Jeff has some really nice ideas about propositional structure, its source in nature, its general character, and its relation to natural language syntax. At least to a first approximation, on Jeff's view, the structure of a (humanly graspable) proposition P *is* the structure of some sentence S of a naturally acquirable human language. This suggests an account of why such sentences are good for *expressing* propositions, and how nature arranged for there to *be* propositions that humans can entertain.

If you're interested in these issues, you really should read Jeff's book. (Some of my grad students are doing just this; and it's good for them.) Unsurprisingly, given the nature of these sessions, I have some questions concerning (i) which theses Jeff is most concerned to defend, and (ii) which theses in the vicinity are really plausible. But I want to start by identifying some assumptions that Jeff defends and I adopt, followed by a bit of (very Whiggish) history, by way of highlighting why Jeff's view is interesting and worth discussing—even if you suspect that in the end, grammatical structure *doesn't* map transparently onto propositional structure.

Early in the book, Jeff makes a good case for appeal to structured propositions. Those who abjure propositions entirely, or who appeal only to unstructured propositions, may not be convinced by his succinct presentation of the main arguments that he and others have pressed in more detail elsewhere. But I happily follow Jeff in saying that there are propositions—potential premises and conclusions, which can be entertained and evaluated—and that these "thought contents" exhibit logical *form*. I have nothing against sets of possible worlds. But like Jeff, I think that *structured* contents have their place, and that one philosophical task is to say which place that is.

I'm also sympathetic to Jeff's sympathy for the Tractarian idea that propositions are *facts* of a special sort, and that the facts in question have something to do with languages. Once you can used to the lingo, I think it's perfectly fine to say things like the following (see p.54, among others):

(L) *the proposition that Rebecca swims is a fact consisting of*
Rebecca standing in a certain relation to the property of swimming;
where this relation encodes the instantiation function, and
this relation can be described in terms of how elements of a language
can be related within a sentence of that language.

I'm less confident about the details. If the syntax/semantics of 'Rebecca swims' is where the action is—and that's fine by me—I wonder if the proposition doesn't end up being more like a fact consisting of *there being something* that is *both* the contextually relevant individual *called* 'Rebecca' *and* the *Agent of* suitably many *events of swimming*. (But maybe that's just because Jeff and I have different hunches about where the action in natural language syntax/semantics is. More about this below). Anyway, such details aside, Jeff has me persuaded on two points that seem important: (L) is the right *sort* of thing to say about what structured propositions are, in light of various "unity of the proposition" considerations that are familiar yet too often underappreciated; and if one says this sort thing, one had better be prepared to defend the idea that the relevant languages are *natural* in some interesting sense. Kudos to Jeff for being clear and persuasive on these points.

Correlatively, I *don't* find it odd to raise the question "what holds the constituents of propositions together," and then quickly turn (as Jeff does) to the topic of natural language syntax (see p.29). I'm sure that Jeff is right that many readers will find it surprising, when

his discussion of Frege and Russell segues into his discussion of Chomsky's "minimalist program" for simplifying prior conceptions of (human) transformational grammar. But perhaps perversely, not only was I unsurprised, I endorse this natural-lingustic-turn. Once you adopt (L), or something like it, this is an option that cries out for exploration.

As it happens, I also accept the (Ockham-Fodor) hypothesis that humans and other animals exhibit languages of thought, at least some of which are independent of the distinctly human languages that children naturally acquire. And these "mentaleses" are, I assume, perfectly natural. So it may be that doesn't Jeff really need propositional structure to match the syntax of a *naturally acquired human* language, as opposed to the syntax of *some* natural mental language (that may have no phonology). But it's not implausible that the kind of syntax human children naturally acquire, given a normal course of experience, is a prerequisite for entertaining the range of (systematically related, nonfirstorderizable) thoughts that humans can regularly entertain. So Jeff's idea is surely worth exploring: the grammatical structures made available by the "human language faculty," which combines lexical items in accord with some recursive procedure, just *are* the propositional structures exhibited by the potential premises/conclusions that human thinkers can entertain and logically relate.

Moreover, this idea is far less wild than it would have been a generation ago, given the developments *within* linguistics (conceived, a la Chomsky, as the study of the human language faculty). In the last 25-30 years, significant strides have been made towards reclaiming the traditional idea that grammatical form and logical form are *not badly* out of sync—*pace* Frege, Russell, Wittgenstein and other heroes of the analytic tradition. Indeed, I think of Jeff as a good example of someone trying to relieve this tradition of an odd tension:

it just won't do (at least not in a naturalistic setting) to say that the way to study thought is by studying language, if you *also* say that natural language (grammatical structure) disguises/distorts the "underlying" thoughts (logical structure). It's worth remembering that the period between Frege and Chomsky was *odd*, in that the study of logic got so far *ahead* of the study of grammar. But now that the gap has at least been narrowed, it becomes possible to explore Jeff's "new analytic" twist on an old idea: sentential structure and propositional structure are in sync…but not because spoken sentences reflect thoughts…according to Jeff, there's an important sense in which structured thought contents (Gedanken) reflect the sentences we speak.

Prior to Frege's *Begriffsschrift*, it was tempting to think that *sentences* of a spoken language—like Greek or Latin—exhibit subject-(copula-)predicate structure because *thoughts* exhibit this kind structure. To be sure, there were plenty of complications. Complex thoughts, the traditional idea went, could be built up from atomic thoughts and *syncategorematic* elements; and sentences might not always indicate this complexity transparently. Consider (1).

(1) Jeff smoked, Seth didn't, many don't smoke, and Jim only smokes Camels

If one of the relevant thoughts has the form indicated in (1a),

(1a) JEFF^(IS-ONE-WHO-)SMOKED

then there's nontrivial work to be done on the other conjuncts. Likewise, fallacious inferences like (2) suggested that in particular cases, grammar might not reveal the underlying propositional structures.

(2) That dog is a father, and that dog is yours, so that dog is your father

(Compare: That dog is a mutt and that dog is yours, so that dog is your mutt.) And sentences with relational verbs and quantificational direct objects, as in (3),

(3) No yellow dog that chased a grey squirrel likes every brown rat

were nightmares for medieval logicians. But given their many successes, and absent a better conception of how language was related to thought, it was a reasonable to treat the nightmares as research problems (as opposed to counterexamples).

Fast forward to Frege's powerful reasons for positing logical forms like (3a), which diverges considerably from the apparent subject-predicate form of (3).

(3a) $\sim\exists x\{Yx \,\&\, Dx \,\&\, \exists y[C(x,y) \,\&\, Gy \,\&\, Sy] \,\&\, \forall z[Bx \,\&\, Rx \supset Lxz]\}$

Moreover, Russell argued, changing 'a' in (3) to 'the' magnifies the mismatch.

(3b) …$\exists y[C(x,y) \,\&\, Gy \,\&\, Sy]$…

(3c) …$\exists y[C(x,y) \,\&\, Gy \,\&\, Sy \,\&\, \forall w[Gy \,\&\, Sy \supset (y = w)]]$…

To be sure, Frege and Russell overstated their case a bit. Given restricted quantifiers, as in (3d) and (3e), some of the apparent mismatches vanish.

(3d) $\sim\exists x{:}[Yx \,\&\, Dx \,\&\, \exists y{:}Gy \,\&\, Sy\{C(x, y)\}]\ \{\forall z{:}Bx \,\&\, Rx[Lxz]\}$

(3d) $\acute{\imath}y{:}Gy \,\&\, Sy\{C(x, y)\}]$

Nonetheless, even given the resources of generalized quantifier theory, traditional puzzles concerning (say) quantificational direct objects remain, as illustrated in (4):

(4) [Jim [smoked [every cigarette]]]

(4a) $\forall x{:}Cx[S(j,x)]$

'smoked every cigarette' is a grammatical constituent of (4), with no analog in (4a).

Enter quantifier raising, early advocates of which are here among us. In its original form, the idea was that the "surface structure" of (4)—a "level" of grammatical structure at

which ‘smoked every cigarette’ is indeed a constituent—could be converted into *both* a “phonological form” (PF) that gets *linearized* in a way that makes ‘every cigarette’ comes last, *and* a distinct “logical form” (LF) like (4b),

(4b) [*every cigarette*]^[Jim [smoked [__]]]
|_____|

which can be mapped transparently to (4a). Independent empirical arguments for such transformations made it possible to rediagnose many apparent examples of grammar/logic mismatches as cases in which PFs differed from LFs. This raised an intriguing possibility: while distinct levels of grammatical structure may diverge—say because the (linear) demands of speaking/signing differ from the (hierarchical?) demands of interpreting logical vocabulary in terms of concepts naturally available to humans—perhaps the grammatical structure *germane to interpretation* does not diverge from propositional structure. In short, maybe (humanly accessible) thought-contents exhibit the structures they do because LFs exhibit the structures they do.

This idea becomes especially attractive in the context of more recent attempts to eschew “deep structures” and “surface structures” as significant levels of grammatical structure, and make do with *just* PFs and LFs (PHONs and SEMs) conceived as “instructions to the articulatory/perceptual and conceptual/intentional systems.” This desire for theoretical economy is typically combined with attempts to posit only the sparest recursive apparatus needed to combine lexical items into phrasal units that can be pronounced and understood. Jeff’s thought is that some such “minimalist syntax,” which still provides enough resources to generate structures like (4b), might be just what the philosopher worried about propositional unity needs. This is all *very* tempting. But…

(In what follows, I've been influenced by many discussions with John Collins, and by his paper "Syntax: more or less." Jim Higginbotham carries much of the blame for my imprinting on "neo-Davidsonian" eventish semantics. Chomsky's influence is ubiquitous.)

0. For purposes of minimizing mismatches between grammatical and logical forms, appeals to LFs surely help. And for present purposes, I concede that examples like (3) pose no special difficulty for Jeff. But lots of examples like (1) and (2) remain. Indeed, just about every interesting construction makes for a hard homework assignment if the hypothesis is that the structure of each thought expressed with a sentence S (used literally) is the structure of S. Does natural language *never* mask/distort the structure of the thought expressed? To mention just one kind of worry: does the grammatical adicity of a predicate (the number of arguments it takes) always reflect the adicity of the concept expressed? Think of 'sell', and not just 'father'. My own inclination is to think that just about every sentence—even 'Rebecca swims'—will present potential difficulties. Of course, given how little we know about conceptual adicities, it's hard to make solid case either way. But that raises questions about what sorts of empirical evidence would count for/against Jeff's view, at least in the foreseeable future.

On the other hand, there's lots of interesting work out there that make the potential nightmares less horrific. So if Jeff's idea works for a substantial fragment of cases, that would be enough (in my book) to make other cases look like research problems, not counterexamples. And it's just boring to reply to interesting new proposals with "but what would you say about the constructions confound experts." But even relatively "simple" cases keeping leading me back to three clusters of questions that I'd like to ask Jeff. In the

book, he says many things that are surely relevant to some of these questions. But there, his primary aims included arguing—in my view successfully—for different (and more obviously philosophical) points. So I'd like to invite him here to speak to some of the questions that occurred to me qua "semanticist from a Minimalism-friendly place."

1. Is the thesis really that a sentential LF *is* the structure of the proposition expressed? Or is the important idea that (in some sense still awaiting full characterization) the LF *determines* the corresponding propositional structure?

(1a) What about the apparent "extra junk" in LFs? For example…

(1a*) What about "formal features" like person ($1^{st}/2^{nd}/3^{rd}$), gender, agreement…

Is the idea that these are *all* invisible at LF? This is an attractive speculation. But then one wants a rationale, in terms of a coherent story about the human language faculty, that doesn't make the correlation between visibility and propositional relevance seem like an eerie pre-established harmony. (Potentially, there are questions related to "autonomy of syntax" theses. But that's a hornet's nest I don't want to disturb today.)

And I don't want to push any of this hard; cp. Collins. It's just a query.

(1a)** Is Jeff's proposal compatible with the copy theory of movement, according to which the LF of (4) is roughly as shown in (4c)?

(4c) [every cigarette]^[Jim [smoked [every cigarette]]]

If so, is the idea that some aspect of LF determines which (*one*) copy of 'every cigarette' is a quantifier, with every other copy being interpreted as a variable? Or is a different kind of semantics for quantifier-copies available? Does the (global) property of "highest copy" have

to come in here? Are there *copies* of quantifiers *in propositions*? More generally, which *aspects of* LFs "count" for these purposes?

Once we start counting relational/global properties exhibited by constituents of LFs, which properties of LFs are/aren't germane to propositional structure? *All and only* those properties that count as important for the study of LFs (and the language faculty), without yet worrying about how LFs map to propositions? *All*? *Only*?

(1b) Relatedly, what about the apparent semantic significance of certain *structural* features of LFs? On the assumption that not all thematic roles get expressed by prepositions, even allowing for covert prepositions, we presumably need to say something like this: "being the internal/external argument of a verb" is a way of introducing a thematic relation like "being the Theme/Agent of an event" into the logical form corresponding to a given LF. One can this while still saying that all expressions (of a human language) are generable via some recursive combination operation MERGE, given a stock of lexical items. Generable expressions can exhibit the relevant structural relations. But to the degree that logical forms are "neo-Davidsonian," how does this fit with the idea of MERGE also being the source of propositional facts (via the instantiation relation). It may be that there's a relatively simple answer here, and that I'm too stuck on the 'Rebecca swims' example, with the gloss in terms of Rebecca and the property of swimming. But the details keep eluding me, when I try to think about 'Rebecca swims', 'swims well', and 'swam the race/Channel' all at the same time.

Examples of what I have in mind…and these count, in many books, as elementary

(5) Noam praised Donald

(5a) Noam praised Donald yesterday

(5b) ∃e[Agent(e, Noam) & Past(e) & Praising(e) & Theme(e, Donald) & Yesterday(e)]

(5c) Jim heard Noam praise Donald

(5d) ∃e∃f{[Agent(e, Jim) & Past(e) & Hearing(e) & Theme(e, f)] &
[Agent(f, Noam) & Praising(f) & Theme(f, Donald)]}

2. Pursuing this thought…Does it matter that the semantic significance of MERGE can't be *uniform* on any plausible view, at least not if MERGE is uniformly the operation that generates predicate-argument structures, predicate-adjunct structures, and quantifier-predicate structures?

(Cp. Higginbotham in "On Semantics," which appeals to three modes of composition;
cp. Heim & Kratzer, which also applies to three modes; cp. classic appeals to saturation, type-shifting, and Tarskian quantification over sequence variants;
cp. my own appeals to predicate-conjunction, thematic-shifting, and Tarski.)

Think about Jeff's example, 'Shane runs down the street' (on p54).
[displayed on my handout as (6)/(6'), as in Jeff's text]
And let's concede, as I'm happy to, that (6')
"is not the property of running down the street. It is a certain fact. It is the fact that there is a context c such that relative to c runs* is the semantic value of a word that occurs in a such-and-such sentential relation in such-and-such a way, etc."

Jeff assumes, perhaps for simplicity, that *runs is "presumably the property of running". But what if 'runs' is not semantically correlated with an abstract entity (like a property) but is rather (more sparely) a predicate satisfied by
(or an instruction to fetch a concept satisfied by) e iff e is an event of running.
Then combining 'down' with 'the street' has to differ somehow from combining 'runs' with 'down the street', which has to be importantly different than combining 'Shane' with 'runs down the street'. Whatever the details, it looks like the logical form has to reflect (one way or another) what Jim called "theta linking" for the adjunct, and "theta marking" for 'Shane' (and perhaps 'the street').

(7) Agent(e, Shane) & Runs(e) & Down(e, the street)

But if MERGE generates both adjunction and complementation structures, then while LFs may determine logical forms (in accord with some Baker-style UTAH-constrained algorithm), it seems like logical forms still *differ in kind* from LFs: logical forms seem to have conjuncts that include monadic predicates and thematic predicates. But do LFs themselves exhibit all of these asymmetries? And if so, do propositions exhibit *other* asymmetries (say, in terms of polyadicities) that LFs don't?
SIDEBAR: I can imagine a response that treats MERGE as a complex operation—involving concatenation and labeling (see Hornstein)—and then correlates theta-linking with label-free adjunction (cp. Chametsky), analyzed as mere concatenation. Then everything else, complementation and labeled adjunction, might be treated as theta-marking (cp. Hornstein and Pietroski). But then I think other questions, about covert event variables, will take an

even sharper form. So I'm inclined not to go down this path unless Jeff wants to. And I'd like to end with a broader issues that strikes me as more important, both for Jeff's project and the study of grammar.

3. What is a *sentential* LF?

In terms of grammatical *theory*, 'sentence' has always been an oddball notion. There are many ways of making this point. Here's one: noun phrases are headed by nouns, verb phrases are headed by verbs, and so on. But is there a grammatical category such that sentences are headed by (projections of) lexical items of that category?

If not, are we really sure that the basic recursion system (MERGE or whatever) generates sentences *per se*, as opposed to generating complex *phrases* that are *not* themselves of the semantic type <**t**>? We know, via Tarski, that "sentence-ish expressions" *can* be analyzed as predicates of individuals (or sequences); at least for lots of cases, the familiar <**t**>-based typology is dispensible, raising the minimalist question of whether it should be dispensed with.

So does it matter to Jeff's project if "sentences" are the result of using an LF to create a concept that is still "unsaturated," and then creating a complete thought by (say) existentially closing that concept *in post-linguistic cognition*. Or what if we humans have a mentalese that lets us make subject-predicate thoughts by using an LF to construct a complex monadic concept, and then saturate it with a singular concept that might be raised to salience in a conversational context?

A possible response: some functional type (perhaps Tense or some flavor of complementizer) is also the "sentence head." But since Jeff is admirably sensitive to

Benaceraff-style worries, he'll presumably want to take this route only if there is a principled reason for saying that some *one* of the animals in the current functional zoo (see Cinque) is the sentence head. Put another way… for any (minimalist) syntactic category X, the corresponding instance of "The sentences are the Xs" may be less plausible than a rejection of each such identification of sentences with Xs. If so, the notion of 'sentence' may have its theoretical home (if anywhere) in discussions of how the language faculty interfaces with other components of human minds—including, perhaps, components that traffic in representations with subject-predicate structures. And if that turned out to be so, would it still be true (at least for Jeff's purposes) that LF structure is propositional structure?

It strikes me as practically definitional that structured propositions are of the Fregean type <**t**>, with constituents of other types. But does that mean that on Jeff's view, it must be possible to create expressions of type <**t**> by MERG-ing lexical items? If so, that's an interesting hypothesis about how the human language faculty interfaces with human thought, and not just a familiar idealization to be made when we're not really worried about what sentences are. Now the interesting hypothesis may well be true. But I see nothing in minimalist syntax to require it, and some pressure to deny it. Maybe that tells against minimalist syntax; maybe not. But since Jeff's view highlights the vexed question of what sentences are, it would be great to hear what he has to about this.