Draft #1 of a paper for the Brown Workshop on Direct Compositionality, June 21, 2003
Please check with the author before citing outside of the Workshop.

A layered semantics for utterance modifiers*

Christopher Potts
UC Santa Cruz and UMass, Amherst
potts@linguist.umass.edu

June 4, 2003

1 Frankly,

we should take seriously the *metalinguistic* feel of sentential modifiers such as *frankly*, *confidentially*, and *honestly*, when intonationally isolated as in (1).

(1) a. Frankly, the soup was too salty.
b. The soup was too salty, {frankly/to be honest}.
c. Confidentially, Ed is a werewolf.
d. Honestly, I got food poisoning and so can't come to class.

Bach (1999) classifies these adverbials as *utterance modifiers*, a term that gets right to the heart of their contribution. One also finds 'second-order speech act' (Bach 1999), 'pragmatic adverb' (Bellert 1977), and 'illocutionary adverbial' (Bach and Harnish 1979). The labels all recognize, in various ways, that these linguistic items permit speakers to qualify, restrict, and modify their relationships to the sentences they utter. This paper takes up the task of making good on the intuitions reflected in such labels while at the same time keeping these expressions *in* the grammar. They might be metalinguistic, but this does not mean that we can allow their descriptions to stand as metagrammatical.

Useful tools and techniques for meeting this requirement are found in recent work on *layered* logics and *layered* structures (Gabbay 1999; Blackburn et al. 1993; Blackburn and Meyer-Viol 1997). In the layered setting defined below, we can talk about discourses and the objects they contain; we can also move down a level, to talk about the meanings of individual sentences. The semantics for utterance modifiers is located mainly in the upper layer of the logic and model theory.

It should be said that the layering idea is not really new in linguistic semantics. Montagovian intensional models contain a lower layer of first-order structures and an upper layer of worlds (atomic

---

[*] My thanks also to Polly Jacobson for organizing this workshop. For conversation and advice, I am grateful to Judith Aissen, Luis Alonso-Ovalle, Ash Asudeh, Kent Bach, James Isaacs, Bill Ladusaw, Line Mikkelsen, Geoff Pullum, ...

entities or tuples thereof). The worlds and the first-order structures are systematically related. The *down* operator, $^\vee$, takes us into the first-order models, and the *up* operator, $^\wedge$, lifts us to the upper, intensional layer.

I regard utterance modifiers as one reason to lift yet again. They indicate that we require a structure full of intensional models and a logical language for talking about it. Moreover, we need these models to contain linguistic objects themselves — words, phrases, and sentences. When I say *Frankly, Ed fled*, I indicate that I stand in the frank-utterance relation to the sentence *Ed fled.* Thus, the frank-utterance relation is a set of individual–sentence pairs. It follows that our structures must contain both individuals and sentences.

Once we have the means to analyze utterance modifiers as genuinely modifying utterances, many observations about similar phenomena are brought into sharper focus. For this paper, I favor breadth over depth, because that seems to be the best way to get a feel for the freedoms and limitations of the approach. Thus, I offer partial analyses of the constructions in (2) in addition to an in-depth look at utterance modifiers.

(2) a. Rising intonation on declaratives (section 6)
b. Performative honorifics in Japanese (section 7)
c. Quotation (section 8)
d. Metalinguistic negation (section 9)
e. Textual organizers like *above* and *later* (section 10)

What unifies (2) is the property of manipulating linguistic objects using linguistic means.

The connections with the program of direct compositionality are numerous. The paper shows that we can do without a complex syntax for utterance modifiers as long as we pay close attention to the structures required for an adequate description of their meanings. For example, as discussed below, the present approach delivers a principled account of why syntacticians always place utterance modifiers in the topmost functional projection. The distributional fact that this is intended to capture follows from the kind of function that utterance modifiers denote, making the functional projection itself superfluous. The models are where the action is.

2 The utterance relation

I center most of the analyses in this paper around the utterance relation. Examples like (3) isolate this relation in English and indicate some of the ways it is expressed:

(3) a. Jed: Ed fled
b. Jed entered into the utterance relation with the sentence *Ed fled.*
c. We were unsure where Ed was, but Jed was confident in his assessment of the situation: “Ed fled”.

In (3a), the colon functions to indicate who is responsible for uttering the sentence *Ed fled.* Example (3b) is a somewhat ponderous statement with the same content. The colon appears again with an utterance-based semantics in (3c), though its contribution is more oblique there.

Most semantic theories are limited in their ability to describe systematically even simple cases of this form. Typically, we capture the notions of utterer and utterance as part of the interpretation procedure. For (3), we would fix Jed as the speaker index on the interpretation function:

(4) $[\![\mathbf{flee}(\mathbf{ed})]\!]^{\mathcal{M},Jed}$

We should regard this speaker index not only as determining the meaning of *I* but also as an indication that interpretation is relative to Jed's belief state. As Schlenker (2003) says, "the set of worlds compatible with the speaker's belief at the time of utterance is *by definition the Common Ground itself*" (p. 87; emphasis in the original). This can be said to get at the heart of the paraphrase in (3b).

On this approach, *say* and *utter* have a much different sort of semantics. Whereas *say* denotes a function in our model, *utter* (and the colons in (3a, c)) has a more abstract, metalogical meaning, one that is wired into the interpretation brackets and our notion of which models count as admissible.

The differences are suspicious; one would expect to find the denotations of both *say* and *utter in* the models for the semantic theory. More importantly, the treatment of *utter* is not flexible enough to provide the basis for utterance modifiers, performative honorifics, and the other expressions listed in (2). Indeed, we lack the means even to capture the semantics of sentences like (5).

(5) a. When Jed said "Ed Fled", I was out by the shed.

b. There is a past time $\mathfrak{t}$ such that Jed uttered the sentence "Ed Fled" at $\mathfrak{t}$ and I was in the shed at $\mathfrak{t}$.

In order to describe such examples in anything like the above terms, we would need a highly flexible theory of the interpretation function. We would need a way to shift the speaker index in mid-discourse (for quotation), and we would need some way to keep track of the interpretations themselves, so that we could refer back to them later. In short, we would need a logic — that is, a grammar — of the interpretation brackets. In many ways, this is where the present paper heads.

The relevant notion is found already in the literature on situation semantics: we must distinguish utterance situations from described situations (Barwise and Perry 1983; Fenstad 1997:§2.2.2). If we develop the tools for making this distinction, then we will be able to analyze (3a) as having two components: the proposition that Jed uttered the sentence *Ed fled* (the utterance situation) as well as the proposition that Ed fled (the described situation).

In (6), I offer a set of tools that permit us to bring utterance situations into the grammar.

(6) a. $\mathcal{L}_U$ is a lambda calculus, the Upper logic.

b. D_e is a set of entities, containing a distinguished subset $A = \{a_1, a_2, \ldots\}$ of discourse participants. The semantic type for D_e is e_U (the Upper entity type).

c. $D_p = \{\mathcal{P}_1, \mathcal{P}_2, \ldots\}$ is a set of well-formed linguistic phrases. The semantic type for D_p is p_U.

d. $D_t = \{1, 0\}$ is a set of truth values. The type for D_t is t_U.

e. $V_{\mathcal{U}}$ is a valuation function, taking constants of $\mathcal{L}_U$ to functions formed from objects in $D_e \cup D_p \cup \{0, 1\}$, constrained so that if α is a term of type σ, then $V_{\mathcal{U}}(\alpha) \in D_\sigma$.

In a sense, these pieces just form a regular extensional model. The twist is the set D_p of linguistic phrases. They are treated as atomic entities in these models — they have the same status as names (unanalyzed, nonfunctional rigid designators). Let's assume that the members of D_p are pairs of trees: for all $\mathcal{P} \in D_p$, we have $\mathcal{P} = \langle \mathcal{T}^s, \mathcal{T}^m \rangle$, where $\mathcal{T}^s$ is a syntactic structure and $\mathcal{T}^m$ is a semantic parsetree. This view of D_p can be refined quite easily; in the discussion of metalinguistic negation in section 9, I add a phonological component, thereby making linguistic objects into triples, as in most work on categorial grammar (e.g., Jacobson 1999, 2000).

Clause (6c) limits the set D_p to only objects that are that 'well-formed'. In this context, 'well-formed' is relative to the logic for the semantic parsetrees. I call this logic $\mathcal{L}_L$. It is for talking about the Lower parts of the models that we end up with. I assume that $\mathcal{L}_L$ is a typed lambda calculus of the sort generally employed in linguistic semantics. Its terms lack raised corner brackets, and its types do not have subscripts at all. These devices distinguish its syntactic elements from those of $\mathcal{L}_U$.

A parsetree $\mathcal{T}^m$, and in turn the whole phrase containing it, is well formed if and only if all the nodes of $\mathcal{T}^m$ are labelled with terms of $\mathcal{L}_L$ and, moreover, each local subtree has the following form:

(7)

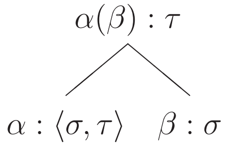

This is just the usual rule of functional application. It supplements the restriction that all nodes have terms of $\mathcal{L}_L$ on them by ensuring that the terms on the leaves fully determine the terms on all nodes above them.

It seems unlikely that functional application alone can provide a theory of semantic composition even for the uses to which we need to put a logic like $\mathcal{L}_L$ (Heim and Kratzer 1998; Chung and Ladusaw 2003; Potts 2003c). But it is the only rule for $\mathcal{L}_L$'s parsetrees that I require for this paper.

I note one other fiction: obviously, semantic constraints are not the only ones that contribute to well-formedness. We must also exclude from D_p phrases that are syntactically and phonologically ill-formed. So the biconditional I used above ($\mathcal{P}$ is in D_p iff $\mathcal{P}$ is well-formed in $\mathcal{L}_L$), should be weakened to an implication ('iff' to 'only if'). But I stick with the stronger statement to avoid avoidable indeterminacy about what is in D_p. In order to define syntactic and phonological constraints with the same level of rigor at which I define semantics constraints in this paper, we would require description logics for those parts of the theory. The result would be ten pages of definitions and a serious loss of clarity.

The logic $\mathcal{L}_U$ provides the means for naming members of D_p. For this, I adopt the convention that the $\mathcal{L}_U$ constant associated with a phrase $\mathcal{P} \in D_p$ is always the surface realization of $\mathcal{P}$ with raised corner brackets around it. For example:

(8) 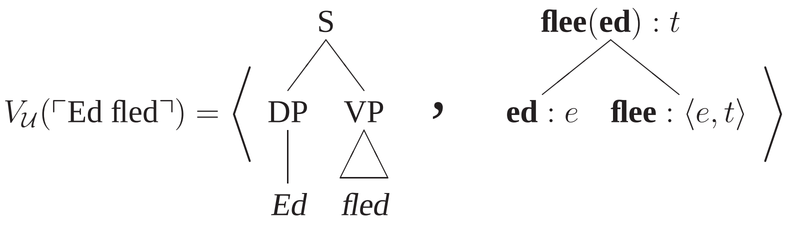

I will mostly be informal about the nature of the syntactic structures. But the semantic parsetrees (depicted on the right) are a major concern of the present approach. Note that the terms that label nodes in these structures lack raised corner brackets. This is because they are drawn from the logic $\mathcal{L}_L$. In essence, $\mathcal{L}_L$ is for semantics as usual; $\mathcal{L}_U$ is for talking about things that might be dubbed 'pragmatic'.

It seems to me that one can usefully think of ⌜Ed fled⌝ as the name of the sentence (pair of trees) in (8). One might worry, though, that this constitutes a philosophical blunder; Searle (1969) warns against such glosses:

(9) "It is generally claimed by philosophers and logicians that in a case like 2 [= *"Socrates" has eight letters*] the word "Socrates" does not occur at all, rather a completely new word occurs, the proper name of that word. [...] I find this account absurd." (Searle 1969:74)

This seems to be an injunction against just the sort of interpretation procedure that $\mathcal{L}_U$ provides. The view that Searle recommends in its place is described in (10).

(10) "But how shall we characterize the utterance of the first word in 2? The answer is quite simple: a word is here uttered but not in its normal use. The word itself is *presented* and then talked about, and that it is to be taken as presented rather than used conventionally to refer is indicated by the quotes [...]" (Searle 1969:74–75)

I believe that we can accept this view without changing the logic; we can regard ⌜Ed fled⌝ as the presentation of the sentence *Ed fled*. If there are differences between naming and presentation in this area, then the logic $\mathcal{L}_U$ is not sensitive to them.

We are now positioned to define the two-place utterance relation:

(11) a. $\ulcorner\mathbf{utter}\urcorner : \langle p_U, \langle e_U, t_U\rangle\rangle$

b. $\ulcorner\mathbf{utter}\urcorner(\ulcorner S\urcorner)(\ulcorner\mathbf{b}\urcorner)$ is true just in case the individual named with $\ulcorner\mathbf{b}\urcorner$ stands in the utterance relation to the sentence named by $\ulcorner S\urcorner$.

This meaning provides the basis for an analysis of examples such as those in (3). In (12b), I offer a semantic parsetree for (12a). If we regard $\ulcorner\mathbf{utter}\urcorner$ as the translation of the colon in examples like (3a), then this structure is appropriate for that example as well.

(12) a. Jed uttered *Ed fled*

b. 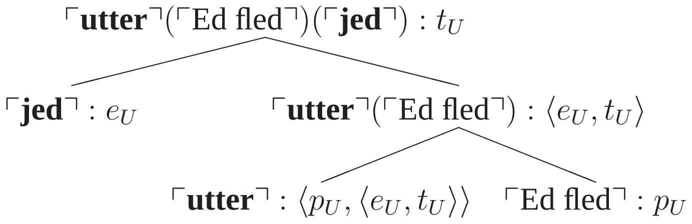

3 The lower layer

The logic presented in section 2 is deficient in one important respect: it treats sentences as though they were atomic objects. Both $\ulcorner$Ali ran the marathon$\urcorner$ and $\ulcorner$Ed fled$\urcorner$ are simply constants, denoting entities in D_p. We need some way to interpret members of D_p. That is, we need some way of doing the usual things we do in semantics. It is only in this way that we'll truly understand what Jed said when he uttered *Ed fled.*

To achieve this added dimension to the logic and its models, I appeal to layering techniques. The new pieces are given in (13).

(13) a. $\mathcal{L}_L$ is a lambda calculus that is distinct from $\mathcal{L}_U$. ($\mathcal{L}_L$ terms lack raised square brackets.)

b. $\mathfrak{M} = \{\mathcal{M}_1, \mathcal{M}_2, \ldots\}$ is a set of models. Each $\mathcal{M}_i$ is a triple $(D_e, \{0,1\}, V_i)$, where D_e is the domain of entities defined in (6b) and V_i is a valuation function, taking formulae of $\mathcal{L}_L$ to functions formed from $D_e \cup \{0,1\}$.

c. $\hbar$ is a function that takes each $a_i \in A$ to the model $\mathcal{M}_i \in \mathfrak{M}$, where $\mathcal{M}_i$ can be viewed as the world-view of a_i.

The set $\mathfrak{M}$ provides us with world-views for our discourse participants. The function $\hbar$ assigns each discourse participant a_i to a model $\mathcal{M}_i$. We can allow that these models might be partial interpretations — that is, they might return 'undefined' for certain expressions of the logic.

I should mention two fictions of this definition: all the models in $\mathfrak{M}$ share the same domain of entities D_e (which is also the entity domain for $\mathcal{L}_U$'s structures); and all the discourse participants "speak" (form meanings from) the same logic $\mathcal{L}_L$. Thus, the only point of variation among the members of $\mathfrak{M}$ is their valuation functions. Meager though this might seem, it allows for quite a lot of variation in the discourse participants' world-views. We can, for instance, define $\mathfrak{M}$ as the pair of models $\mathcal{M}_{lois}$ and $\mathcal{M}_{ali}$ in (14), which are defined for a version of $\mathcal{L}_L$ with just three constants and one predicate. (In the representations of these models and those to come, I use set-notation rather than functional notation. This is merely a presentational choice.)

(14)

$$\begin{array}{|rcl|}
\hline
D_e & = & \{clark, lois, ali\} \\
D_t & = & \{0,1\} \\
\hline
V_{lois}(\textbf{superman}) & = & ali \\
V_{lois}(\textbf{ali}) & = & ali \\
V_{lois}(\textbf{clark}) & = & clark \\
V_{lois}(\textbf{hero}) & = & \{ali\} \\
\hline
\end{array}
\qquad
\begin{array}{|rcl|}
\hline
D_e & = & \{clark, lois, ali\} \\
D_t & = & \{0,1\} \\
\hline
V_{ali}(\textbf{superman}) & = & clark \\
V_{ali}(\textbf{ali}) & = & ali \\
V_{ali}(\textbf{clark}) & = & clark \\
V_{ali}(\textbf{hero}) & = & \{clark\} \\
\hline
\end{array}$$

For whatever reason, Lois thinks that **superman** names Ali (ali), whereas Ali thinks (knows) that **superman** names Clark ($clark$). We can define a common ground model for these two — the valuation V_{cg} such that $V_{cg}(\varphi)$ is defined for a term φ of $\mathcal{L}_L$ only if $V_{lois}(\varphi) = V_{ali}(\varphi)$, and such that $V_{cg}(\varphi) = V_{lois}(\varphi)$ for all φ of $\mathcal{L}_L$. For this pair, we have (15).

(15)

$$\begin{array}{|rcl|}\hline D_e & = & \{clark, lois, ali\} \\ D_t & = & \{0,1\} \\ \hline V_{cg}(\textbf{superman}) & = & \text{undefined} \\ V_{cg}(\textbf{ali}) & = & ali \\ V_{cg}(\textbf{clark}) & = & clark \\ V_{cg}(\textbf{hero}) & = & \text{undefined} \\ \hline \end{array}$$

It is probably evident at this point that the first-order models we associate with each discourse participant have the same information content as the worlds in an intensional model. Indeed, we could redefine the lower parts of these structures to be the usual sort of intensional model, and we could represent belief states with sets of valuations (equivalently, sets of worlds; Gunlogson 2001) rather than as individual valuations. This would provide another area in which to locate variation if we needed further ways of representing discourse participants' varying belief states.

For the purposes of this paper, the structures suggested by (13) provide us with the right amount of information, and they package it usefully. The advantages of this approach will, I hope, become clear in the next section, where I define the interpretation function.

4 Interpretation

I've essentially defined two logics and two classes of models. They are brought together by the interpretation function, which is sensitive to the sort of object it is applied to. In essence, we can feed it either expressions of $\mathcal{L}_U$ or expressions of $\mathcal{L}_L$.

First, I bring together the above ideas in a definition of *discourse structure*:

(16) A *discourse structure* for $\mathcal{L}_U$ and $\mathcal{L}_L$ is a tuple $\mathcal{D} = (D_e, D_p, \{0,1\}, V_{\mathcal{U}}, \mathfrak{M}, \hbar)$, where

a. D_e is a set of entities, containing a distinguished subset $A = \{a_1, a_2, \ldots\}$ of discourse participants. The semantic type for D_e is e_U.

b. $D_p = \{\mathcal{P}_1, \mathcal{P}_2, \ldots\}$ is a set of linguistic phrases. The semantic type for D_p is p_U. Each $\mathcal{P} \in D_p$ is a pair $\langle \mathcal{T}^s, \mathcal{T}^m \rangle$, where $\mathcal{T}^s$ is a syntactic structure and $\mathcal{T}^m$ is a semantic parsetree. Each $\mathcal{P} \in D_p$ is such that its parsetree is decorated entirely with terms of $\mathcal{L}_L$ according to the condition in (7).

c. $D_t = \{1, 0\}$ is a set of truth values. The semantic type for D_t is t_U.

d. $V_{\mathcal{U}}$ is a valuation function, taking constants of $\mathcal{L}_U$ to functions formed from objects in $D_e \cup D_p \cup \{0,1\}$, constrained so that if α is a term of type σ, then $V_{\mathcal{U}}(\alpha) \in D_\sigma$.

e. $\mathfrak{M} = \{\mathcal{M}_1, \mathcal{M}_2, \ldots\}$ is a set of models. Each $\mathcal{M}_i$ is a triple $(D_e, \{0,1\}, V_i)$, where D_e is the domain of entities defined in (16a) and V_i is a valuation function, taking formulae of $\mathcal{L}_L$ to functions formed from $D_e \cup \{0,1\}$.

f. $\hbar$ is a function that takes each $a_i \in A$ to the model $\mathcal{M}_i \in \mathfrak{M}$, where $\mathcal{M}_i$ can be viewed as the world-view of a_i.

I henceforth extend the basic valuations $V_{\mathcal{U}}$ and V_i so that they take complex formulae as arguments. For instance, $V_{\mathcal{U}}(\ulcorner\textbf{utter}\urcorner(\ulcorner\text{Ed fled}\urcorner))$ abbreviates $(V_{\mathcal{U}}(\ulcorner\textbf{utter}\urcorner))(V_{\mathcal{U}}(\ulcorner\text{Ed fled}\urcorner))$.

The definition appears rather elaborate. I assure readers that it is merely long. Recall the pair of extensional models defined in (14). Let's use them as a basis for a discourse structure $\mathcal{D}$. That is we set $\mathfrak{M}$ equal to $\{\mathcal{M}_{lois}, \mathcal{M}_{ali}\}$. The set of discourse participants A is $\{ali, lois\}$. The world-view models fix the set D_e as $\{clark, lois, ali\}$. And the mapping $\hbar$ works in the obvious way, taking each discourse participant to the model with her "subscripted" to it. The set of sentences in D_p is meager, fixed as it is by the rather meager version of $\mathcal{L}_L$ we are working with: $\{\ulcorner\text{ali}\urcorner, \ulcorner\text{is}\urcorner, \ulcorner\text{ali is superman}\urcorner, \ldots\}$. This picture is given in full in (17).

(17)

$$\begin{aligned}
A &= \{ali, lois\}\\
D_e &= A \cup \{clark\}\\
D_t &= \{0,1\}\\
\hbar(ali) &= \mathcal{M}_{ali}\\
\hbar(lois) &= \mathcal{M}_{lois}\\
D_p &= \left\{\begin{array}{ll}\ulcorner\text{Ali}\urcorner, & \ulcorner\text{Lois}\urcorner,\\ \ulcorner\text{Clark}\urcorner, & \ulcorner\text{is a hero}\urcorner,\\ \ulcorner\text{Ali is a hero}\urcorner, & \\ \ulcorner\text{Lois is a hero}\urcorner, & \\ \ulcorner\text{Clark is a hero}\urcorner & \end{array}\right\}\\
V_{\mathcal{U}}(\ulcorner\textbf{left-of}\urcorner) &= \{\langle lois, ali\rangle\}\\
V_{\mathcal{U}}(\ulcorner\textbf{pointing}\urcorner) &= \{lois\}\\
V_{\mathcal{U}}(\ulcorner\textbf{utter}\urcorner) &= \{\langle lois, \ulcorner\text{ali is a hero}\urcorner\rangle\}
\end{aligned}$$

$\mathcal{M}_{ali} =$

$$\boxed{\begin{aligned}
D_e &= \{clark, lois, ali\}\\
D_t &= \{0,1\}\\
\hline
V_{lois}(\textbf{superman}) &= ali\\
V_{lois}(\textbf{ali}) &= ali\\
V_{lois}(\textbf{clark}) &= clark\\
V_{lois}(\textbf{hero}) &= \{ali\}
\end{aligned}}$$

$\mathcal{M}_{lois} =$

$$\boxed{\begin{aligned}
D_e &= \{clark, lois, ali\}\\
D_t &= \{0,1\}\\
\hline
V_{ali}(\textbf{superman}) &= clark\\
V_{ali}(\textbf{ali}) &= ali\\
V_{ali}(\textbf{clark}) &= clark\\
V_{ali}(\textbf{hero}) &= \{clark\}
\end{aligned}}$$

I filled out the structure with a small valuation function $V_{\mathcal{U}}$, defined for $\mathcal{L}_U$ with just two predicates. The structure models a discourse situation in which Lois is standing to the left of Ali, pointing, and uttering the sentence *Ali is a hero.* Lois believes that Ali is superman. Ali knows what's actually the case, though: she takes **superman** to Clark (who is not a discourse participant). The diagram is divided up conceptually: the left cell contains the basic elements of the upper model. The right cell contains the lower part of the structure — the two world-view models in the range of $\hbar$.

What we require now is a unified way of moving from our logics to these structures. This is the job of the interpretation function, which I define in such a way that it has both $\mathcal{L}_U$ and $\mathcal{L}_L$ terms in its range. Specifically, the interpretation function $[\![\cdot]\!]^{\mathcal{D},s,a}$ for a discourse structure $\mathcal{D}$ is relativized to a speaker s and an addressee a, both members of A. It is defined, in (18), as a function that takes pairs of trees to model-theoretic denotations.

(18)
$$[\![\langle \mathcal{T}^s, \mathcal{T}^m\rangle]\!]^{\mathcal{D},s,a} = \begin{cases} V_{\mathcal{U}}(\varphi) & \text{if } \varphi \text{ is a formula of } \mathcal{L}_U \text{ and}\\ & \varphi \text{ is the term that labels the root node of } \mathcal{T}_m\\ V_s(\varphi) & \text{if } \varphi \text{ is a formula of } \mathcal{L}_L \text{ and}\\ & \varphi \text{ is the term that labels the root node of } \mathcal{T}_m \end{cases}$$

To keep things simple, I assume that a parsetree's denotation is determined entirely by the term that decorates its root node (cf. Potts 2003c).

In brief, $[\![\langle \mathcal{T}^s, \mathcal{T}^m \rangle]\!]^{\mathcal{D},s,a}$ provides the $\mathcal{L}_U$ interpretation of $\mathcal{T}^m$ if its root node is decorated with a term of $\mathcal{L}_U$, whereas $[\![\langle \mathcal{T}^s, \mathcal{T}^m \rangle]\!]^{\mathcal{D},s,a}$ provides the $\mathcal{L}_L$ interpretation inside the speaker's model $\mathcal{M}_s$ if its root node is decorated with a term of $\mathcal{L}_L$.

In the remainder of this paper, I refer to the logic that we get when we combine $\mathcal{L}_U$ and $\mathcal{L}_L$ using the above interpretation brackets as $\mathcal{L}_U{+}\mathcal{L}_L$, the layering of $\mathcal{L}_U$ over $\mathcal{L}_L$.

5 Utterance modifiers

The above sections, taken together, provide us with the tools we need for analyzing utterance modifiers in a way that remains true to their name. The guiding idea is that utterance modifiers contribute something to the specification of how the sentence is uttered — they are properly conceived of as restricting the utterance relation:

(19) a. Frankly, Ed fled.
b. ≈I utter this sentence frankly: *Ed fled*

(20) a. Honestly, Ed fled.
b. ≈I utter this sentence honestly: *Ed fled*

One might at first be suspicious of the appearance of *utter* in the paraphrases. But Bellert (1977) offers rather direct evidence for a role for this morpheme in the semantics of utterance modifiers and similar expressions. She writes:

(21) "All pragmatic adverbs cooccur with the participle *speaking*, which is implicit in the sentences containing those adverbs." (Bellert 1977:349)

This generalization is based on contrasting paradigms like the following:

(22)	a.	frankly speaking	speaking frankly
	b.	confidentially speaking	speaking confidentially
	c.	honestly speaking	speaking honestly

(23)	a.	*possibly speaking	*speaking possibly
	b.	*obviously speaking	*speaking obviously
	c.	*amazingly speaking	*speaking amazingly

We can regard *speaking* as the (optional) syntactic realization of the term $\ulcorner\textbf{utter}\urcorner$, defined in (11) as a function from sentences to predicate-level functions. Working from the premise that we want the translation of *frankly* and the like to apply to $\ulcorner\textbf{utter}\urcorner$ (type $\langle p_U, \langle e_U, t_U\rangle\rangle$) we arrive at the conclusion that utterance modifiers should be functions from objects of type $\langle p_U, \langle e_U, t_U\rangle\rangle$ into some other kind of function.

What is the result of combination with $\ulcorner\textbf{utter}\urcorner$? To properly describe the distribution of these modifiers, we seem to require a more articulated view of D_p than is afforded us by the definition in (16b). Bach (1999) observes that utterance modifiers do not usually embed; the examples in (24) are Bach's (32)–(34) (p. 358).

(24) a. $^{\#}$Bill said that, confidentially, Al's wife is having an affair.
b. Man to man, your wife is having an affair.
c. $^{\#}$Bill said to Al that, man to man, his wife was having an affair.
d. In case you're interested, your wife is having an affair.
e. *Bill said to Al that, in case he was interested, Al's wife was having an affair.

This indicates that the meaning of, for example, $\ulcorner\textbf{frankly}\urcorner(\ulcorner\textbf{utter}\urcorner)$ should be a function that takes full, main-clause sentences into the set of truth values. It is not precise enough to say $\ulcorner\textbf{frankly}\urcorner(\ulcorner\textbf{utter}\urcorner)$ is merely a function in $\langle p_U, t_U\rangle$. The set D_p has a rich internal structure; we need some way to access that structure in $\mathcal{L}_U$. We can do this by dividing the set of expressions in p_U into the following subtypes:

(25) a. r_U for $r$oot clauses
b. s_U for $s$ubordinate clauses
c. $p_U = r_U \cup s_U$

This permits us to limit the range of $\ulcorner\textbf{frankly}\urcorner(\ulcorner\textbf{utter}\urcorner)$ to members of D_p that are root clauses: we say that $\ulcorner\textbf{frankly}\urcorner(\ulcorner\textbf{utter}\urcorner)$ is of type $\langle r_U, t_U\rangle$. The details for *frankly* are in (26).

(26) a. $[\![\ulcorner\textbf{the-speaker}\urcorner]\!]^{\mathcal{D},s,a} = s$
b. $\ulcorner\textbf{frankly}\urcorner : \langle\langle p_U, \langle e_U, t_U\rangle\rangle, \langle r_U, \langle e_U, t_U\rangle\rangle\rangle$
c. *frankly* $\rightsquigarrow$
$\lambda\ulcorner\mathcal{U}\urcorner\lambda\ulcorner S\urcorner.\ulcorner\textbf{frankly}\urcorner(\ulcorner\mathcal{U}\urcorner)(\ulcorner S\urcorner)(\ulcorner\textbf{the-speaker}\urcorner) : \langle p_U, \langle e_U, t_U\rangle\rangle, \langle r_U, t_U\rangle$
d.

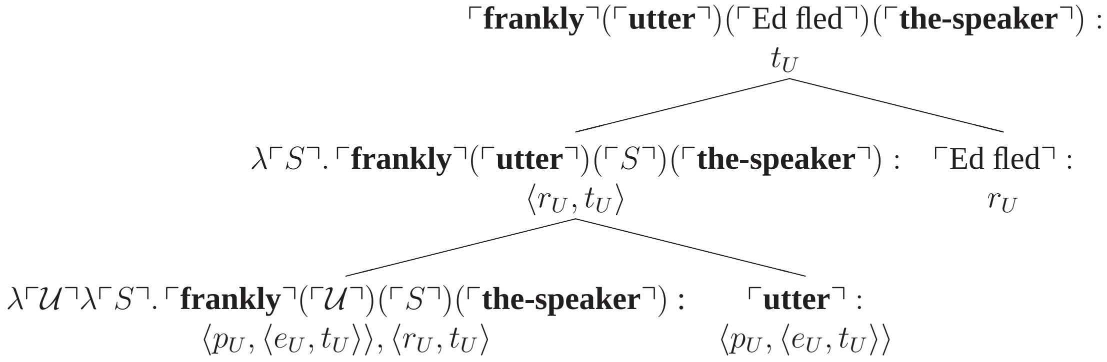

The truth conditions are then given in the expected manner:

(27) $[\![\ulcorner\textbf{frankly}\urcorner(\ulcorner\textbf{utter}\urcorner)(\ulcorner\text{Ed fled}\urcorner)(\ulcorner\textbf{the-speaker}\urcorner)]\!]^{\mathcal{D},s,a}$
$= 1 \text{ iff } \langle s, \textit{Ed fled}\rangle \in \big([\![\ulcorner\textbf{frankly}\urcorner]\!]^{\mathcal{D},s,a}([\![\ulcorner\textbf{utter}\urcorner]\!]^{\mathcal{D},s,a})\big)$

The interpretation of $\ulcorner$Ed fled$\urcorner$ works as follows: first, $[\![\cdot]\!]^{\mathcal{D},s,a}$ applies to the constant itself:

(28) $[\![\ulcorner\text{Ed fled}\urcorner]\!]^{\mathcal{D},s,a} =$

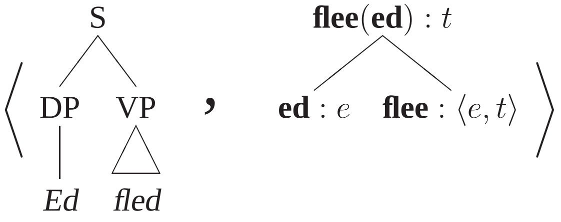

The output is itself a possible input to $[\![\cdot]\!]^{\mathcal{D},s,a}$. Because the root node of the output's parsetree is decorated with a term of $\mathcal{L}_L$, we interpret the sentence inside the speaker's model:

(29)

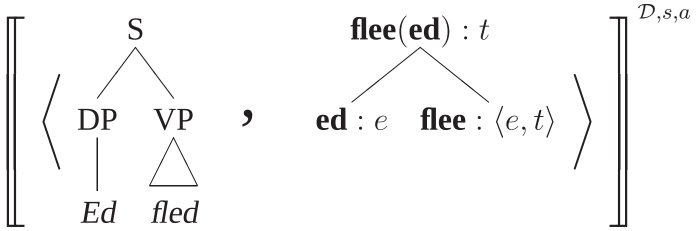

$= V_s\,(\textbf{flee}(\textbf{ed}))$

So, in sum, *Frankly, Ed Fled*, as uttered by the individual s, means that s utters the sentence *Ed Fled* frankly, and that he proposes (offers into the common ground) the proposition that Ed fled.

On this treatment, we achieve the syntactic unembeddability of utterance modifiers in their type-specification: we interpret r_U as the set of main-clause objects in D_p, and we specify that utterance modifiers are functions from these main-clauses to truth values. This is, in effect, a semantic statement of the idea that utterance modifiers are housed syntactically in the highest available functional projection in the syntactic structures. On the present treatment, we can allow them to roam free syntactically, as long as we enforce some conditions on how the syntactic and semantic structures relate to each other. The actual work of ensuring that they appear only in the leftmost (highest) syntactic position is taken up by the type theory for $\mathcal{L}_U$.

I note, though, that this might not be the best solution. The potential difficulty is that utterance modifiers seem to be embeddable when the matrix subject matches the speaker index, as in (30).

(30) a. I think that, just between you and me, Ed fled with the winnings.
b. I feel that, quite frankly, Ed is not trustworthy.
c. I swear that, man to man, I did not sell your chihuahua into slavery.

One senses that these cases work because of the close semantic similarity between the sentence *I think that S* as uttered by a speaker d, and the sentence *S* as uttered by that same speaker d. In general, this suggests that it might be too strong to build the unembeddability condition into the types. It might be preferable to say that examples such as (24) are marked because they are invariably *false*. We can achieve this result with the following conditions on admissible models:

(31) a. Assume that a speaker s utters a phrase $\mathcal{P}$.

b. Assume that $\mathcal{P}$ has subconstituents $P_1, \ldots, P_n$.

c. It does *not* follow that s uttered any of $P_1, \ldots P_n$ that is not equal to $\mathcal{P}$.

This is a rather weak condition. It is perhaps best regarded as a warning against potentially dangerous meaning postulates. It is certainly easy enough to support independently of utterance modifiers: it would be disastrous if a speaker's uttering *It's false that pigs fly* entailed that he had uttered *pigs fly*.

One thing seems certain: the root/subordinate distinction is not the only one that exists in D_p. It seems that we also need to distinguish assertive from declarative root clauses. When utterance modifiers take interrogative sentences as their arguments, they modify not the question, but rather the expected answer, an observation that I owe to Bill Ladusaw.

(32) a. Confidentially, is Al having an affair?

$\approx$I promise to keep the answer to *Is Al having an affair?* a secret.

b. Honestly, has Ed fled?

$\approx$Provide me with an honest answer to the question *Has Ed fled?*

The most important step is to divide up the set denoted by r_U into two subsets: a set $\{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$ of assertions and a set $\{\mathcal{Q}_1, \ldots, \mathcal{Q}_n\}$ of interrogatives. Let the type of assertions be a_U and the type of interrogatives be q_U. This yields the following structure for the type p_U:

(33) a. $r_U = a_U \cup q_U$

b. $p_U = r_U \cup s_U$

We then say that the translation of an utterance modifier is partially dependent upon what kind of complement it has after combining with $\ulcorner$**utter**$\urcorner$. In (34), I provide a meaning for *honestly* that achieves this (here, $\ulcorner Q \urcorner$ is a variable of type q).

(34) a. *honestly* $\rightsquigarrow$

$\lambda\ulcorner\mathcal{U}\urcorner\lambda\ulcorner Q\urcorner.\ulcorner\textbf{honestly}\urcorner(\ulcorner\mathcal{U}\urcorner)(\ulcorner\textbf{answer}\urcorner(\ulcorner Q\urcorner))(\ulcorner\textbf{the-addressee}\urcorner) : \langle\langle p_U, \langle e_U, t_U\rangle\rangle, \langle q_U, t_U\rangle\rangle$

b. *honestly speaking* $\rightsquigarrow$

$\lambda\ulcorner Q\urcorner.\ulcorner\textbf{honestly}\urcorner(\ulcorner\textbf{utter}\urcorner)(\ulcorner\textbf{answer}\urcorner(\ulcorner Q\urcorner))(\ulcorner\textbf{the-addressee}\urcorner) : \langle q_U, t\rangle$

The term **answer** should take an interrogative meaning to its answer (in the world of evaluation). There is an imperative element to the nature of *honestly* (and other utterance modifiers) in this guise. Such *hidden imperatives* are the subject of ongoing work by James Isaacs and myself.

6 Rising intonation on declaratives

Gunlogson (2001) develops a rich, detailed account of the semantics and pragmatics of, among other things, rising intonation on declaratives. I offer rising declarative in (35), in which the closing question mark is a signal that the sentence has a rising intonation.

(35) Ed dated a gorilla?

In brief, Gunlogson concludes that rising intonation on a declarative has the effect of attributing its propositional content to the hearer. So, with (35), a speaker commits his addressee to the proposition that Ed dated a gorilla. This sounds at first like a rather amazing performative. But one cannot in fact use rising declaratives to enforce beliefs and commitments upon others, at least not felicitously. Gunlogson shows that such intonation is subject to a felicity condition: the addressee must have given some prior indication of commitment to the propositional content. The speaker is generally not committed to the truth of the propositional content, which contributes to the sense that these examples are related to interrogatives.

On Gunlogson's analysis, rising (and falling) intonation belongs to an intonational lexicon. The logic $\mathcal{L}_U$+$\mathcal{L}_L$ provides a useful setting in which to view this intonational lexicon. With $\mathcal{L}_U$+$\mathcal{L}_L$, we can translate rising intonation as the same sort of term that we use for utterance modifiers, except that in this case we have invariable relativization to the addressee's index (somewhat like the interrogative version of utterance modifiers in (34). The details are given in (36); I ignore the felicity condition for the sake of simplicity (it could be added by partializing $\ulcorner$**rise**$\urcorner$).

(36) a. rising intonation $\rightsquigarrow$ $\ulcorner$**rise**$\urcorner$: $\langle a_U, t_U \rangle$

b. $[\![\lambda\ulcorner S\urcorner.\ulcorner\textbf{rise}\urcorner(\ulcorner S\urcorner)]\!]^{\mathcal{D},s,a} = [\![[\![\ulcorner S\urcorner]\!]^{\mathcal{D},s,a}]\!]^{\mathcal{D},a,a}$

c. *Ed fled?* $\rightsquigarrow$ $\ulcorner$**rise**$\urcorner$($\ulcorner$Ed fled?$\urcorner$) : t_U

d. $[\![\ulcorner\textbf{rise}\urcorner(\ulcorner\text{Ed fled?}\urcorner)]\!]^{\mathcal{D},s,a} = [\![[\![\ulcorner\text{Ed fled}\urcorner]\!]^{\mathcal{D},s,a}]\!]^{\mathcal{D},a,a} = [\![\textbf{flee}(\textbf{ed})]\!]^{\mathcal{D},a,a}$

In essence, $\ulcorner$**rise**$\urcorner$ simply changes the speaker index to the addressee index. As a result, the interpretive procedure for the logic moves the interpretation to the addressee's world-view model. (In the interpretation, the addressee a is both speaker and hearer. This is strange, but it should be kept in mind that it is only for purposes of interpretation.) This seems to do justice to Gunlogson's treatment, at least to the extent that an extensional, presupposition-free analysis can do that.

The entirely new element of the analysis is that rising intonation takes a sentence, rather than a proposition, as its argument. The effect of this is that it locates the contribution of the intonation in the utterance situation — the key terms in (36) are drawn from $\mathcal{L}_U$, which talks about the upper, discourse portion of the structures.

This has at least one welcome consequence: an account of why rising intonation is not semantically embeddable:

(37) Ed said he dated a gorilla?

$\not\Rightarrow$ addressee is publicly committed to: the proposition that Ed dated a gorilla

$\not\Rightarrow$ Ed is publicly committed to: the proposition that Ed dated a gorilla

$\Rightarrow$ addressee is publicly committed to: the proposition that Ed said Ed dated a gorilla

The behavior parallels that of utterance modifiers. This restriction follows from the typing specification for $\ulcorner$**rise**$\urcorner$: it takes only root-level, assertive members of D_p as possible arguments.

7 Performative honorifics in Japanese

The previous section treats rising intonation on a declarative as an abstract sentential adverb. Performative honorifics seem to be another instance of such items. The classic description of honorification in Japanese is Harada 1976, which offers the following characterization of performative honorification:

(38) "Very roughly, we can say that one uses performative honorifics in order to talk 'politely' to the addressee, to make one's speech sound 'milder'." (Harada 1976:507)

Performative honorification is distinguished from propositional ('argument-oriented') honorification in that it does not require the presence of a syntactic phrase whose denotation could be the subject of honorification. Here's a typical example:

(39) Ame ga huri-masa-ta.
rain SUBJ *fall*-HON-PAST
'It rained.' (Harada 1976:502; my transliteration)

The honorific morpheme is realized on the verb form, as it is for all kinds of honorification in Japanese. But its semantic force is sentence-level. To achieve this result, we can assume that the performative morpheme in the syntax merely licenses the presence of a term $\ulcorner$**perf-hon**$\urcorner$ in the semantic parsetree. I offer a definition of this term in (40).

(40) $\ulcorner\textbf{perf-hon}\urcorner \overset{def}{=} \lambda\ulcorner S\urcorner.\ulcorner\textbf{politely}\urcorner(\ulcorner\textbf{utter}\urcorner)(\ulcorner S\urcorner)(\ulcorner\textbf{the-speaker}\urcorner) : \langle r_U, t_U\rangle$

As with utterance modifiers and (my treatment of) rising intonation on declaratives, this is a function that take a sentence to return a truth-value. The treatment is quite close to that of utterance modifiers, in that it involves a restriction on the relation named by $\ulcorner$**utter**$\urcorner$. Using this meaning, we can have parsetrees such as (41).

(41)

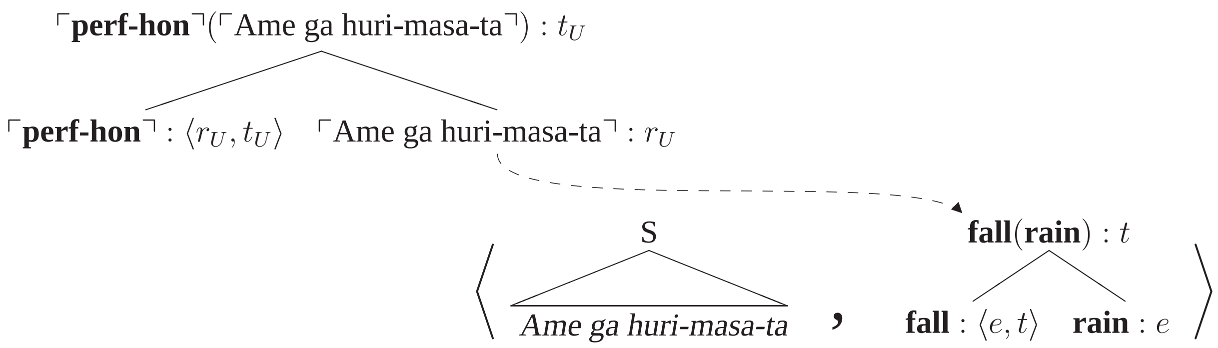

The dashed line represents the interpretation function. The upper parsetree limits the admissible utterance situations to just those in which the speaker says *Ame ga huri-masa-ta* politely. The lower parsetree expresses the proposition that it is raining.

Like utterance modifiers, performative honorifics are unembeddable. Harada (1976) is firm on this point. He writes that, "performative honorifics are not allowed to occur in nondirect discourse clausal complements" (p. 503). The examples that accompany the generalization include (42).

(42) *Boku wa [kyoo Yamada sensei ga **ki** **-mas-u** / **o-ide ni nari- mas-u**] koto o
I SUBJ *today* *Yamada* *teacher* TOP HON *come* / *come* COMP OBJ
sukkari wasure-te i-ta.
entirely *forget*-GERUND *be*
'I completely forgot that Professor Yamada is coming today.' (Harada 1976:544)

Once again, we achieve the limitation via the typing specification for $\ulcorner$**perf-hon**$\urcorner$. It names a function taking only root-level sentences into same (type $\langle r_U, t_U \rangle$). Thus, embedded clauses, being subparts of full sentences, are not admissible arguments.

8 Quotation

The logic $\mathcal{L}_U + \mathcal{L}_L$ excels at describing phenomena that are tied to speech acts. Thus, not surprisingly, quotation provides useful, complex motivation for it. Though quotation does not have a uniform semantics (Predelli 2003; and below), we can capture a class of cases in which quotation marks delimit a speech act. The essential step is the idea that when *say* takes a quotative complement, it translates as $\ulcorner$**utter**$\urcorner$, which is defined in (11) as a function from sentences to functions from entities to truth values ($\mathcal{L}_U$ type $\langle p_U, \langle e_U, t_U \rangle\rangle$). On this hypothesis, we have the following as a typical translation:

(43) *Jed said, "Ed fled"* $\rightsquigarrow$ $\ulcorner$**utter**$\urcorner$($\ulcorner$Ed fled$\urcorner$)($\ulcorner$**jed**$\urcorner$) : t_U

One might be initially suspicious of the decision to treat this *say* as $\ulcorner$**utter**$\urcorner$, rather than as, say, **say**. But there is a genuine lexical ambiguity between *say* when it has a clausal complement and *say* when it has a quotative complement. The differences show up clearly when one looks at inversion. With quotation, inversion is possible, as seen in (44a); with indirect quotation, it is impossible, as seen in (44b).

(44) a. "Ed fled", said Jed.

b. *(That) Ed fled said Jed.

The lexical ambiguity claim is further supported by the many languages that employ different morphemes for the two kinds of construction. (My thanks to Judith Aissen for discussion of this point, May 30, 2003.) But even internal to English we find good evidence for an ambiguity. So I'll regard the distinguished translation for *say* with a quotative complement as justified.

The translation in (43) gives us access to all the information about the sentence *Ed fled.* The interpretation is spelled out in (45):

(45) a.

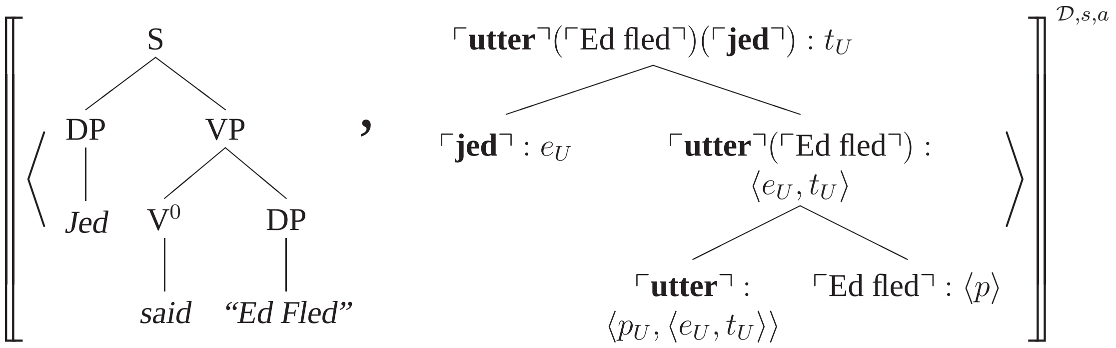

b. $= V_{\mathcal{U}}\big(\ulcorner\textbf{utter}\urcorner(\ulcorner\text{Ed fled}\urcorner)(\ulcorner\textbf{jed}\urcorner)\big)$

c.

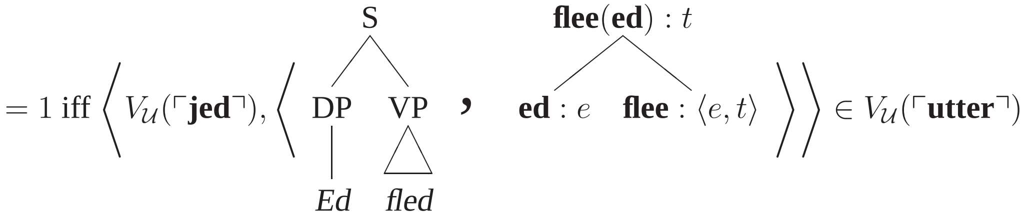

To determine what Jed said when he said *Ed fled*, we move to Jed's model to interpret the interpretation of ⌜Ed fled⌝:

(46)

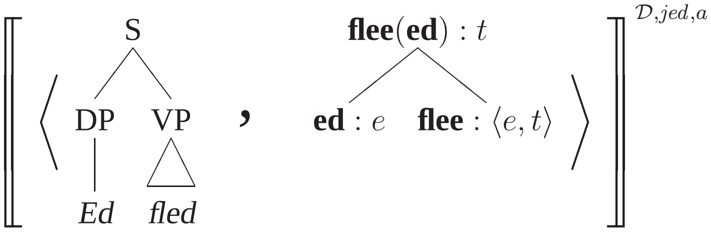

$= V_{jed}(\textbf{flee}(\textbf{ed}))$

That is, Jed informed the other members of A that, in his model $\mathcal{M}_{jed}$, it is true that Ed fled. This shift in the speaker parameter (from s to jed) is, I assume, engendered by the presence of the quotation marks. I cannot, at present, state this condition *in* the terms of $\mathcal{L}_U{+}\mathcal{L}_L$. But we can enforce it metalogically:

(47) If

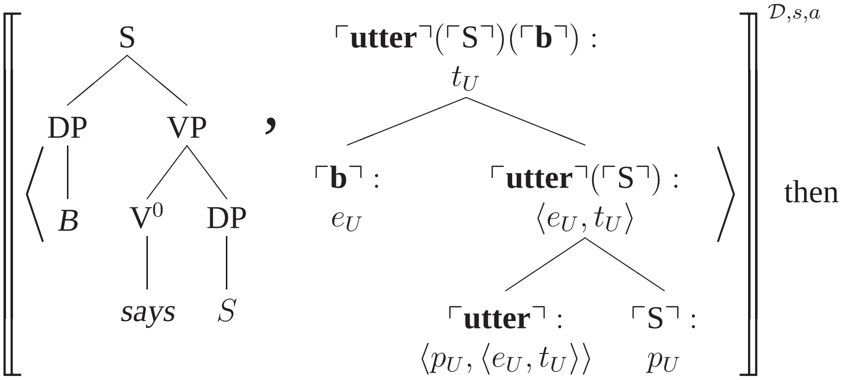

then $[\![\,[\![\ulcorner \mathrm{S}\urcorner]\!]^{\mathcal{D},s,a}\,]\!]^{\mathcal{D},[\![\ulcorner \mathbf{b}\urcorner]\!]^{\mathcal{D},s,a},a}$

The mess of interpretation brackets on the right is fairly easily sorted out. First, we apply the interpretation function $[\![\cdot]\!]^{\mathcal{D},s,a}$ to the one-node parsetree decorated with $\ulcorner$S$\urcorner$ (these are the innermost brackets). The result is a pair of trees in D_p. Next, we interpret the parsetree of that pair using $[\![\cdot]\!]^{\mathcal{D},s,a}$. When we do that, the interpretation is relativized to the second argument to $[\![\ulcorner \mathbf{utter}\urcorner]\!]^{\mathcal{D},s,a}$, here, $[\![\ulcorner \mathbf{b}\urcorner]\!]^{\mathcal{D},s,a}$.

This is an extremely fine-grained view of quotation. But that seems entirely correct. It meets the basic requirement that the speaker of a sentence containing a clausal quotation need not know the meanings of the words inside that quotation. Consider, for instance, a discourse structure in which Jed thinks that *flee* means *confess*, whereas the speaker, call her Ali, thinks that *flee* means what it actually means in English:

(48)

D_e	$=$	$\{nancy, ed, jed, ali\}$
D_t	$=$	$\{0, 1\}$
$V_{jed}(\mathbf{confess})$	$=$	$\{d \in D_e \mid d$ flees$\} = \{nancy\}$
$V_{jed}(\mathbf{flee})$	$=$	$\{d \in D_e \mid d$ confesses$\} = \{ed\}$
$V_{jed}(\mathbf{ed})$	$=$	ed
$V_{jed}(\mathbf{nancy})$	$=$	$nancy$

D_e	$=$	$\{nancy, ed, jed, ali\}$
D_t	$=$	$\{0, 1\}$
$V_{ali}(\mathbf{confess})$	$=$	$\{d \in D_e \mid d$ confesses$\} = \{ed\}$
$V_{ali}(\mathbf{flee})$	$=$	$\{d \in D_e \mid d$ flees$\} = \{jed\}$
$V_{ali}(\mathbf{ed})$	$=$	ed
$V_{ali}(\mathbf{nancy})$	$=$	$nancy$

In this model, (43) is true, because when we interpret $\mathbf{flee}(\mathbf{ed})$ in Jed's model, we actually ask whether ed is a member of the set of all things that confessed.

But the semantics for quotation is even more fine-grained than that. Consider, for example, that we regard (49b) true, but (49c) false, in the situation set up with (49a).

(49) a. Jed: “The police arrested Ed”.

b. Jed said, “The police arrested Ed”.

c. Jed said, “Ed was arrested by the police”.

The propositional content of the direct quotations is the same throughout (49). Thus, a view of quotation that regards it as roughly like a regular propositional attitude will derive the incorrect result that both (49b, c) are equally felicitous.

We can expand the account so that it is sensitive even to phonological details. Consider, for example, (50).

(50) a. When in Santa Cruz, I say “[eɪ]pricots are delicious”.
b. When New York, I say “[æ]pricots are delicious”.

In order to keep the exposition simply, I’ve so far ignored the phonology. But it would be easy to move to a view on which sentences are triples: for all $\mathcal{S} \in D_p$, we say that $\mathcal{S} = \langle \mathcal{T}^p, \mathcal{T}^s, \mathcal{T}^m \rangle$, where $\mathcal{T}^p$ is a phonological representation of some kind. Then we could add to $\mathcal{L}_U$ predicates like ⌜**ash-initial**⌝, defining them as follows:

(51) $[\![\ulcorner \textbf{ash-initial} \urcorner]\!]^{\mathcal{D},s,a} = \{ \langle \mathcal{T}^p, \mathcal{T}^s, \mathcal{T}^m \rangle \in D_p \mid$ the leftmost node in $\mathcal{T}^p$ is labelled with æ$\}$

The phrase [æ]*pricots* is a member of the set picked out by ⌜**ash-initial**⌝, whereas eɪ*pricots* is not.

The logic $\mathcal{L}_U$ can talk about all aspects of sentences. Thus, we have a maximally fine-grained view of quotation if we base it in ⌜**utter**⌝.

8.1 Some quotation-marks, not all of them

It would be a mistake to think that all instances of quotation marks (or the intonation that signals them in speech) are properly analyzed in the above terms. Consider, for instance, the example in (52a), and its quotation-free counterpart (52b).

(52) a. The answer is “yes”.
b. The answer is yes.

In many cases, these do not differ in truth conditions. But they can part company. Suppose, for instance, that we are taking a yes–no yes. Suppose the correct answer to question 7 is *yes*. Then we would use (52a). In contrast, we use (52b) when we wish to answer a question positively. For instance, if someone asks me whether I am called Chris, I would reply with (52b). I sense that (52a) would actually be false in this situation.

But it seems unwise to treat (52a) using the theory of quotation described above. The example does not house a distinct speech act inside its quotation marks. It seems to me that we simply want to say, for these cases, that the domain of entities D_e can contain linguistic objects. Thus, in the context of the yes–no test described above, it could be that *answer* denotes the set $\{$*yes*, *no*$\}$. Then (52a) could simply be a matter of ascribing the meaning of *the answer* to *yes* (or equating the two; the details depend on one’s theory of these copular constructions).

This is not the only instance of quotation that lacks its own distinguished speech act. Some other examples are given in (53).

(53) a. The sentence *Ed fled* is annoyingly alliterative.

b. *Ed ran off* is longer than *Ed fled.*

(The desk is longer than the end table.)

c. Chris's favorite word is *salmagundi.*

d. I have never uttered the word *recalcitrant.*

e. No syllable in Hawaiian has a coda.

We could, of course, analyze these sentences using $\mathcal{L}_U$ terms. But this seems an abuse of the system; these examples are not inherently tied to the utterance situation. It is perhaps best, in light of these cases, to assume that the lower parts of our structures can also contain linguistic objects.

I close this section by pointing out an area of this theory of quotation that still requires work. The important generalization is that speech-act quotation is not limited to full sentences. The following example is taken from Potts 2003b.

(54) a. Ellen: *The Godfather II* is a total snooze.

b. Frank: Well, Pauline Kael said that this "total snooze" is a defining moment in America cinema.

This example is useful because it is so clear that the phrase *total snooze* is quotative. One could not easily argue that the quotation marks were, for instance, merely a device for signalling that the phrase is evaluated at a particular, sentence-internal intensional index. Rather, we really do need a theory of quotation to understand how *total snooze* — better, *"total snooze"* — is evaluated.

This seems straightforward. What does it mean, though, for the combinatoric system? If, in (54), *"total snooze"* translates as ⌜total snooze⌝, then we are in danger of deriving the wrong meaning. After all, we do not, with this example, claim that the movie *the Godfather II is* the linguistic object named by ⌜total snooze⌝. Rather, we claim that it has the property named by *total snooze.* We need some way to get from the quotation to the property that it names. I believe that the tools for doing this lie hidden in some combination of the conventional implicature logic of Potts 2003c and the present descriptive apparatus. Some additional comments appear in section 9 below. I leave the details to a later draft of this paper.

9 Metalinguistic negation

As noted briefly above, in $\mathcal{L}_U$ we can define properties of sentences that reference their phonology, their syntax, or their semantics. Here are some examples

(55) a. ⌜**stress-initial**⌝ (a property of words and sentences)

b. ⌜**composed-of-open-syllables**⌝ (a property of words and sentences)

c. ⌜**deems-inappropriately-blunt**⌝ (a function from sentences to functions from entities to truth values)

Predicates like these are helpful in formalizing and understanding existing observations about the way that metalinguistic negation works. Consider, for instance the following example, which is due to Horn (1989:371).

(56) He didn't call the POlice, he called the poLICE.

The small capitals illustrate the stressed syllable. In the case at hand, we are dealing with the following two lexical items (I use boldface on a node's label to indicate that it is stressed):

(57) $[\![\ulcorner \text{POlice} \urcorner]\!]^{\mathcal{D},s,a} =$

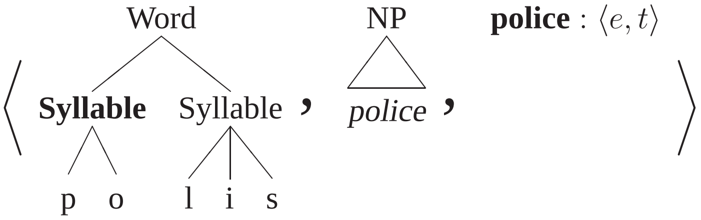

$[\![\ulcorner \text{poLICE} \urcorner]\!]^{\mathcal{D},s,a} =$

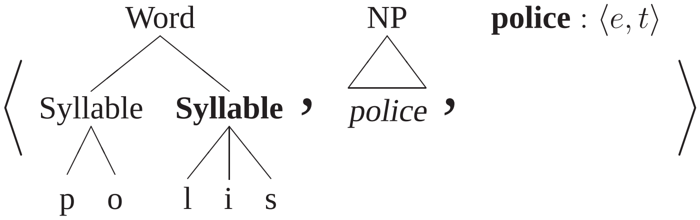

The first of these is in the set defined by the meaning of ⌜**stress-initial**⌝. The second is not.

The second conjunct of the example guarantees that we should not allow the negation to act here as a function on propositions. That is, it should not have its classical meaning (as a function reversing truth values or as a function taking a set of worlds to its complement). Rather, it should act as a restriction on appropriateness; in (56), it says that initial stress on *police* is inappropriate. This functor is definable in $\mathcal{L}_U$:

(58) *not* (metalinguistic use) $\rightsquigarrow$ ⌜**inappropriate**⌝ : $\langle p_U, t_U \rangle$

It seems, then, that *police* must appear in two capacities in these examples. One the one hand, we want to derive **call**(**the**(**police**))(x_7), where x_7 is the meaning of *he*. This means that *police* should have its usual $\mathcal{L}_L$ translation. On the other hand, we clearly need access to a translation of *police* as a constant of $\mathcal{L}_L$ of type s_U — a subordinate phrase. This will permit us to manipulate the objects in (57) using the functor in (58). Our goal, then, should be the following pair of propositions:

(59) a. **call**(**the**(**police**))(x_7) : t

b. ⌜**inappropriate**⌝(⌜POlice⌝) : t_U

We've seen meanings like this throughout this paper. But this is the first instance in which we need to appeal to a single element in two guises internal to a single composition. It is clear that we are dealing with a pair of structures. But we can overlay them fairly easily, suggesting that a unified treatment is possible; I provide an accurate representation in (59).

(60)

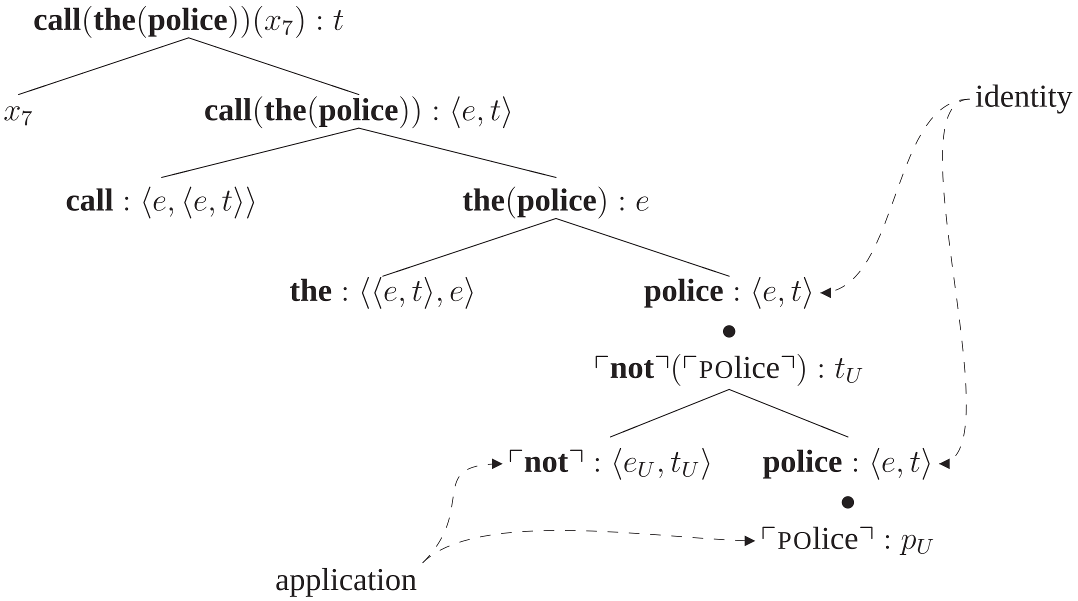

Composition schemes very much like this are the subject of much of Potts 2003a,b,c. I am as yet unsure how best to derive structures like this in a principled fashion. But I sense that the job is doable. The solution is also likely to yield a treatment of subsentential quotations such as (54), which I left unanalyzed (for now) in section 8.1 above.

10 Metadiscourse

Natural languages are rich with expressions that organize other expressions — items like *firstly* and *above* when used internally to a text to refer to various parts of that text or indicate its direction. If we add a temporal index to the upper parts of our models, then we can describe their effects in a way that makes good on the intuition that they belong to the metadiscourse, a feature of their semantics that is recognized even by good, nontechnical style manuals like Williams 1990:

(61) "Metadiscourse is the language we use when, in writing about some subject matter, we incidentally refer to the act and to the context of writing about it. [. . .] We use metadiscourse to list the parts or steps in our presentation: *first, second, third, finally* [. . .]" (Williams 1990:40)

The logic $\mathcal{L}_U + \mathcal{L}_L$ can get a grip on this kind of assessment. In this section, I develop a semantics for the following abbreviated discourse:

(62) The book *Babar* is a classic of children's literature.
⋮
Above, I mentioned the children's classic *Babar*.

In contexts such as this, *above* does not have a spatial interpretation; writing the entire text out on a single line would not render the use of *above* false or infelicitous. This is because *above*, in its capacity as a metadiscourse operator, refers to way the text containing it is organized. The use would be false if the mention of the children's classic occurred after it in the temporal order of the text.

To describe this, we need only add to the structures for $\mathcal{L}_U + \mathcal{L}_L$ a temporal index. We can assume, as is common in tense logic, that this temporal ordering is modelled by a set of numbers. In this case, I choose the set of natural numbers on their natural ordering, to avoid problems associated with temporal adjacency. Textual time, unlike real time, does seem to operate in a series of discrete, countable steps.

So we enrich the discourse structures with a set $\mathbb{N}$ of natural numbers. Actually, so that we don't get the temporal indices confused with the truth values $\{0, 1\}$, I assume for this presentation that $\mathbb{N}$ is the set of all natural numbers greater than 2.

I use $\mathbf{t}, \mathbf{t}_3, \mathbf{t}_4, \ldots$ to name these entities in the expected way. Their type is j_U. Furthermore, I temporalize the meaning of $\ulcorner$**utter**$\urcorner$, as in (63).

(63) a. $\ulcorner\textbf{utter}\urcorner : \langle p_U, \langle e_U, \langle j_U, t_U\rangle\rangle\rangle$

b. $\llbracket \ulcorner\textbf{utter}\urcorner(\ulcorner S\urcorner)(\ulcorner\textbf{the-speaker}\urcorner)(\ulcorner\mathbf{t}\urcorner)\rrbracket^{\mathcal{D},s,a} = 1$ iff s uttered $\ulcorner S\urcorner$ at $\llbracket\ulcorner\mathbf{t}\urcorner\rrbracket^{\mathcal{D},s,a}$

Thus, if we suppose that the sentences in (62) are the first and fourth sentences in the text, and if we ignore, for now, the presence of *above*, then we have the following simplified analyses:

(64) a. $\llbracket \ulcorner\textbf{utter}\urcorner(\ulcorner\text{Babar is a classic}\urcorner)(\textbf{the-speaker})(\ulcorner\mathbf{t}_3\urcorner)\rrbracket^{\mathcal{D},s,a}$
$= 1$ iff s uttered *Babar is a classic* at 3

b. $\llbracket \ulcorner\textbf{utter}\urcorner(\ulcorner\text{I mentioned Babar}\urcorner)(\textbf{the-speaker})(\ulcorner\mathbf{t}_{16}\urcorner)\rrbracket^{\mathcal{D},s,a}$
$= 1$ iff s uttered *I mentioned Babar* at 16

In this setting, *above* emerges as a temporal modality, equivalent in force to the tense-logical operator generally symbolized with P (the dual of H). For the sake of perspicuity, I symbolize this temporal operator with $\ulcorner$**above**$\urcorner$, which I defined as in (65)

(65) a. *above* (metadiscourse) $\rightsquigarrow \ulcorner\textbf{above}\urcorner : \langle p_U, \langle j_U, t_U\rangle\rangle$

b. $\llbracket \ulcorner\textbf{above}\urcorner(\ulcorner S\urcorner)(\mathbf{t}_k)\rrbracket^{\mathcal{D},s,a} = 1$ iff
   i. there is a time $i \in \mathbb{N}$ such that $i < k$ and
   ii. there is an $\ulcorner S'\urcorner$ such that $\llbracket \ulcorner\textbf{utter}\urcorner(\ulcorner S'\urcorner)(\mathbf{t}_i)\rrbracket^{\mathcal{D},s,a} = 1$ and
   iii. $\llbracket\ulcorner S'\urcorner\rrbracket^{\mathcal{D},s,a}$ entails $\llbracket\ulcorner S\urcorner\rrbracket^{\mathcal{D},s,a}$

Thus, we arrive at a complete analysis of the second sentence in (62)

(66) a. Above, I mentioned the children's classic *Babar*. $\rightsquigarrow$

b. $\llbracket \ulcorner\textbf{above}\urcorner(\ulcorner\text{I mentioned Babar}\urcorner)(\ulcorner\textbf{the-speaker}\urcorner)(\ulcorner\mathbf{t}_7\urcorner)\rrbracket^{\mathcal{D},s,a} = 1$ iff
   i. there is a time $i \in \mathbb{N}$ such that $i < 7$ and
   ii. there is an $\ulcorner S'\urcorner$ such that $\llbracket \ulcorner\textbf{utter}\urcorner(\ulcorner S'\urcorner)(\mathbf{t}_i)\rrbracket^{\mathcal{D},s,a} = 1$ and
   iii. $\llbracket\llbracket\ulcorner S'\urcorner\rrbracket^{\mathcal{D},s,a}\rrbracket^{\mathcal{D},s,a}$ entails $\llbracket\llbracket\ulcorner\text{I mentioned Babar}\urcorner\rrbracket^{\mathcal{D},s,a}\rrbracket^{\mathcal{D},s,a}$
   $= \llbracket\llbracket\ulcorner S'\urcorner\rrbracket^{\mathcal{D},s,a}\rrbracket^{\mathcal{D},s,a}$ entails $\llbracket\textbf{mention}(\textbf{babar})(\textbf{the-speaker})\rrbracket^{\mathcal{D},s,a}$

Additional metadiscourse operators can be defined by adapting the above illustrative example.

11 Conclusion

The layered semantics defined here provides a treatment of utterance modifiers on which they literally modify utterances, by modifying the utterance relation. They are ‘second-order’ in the sense that they operate in a universe full of sentences. We move a formerly metagrammatical perspective into the grammar itself. This perspective leads to suitable analyses of a diverse group of sentential modifiers (rising intonation on declaratives, performative honorifics) as well as quotation, metalinguistic negation, and metadiscourse items like *above*. Some important kinds of quotation and metalinguistic negation are not yet fully accounted for, but the outstanding problems appear solvable. It seems likely that one can use $\mathcal{L}_U + \mathcal{L}_L$ to formalize a wide range of pragmatic conditions, especially those that concern the conditions on felicitous belief reports (e.g., Berg’s (1988) *assentialist principle*). This is the direction I plan to head next.

References

Bach, Kent. 1999. The myth of conventional implicature. *Linguistics and Philosophy* 22(4):367–421.

Bach, Kent and Robert M. Harnish. 1979. *Linguistic Communication and Speech Acts*. Cambridge, MA: MIT Press.

Barwise, Jon and John Perry. 1983. *Situations and Attitudes*. Cambridge, MA: MIT Press.

Bellert, Irena. 1977. On semantic and distributional properties of sentential adverbs. *Linguistic Inquiry* 8(2):337–351.

Berg, Jonathan. 1988. The pragmatics of substitutivity. *Linguistics and Philosophy* 11(3):355–370.

Blackburn, Patrick, Claire Gardent, and Wilfried Meyer-Viol. 1993. Talking about trees. In *Proceedings of the Sixth Conference of the European Chapter of the Association for Computational Linguistics*, 21–29. San Francisco: Morgan Kaufmann.

Blackburn, Patrick and Wilfried Meyer-Viol. 1997. Modal logic and model-theoretic syntax. In Maarten de Rijke, ed., *Advances in Intensional Logic*, 29–60. Dordrecht: Kluwer.

Chung, Sandra and William A. Ladusaw. 2003. *Restriction and Saturation*. Cambridge, MA: MIT Press.

Fenstad, Jens Erik. 1997. Partiality. In Johan van Benthem and Alice ter Meulen, eds., *Handbook of Logic and Language*, 649–682. Amsterdam and Cambridge, MA: North-Holland and MIT Press.

Gabbay, Dov M. 1999. *Fibring Logics*. Oxford: Clarendon Press.

Gunlogson, Christine. 2001. *True to Form: Rising and Falling Declaratives as Questions in English*. Ph.D. thesis, UC Santa Cruz, Santa Cruz, CA.

Harada, S. I. 1976. Honorifics. In Masayoshi Shibatani, ed., *Syntax and Semantics*, Volume 5: Japanese Generative Grammar, 499–561. New York: Academic Press.

Heim, Irene and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell Publishers.

Horn, Laurence R. 1989. *A Natural History of Negation*. Chicago: University of Chicago Press. Reissued 2001 by CSLI.

Jacobson, Pauline. 1999. Towards a variable-free semantics. *Linguistics and Philosophy* 22(2):117–184.

Jacobson, Pauline. 2000. Paycheck pronouns, Bach-Peters sentences, and variable-free semantics. *Natural Language Semantics* 8(2):77–155.

Potts, Christopher. 2003a. Conventional implicatures, a distinguished class of meanings. In Gillian Ramchand and Charles Reiss, eds., *The Oxford Handbook of Linguistic Interfaces*, Studies in Theoretical Linguistics. Oxford: Oxford University Press.

Potts, Christopher. 2003b. Expressive content as conventional implicature. In Makoto Kadowaki and Shigeto Kawahara, eds., *Proceedings of the North East Linguistics Society 33*. Amherst, MA: GLSA.

Potts, Christopher. 2003c. *The Logic of Conventional Implicatures*. Ph.D. thesis, UC Santa Cruz.

Predelli, Stefano. 2003. *Linguistics and Philosophy* 26(1):1–28.

Schlenker, Philippe. 2003. A plea for monsters. *Linguistics and Philosophy* 26(1):29–120.

Searle, John. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge, UK: Cambridge University Press.

Williams, Joseph M. 1990. *Style: Toward Clarity and Grace*. Chicago: University of Chicago Press. With two chapters coauthored by Gregory G. Colomb.