's so many things in the world that can be named that it might be tempting to think that the names ARE the thing, or directly denote it. But the Hesperus/Phosphorus case shows us that that's not so. In the case of sentences, though, there's only two truth values, and yet there's an infinite number of propositions that may have them as their reference.

For the answer to number 2 above, see the answer below to Katie's #1.

Katie

#1 (Editor's note: see Erika #2 above)
1. Has anyone ever thought that maybe the things which need to be type-shifted (on that analysis) might actually just be coincidental homonyms? My guess is no because it doesn't seem like it would get you anywhere except in trouble with your theory of the (size of the) lexicon. I'm not sure what other kinds of trouble it might get you in, though.

2. We are doing semantics couched in a certain theory of syntax. How much is dependent upon that? Also, I was wondering if you could say a little about the state of the field of semantics--do most people do it the way we are?

Heidi:
1. The question of whether type-shifted versions of e.g. adjectives, quantifiers or sentential connectives are 1 lexical item with three or four different types made available by a generalizable type-shifting rule, or simply three or four different lexical items, with different selectional co-occurence restrictions, is one we've talked about a bit but not very much. Basically, the conclusion for (at least) adjectives and quantifiers is that it's considerably more parsimonious, theoretically, to posit one lexical item for each, and several generalized type-shifting rules that can apply to the whole class of adjectives and quantifiers to get the correct types in the relevant syntactic configurations.

This conclusion is borne out by our intuitions about these items in their various syntactic positions. Since the "lexical" content of the item (i.e. the part that's the English summary of its truth-conditions) remains the same across all items, no matter how complex their type gets, it seems likely that the generalized-rule approach is the correct one (insofar as type-shifting is the correct approach at all). If they were separate lexical entries, or coincidental homonyms, it might be expected that a quantifier in object position, say, or an adjective in predicate postition could undergo semantic drift independently of the meaning of their counterparts in subject position or modifer position respectively. However, as far as I know, this doesn't happen, and seems fairly unlikely. All intersective adjectives that are modifiers mean the same thing as predicates that they do as modifiers, and vice versa.

With the sentential connectives, however, the case is somewhat trickier. I'm rather inclined to think that for "and" and "or", it might be the case that there are several lexical items rather than a type-shifting rule that applies to connectives generally. My main reason for this is that type shifting in connectives seems to be restricted to "and" and "or"; as we saw in class last time, "but",

"although", and even "not" don't type-shift; they always have to apply to items of type t. We don't save any theoretical ink by saying there's a general rule and then restricting its application to two lexical items — unlike the case of adjectives and GQs above, where type-shifting, if it's appropriate, is appropriate quite generally, applying to all the lexical items of the class.

The fact that the only case where it seems super clear that we need to allow an item to have different types ("and" and "or") is the case where it seems that we only need to allow two lexical items (not a whole category of items) to have different types suggests to me that in fact, type-shifting rules are undesirable in natural language analysis, if we can get away without them. As we've seen, there are good theoretical alternatives to type-shifting for adjectives and GQs (contextually-restricted PM and Quantifier Raising, respectively), and so it seems likely that we **can** get away without them — except in the case of "and" and "or". For "and" and "or", then, the long and roundabout answer to Katie's and Erika's questions might well be yes -- their type-shifted alternatives are actually separately listed lexical items.

2. With respect to the theory-specificity of our particular semantics:

(a) yes, it does matter, for the particular analyses I've (H&K've) been defending, that we're using a particular syntactic theory. For example, in a syntactic theory without movement operations, the QR arguments we've been discussing simply disappear, because there's no possibility of moving the quantifiers (or questions or anything else).

(b) No, it doesn't matter, or matters less than it might appear to. For some of the specific analyses, a particular syntactic theory is necessary. But for much (indeed, so far, most) of the material we've covered, we've been dealing with fairly unsophisticated syntactic trees, of a sort that nearly any syntactic theory will assume represents the basic constituency facts of sentences, and all our discussion of interpretation will carry over into those theories. For any syntactic theory, gradability of adjectives is a question, as is the interpretation of relative clauses, and the proper treatment of quantification. H&K try to give at least the feel of what some alternative theories would feel like (e.g. they outline the type-shifting approach to quantification in object position, and the type-shifting approach to gradable adjectives), and by starting with deSwart, I've tried to emphasize that as well (deSwart works largely within the syntactic assumptions of the HPSG framework, not a Minimalist/GB one like H&K). So certainly many of the issues remain the same across syntactic theories, and many approaches to solving the problems have interesting translations into other theories.

Bob

(1) Is it fair to say that treating "everything" and "nothing" as intransitive quantifiers is a shortcut for a more compositional analysis, in which "thing" is the argument of the quantifier? eg [[everything]] = [[every]]([[thing]])

(2) If we argue, as we have, that verbs are arguments of DPs, does that force us to conclude sentences are really headed by D? Does that really pose any problem at all for syntax, since we see sentences sometimes patterning with DPs (as objects of verbs...I wrote a paper for Andrew last term in which I ended up analysing embedded questions as projections of a null D)?

Heidi:

1. Yep, it's fair. However, it's possible to make an argument that "thing" has lost its truth-conditional force, and the whole lexeme is analysed as a monomorphemic whole. Consider: "thing" in its stand-alone sense doesn't include people. However, if I said, "Everything is magenta!" or "Everything sucks", it's likely that I would mean everything, including people, is magenta (or sucks). So in that respect, it doesn't look like "everything" is really "every thing". On the other hand, it's inappropriate to say after dinner, "I really liked everything!" meaning that you liked the people as well as the food. Perhaps. Or maybe it's not. But you can see the shape the argument would take.

2. That's an interesting question, and one that I'm not quite sure how to answer. If VPs are arguments of the DP subject, that doesn't really mean that the whole S is a projection of D, necessarily. As we've seen, the syntax must be to some degree autonomous (ask me more about this if you like). And so far, we've got a variable semantics, where sometimes the VP is the argument of DP, when DP is headed by a quantifier, and sometimes the DP is the argument of VP (in particular, when the DP is a name). In order to make a uniform semantics-driven story about the projection properties of predicates, we'd have to modify our analysis of names to also make them generalized quantifiers (so, e.g., John would be a function from PowD that gives a function from PowD to truth values, such that the sentence "John likes Mary" is true iff {John} ∩ {likes Mary} ≠ .) This is far from a crazy move; it's been proposed in the literature. However, the payoff in syntactic terms is not clear. If someone can show that in fact any sentence has, in some significant respect, D features (e.g. your analysis of embedded questions), then there might be a reason to pursue it. But the question then becomes, where do inflectional features fit in (tense, aspect, mood)? They are usually assumed to be the head of a sentence in syntactic terms, probably precisely because they cannot occur in DPs.

But my main uncertainty arises from the fact that I'm not sure what it would mean to say that sentences are headed by D. Are they then themselves DPs, as you seem to suggest is the case with embedded questions? Then there must be some explanation for the differences in syntactic distribution of DPs and sentences (case-requirements, for example). If they're not DPs, but rather some bigger thing, then we have to explain why DPs behave as constituents in subject and object and adjunct positions (with, e.g., rampant quantifier raising). It seems like a tricky enterprise, but an interesting one. Perhaps I'll have more insight into it next week.

Travis

I hope you don't mind if I send you my homework question early. I was reading this Russell article this morning and I gather from the intro that it is just the first of two chapters about [[the]]. I guess the chapter I'm reading is about singular the and the second talks about plural the. From yesterday's lecture and what I've read in H & K and Russell it seems like a good idea to treat singular [[the]] as a presuppositional determiner, the presupposition being that there exists only one contextually relevant x s.t. f(x) is true. This brings me to the first part of my question. Will it work to treat plural [[the]] in a similar fashion, with the presupposition that there are at least two x s.t. f(x) is true? It seems to me that this should work. The second part of my question: How do we interpret [[the]] when it takes a mass noun as its complement? It doesn't look like either of the previous presuppositions will work.

Heidi: Hmm. Interesting thought. The answer to the first part of your question is, it seems to me, yes: if you just change the presupposition in [[**the**]] it ought to work out that it combines appropriately with plural Ns. On that analysis, though, you end up with two coincidentally homophonous determiners [[**the**]] which differ only in their presuppositions.

This makes me think that there must be a nicer way to do it, with only one [[**the**]], How about the following:

[[**the**]] = [λf D_{et} and |{x:f(x)=1}| = ±*pl* **.** [λg D_{et} **.** for all x {f(x)=1}, g(x)=1]]

(Note, of course, that this is a GQ version of **the**).
where ±pl = either 1 or more than 1, depending on the value of the number marker in the N head that denotes f(x). What I'm trying to implement is that the [-pl] or [+pl] feature of the N head gets incorporated into the presuppositional restriction on the meaning of **the**. I don't quite know how to inject the ±pl feature into our semantics for N heads, and ensure that it gets passed up to **the**, but the intuition should be clear: there's only one **the** and the presupposition associated with it changes depending on the number feature of the N head. It shouldn't be too tricky to work out an appropriate formalism.

Then, of course, the problem with mass nouns that you point out comes down to deciding what their status with respect to the number marker ±pl is. It seems likely that they're *neither*, i.e. that they're unmarked for number. Then, perhaps, a default interpretation takes over. The presupposition is no longer that the cardinality of the set is ±pl, but rather that there *is* a cardinality to the set — that is, the presupposition on **the** simply becomes an existence presupposition. If you're going to say "the water", then you presuppose that the set {x: x is water} is not empty. The question of what the content of such a set is (it can't be a numbered bunch of things!) can then safely be left for someone else to sort out.

I should add to all this that, with respect to specific analyses of **the**, I'm hardly an expert, so take my musings above with several heaping teaspoons of salt. It's interesting enough that it might be worth some more investigation, though; the place to start is Irene Heim's thesis on definite descriptions (and, presumably, the second chapter of Russell!). If I'd had time, I might have done a little research to help me answer this question more confidently, but so far you'll have to be content with my intuitions.