Expressive content as conventional implicature[*]

Christopher Potts

UC Santa Cruz

1.   A second dimension of meaning

In defining *conventional implicature*, Grice (1975) supplied terminology for reasoning about a group of secondary lexical entailments indicating speakers' attitudes toward (parts of) their main assertions. The passage in question is potentially ground breaking: it draws a restrictive boundary around (Grice's) pragmatic theory, and it homes in on a complex mode of expression that operates under the radar of most linguistic operators (e.g., verbs of saying, negation). Unfortunately, Grice's (1975:34–35) discussion is based on an unconvincing example (*therefore*). Subsequent research enriched the factual basis only slightly (*still*, *but* and its synonyms). The evidence seemed not to match the potential of the original definition; at least one commentator (Bach 1999) has challenged the very existence of CIs, based largely on the apparent meagerness of the support.

This bleak picture does not accurately reflect the status of CIs in natural language. In truth, Grice's (1975) original definition picks out a broad range of expressions (Potts 2003b). The concern of the present paper is *expressive* (*emotive, affective*) *modification*. Many items falling under this heading match perfectly the definition of CIs that Grice laid out. Attention to their semantics validates Grice's original definition and answers the question of how to manage expressive content, which is ubiquitous in discourse but challenging to assimilate to existing semantic theories.

The class of expressive modifiers is itself huge and diverse. I limit attention mostly to *expressive attributive adjectives* (EAs) such as *damn* in (1a) and epithets like the expressions highlighted in (1b–d). Huddleston and Pullum (2002:36) identify the content of EAs as conventionally implicated; extending this insight to epithets systematizes existing observations about their interactions with commanding operators (Aoun et al. 2001).

(1)   a.   Ed refuses to look after Sheila's <u>damn</u> dog.

          b.   Right after Chuck agreed to help out, <u>the jerk</u> boarded a plane for Tahiti.

          c.   Right after he agreed to help out, <u>that jerk Chuck</u> boarded a plane for Tahiti.

          d.   Every Democrat with [a proposal for reform]$_1$ claims [<u>the stupid thing</u>]$_1$ deserves public support.

We can characterize the content of these items as speaker-oriented (nonpresuppositional) assertions that are, in an intuitive sense, independent of the regular semantic content, which I henceforth refer to as the *at-issue entailments*.[1] Section 3 takes up the task of working systematically through this characterization, amplifying each of its points and bringing them back to the original definition of CIs, summarized below in (2).

The main theoretical goal is a description logic that is capable of modelling the properties described in definition (2) and, in turn, capable of a transparent formalization of the semantics of EAs and epithets. Karttunen and Peters (1979) identify CIs as multidimensional phenomena; the proposal in section 5 seeks to draw out the insights embodied in their formalism and transfer them to a type-driven setting (Klein and Sag 1985). A consideration of epithets that are dependent upon quantifiers leads, in section 7, to a novel application for the insights of Gawron's (1996) (multidimensional) *restriction logic*.

Alternative classifications of claimed CI content always loom large, especially those that seek to assimilate it to at-issue entailment. I close by spelling out the challenges facing an account in this vein, arguing that important aspects of Grice's (1975) definition cannot be captured in these terms without conceding all points of substance to the CI approach. The main utility of this discussion is not to discourage alternative formalizations, but rather to show that a few key concepts are bound to turn up in any adequate description.

2. Conventional implicature

As noted, Grice's (1975:34–35) original discussion of CIs is entwined with informal assertions about the semantics of *therefore*. Since the utility of this example has often been questioned (e.g., Bach 1999), I extract from the passage a more abstract set of properties:

(2)   a.   CIs are part of the conventional (lexical) meaning of words.

          b.   CIs are commitments, and thus give rise to entailments.

          c.   These commitments are made by *the speaker of the utterance* "by virtue of the meaning of" the words he chooses.

          d.   CIs are logically and compositionally independent of what is "*said* (in the favored sense)".

---

[1] The term 'at-issue entailment' permits us to use *assertion* and *semantic content* in a way that includes CIs, while at the same time suggesting the secondary, deemphasized role CIs play in discourse.

Potts 2003a,b explore the contrasts between this definition and those standardly assumed for conversational implicatures, presuppositions, at-issue entailments, and intonational meanings. Here, I offer only a fast summary of those discussions, so that we can proceed confident that Grice did not merely happen upon a previously known commodity.

Clauses (2a) and (2b) effectively define CIs as beyond the reach of Grice's (1975) theory of pragmatics: their content traces back to arbitrary features of lexical items and cannot be nullified by contextual factors. After introducing CIs in a short paragraph, Grice turns his attentions away from them and to "nonconventional implicatures" — conversational implicatures, which are relations among propositions and do not (for Grice) trace back to specific lexical items, or even always to linguistic stuff.

This definition is also quite different from that of *presupposition*, though 'conventional implicature' and 'presupposition' are often used interchangeably (e.g., Beaver 1997, 2001; Krahmer 1998; Gamut 1991). But one finds only contrasts between (2) and the usual conception of presuppositions. As Green (2000:461) notes, essentially all definitions of presupposition maintain that they are discourse-conditioned. By clause (2b), the presence or absence of a CI is not negotiable if the lexical content is fixed.

Traditionally, presuppositions must also, absent nullifying contextual factors, be entailed by the context at the point of utterance. If they are not, then accommodation is required. Definition (2) makes no mention of this backgrounding requirement (van der Sandt 1988:74). This is telling, but it must be noted that recent definitions of 'presupposition' depart from the word's ordinary-language meaning, allowing, for instance, that "the listener rapidly and unconsciously adjusts his or her model of the domain of discourse to support the presuppositions of the speaker" (Steedman 2000:654). This reduces somewhat the difference between CIs and presuppositions, but even this revised definition does not collapse with (2). A look at presupposition logics in general reveals a more significant point of contrast. Clause (2d) defines CIs as independent of the at-issue content, whereas all presupposition logics create a strict dependency between the two kinds of meaning: the presupposition must be true for definedness. This is the heart of the reconstruction of presuppositions in terms of partiality (Heim 1983; Beaver 2001; Krahmer 1998; and others).

CIs' disjointness from at-issue entailment is stipulated in (2d). This stipulation is given substance with (2c): at-issue entailment often involves content that is relativized to some entity other than the speaker; in *Sue believes it is raining*, the proposition that it is raining is asserted to hold only in (all of) Sue's belief worlds.

Finally, it is easy to see that (2) bears no resemblance to any existing construal of intonational meaning. There are abstract theoretical connections between focus and CIs — alternative semantics is a multidimensional theory — and CI constructions often trace back to intonational properties rather than traditional lexical items (Potts 2003a). But there is no sense in which (2) recasts an existing or viable conception of intonational meaning.

In sum: we've run the gamut of accepted classes of meaning. None specifies the same set of properties as (2). But one can define new abstract classes of meaning on a whim. The challenge is in showing that the definition matches something in natural language.

3. Expressive content

EAs and epithets provide data showing that definition (2) picks out a nonempty class of linguistic phenomena. An important clause of (2) is (2c), which relativizes CI content to the speaker of the utterance. I offer two naturally occurring cases in which this is undoubtedly the intended interpretation of the EA:

(3)   a.   "We bought a new electric clothes dryer [. . .] Nowhere did it say that the <u>damn</u> thing didn't come with an electric plug!"[2]

        b.   "I remember practicing for my first Confession in the second grade and of course Sister role-played the priest. Trying to do a good job, I told her all the big sins [. . .] Never again!!!! For my Penance she made me say the <u>damn</u> rosary."[3]

These show that an embedded EA *can* be interpreted with widest-scope. They do not, however, show that it *must* be so interpreted. Two sorts of arguments fill this gap. The first appeals to standard presupposition holes (negation, modalization, conditionals, and questioning). Example (4) is representative:

(4)   It just not true that Sheila's damn dog is on the couch!

This sentence cannot be read as negating the speaker's disapprobation of Sheila's dog; it is judged false if and only if Sheila's dog is not on the couch. Testing with the other holes reveals the same invariance. But, unlike presuppositions, EA content need not be entailed by the input context for felicitous use — we process it as we do at-issue content — and (3) shows that EA content need not embed under presupposition *plugs* (verbs of saying and other performatives). The lesson seems to be that we should deny (with Beaver (2001:19ff)) that invariance under the holes provides a sufficient condition for presuppositionhood.

But presupposition holes are far more likely to project semantic content than presupposition plugs. So we need to look to the plugs to establish that EAs are widest-scope operators. Pairs like (5) fill out the picture.

(5)   a.   Clinton: The damn Republicans should be less partisan.

        b.   Bush: Clinton says the damn Republicans should be less partisan.

The sentence in (5b) is an unlikely report of Clinton's utterance (5a). Even those with a limited grasp of the language recognize that *damn*, even inside an indirect quotation, is heard as a contribution of the speaker of the utterance. Though *Clinton* is the subject of the propositional attitude verb in (5b), the content of *damn* is not relativized to Clinton, but rather to Bush, the speaker. The meaning of (5b) is roughly given by the pair of propositions in (6).

---

[2] <http://jjdavis.net/blog/arc20010325.html>
[3] <http://www.nunstories.com/SampleStories/SampleStories.html>

(6)   a.   Clinton says the Republicans should be less partisan

          b.   Bush looks with disapprobation upon Republicans

I refer to Dong 1971 and Cruse 1986:271–272 for the same generalization based on similar examples.

To report the content of *damn* in (5a), one must resort to a paraphrase (*Contemptuous of Republicans, Clinton says. . .*) or assign the EA a special intonation contour indicating that it is a quotative utterance: one signals this with heavy emphasis on the EA in speech, quotation marks in print; the result is subject to felicity conditions parallel to those of anaphoric resolution. It's worth stressing that, though quotative utterances can give rise to what appear to be embedded readings, a general analysis should treat them as scopeless. For instance, some speakers find, contrary to the judgment of Dong (1971), that *fucking* in (7) can be read as relativized to John's beliefs.

(7)   John says that his landlord is a fucking scoutmaster.

Such non-speaker-oriented readings require heavy emphasis on *fucking*, an indication that they are quotative. In virtue of being in this quotative form, an EA can be attributed to an entity not mentioned in the sentence, as in (8).

(8)   a.   Sue: John's landlord is a fucking scoutmaster.

          b.   Eddie: Well, John wouldn't say that his landlord is a "fucking" scoutmaster. He rather admires scoutmasters, and so do I.

If these readings were a matter of scope, the EA could be attributed only to Eddie or the speaker. Neither scoping gives us the intuitively correct reading. One might think that the addressee is another index we could exploit for (8). But suppose Eddie is speaking to a crowd. He reads Sue's claim aloud, then says (8b). Nothing changes about how we read "*fucking*". Potts 2003b contains an in-depth analysis of quotative cases like this. I mention them here only to set them aside; their special, anaphoric nature indicates that they do not even constitute narrow-scope readings of epithets, for which widest-scope appears to be the only possibility.

In a phrase, EAs are syntactically embeddable but semantically unembeddable. They are not alone among expressive modifiers in displaying this mix of properties. Epithets are also speaker-oriented even when embedded below propositional attitude verbs. I offer first an attested example; nowhere does the news story mentioned in (9) characterize the person who broke into Clements' house negatively.

(9)   "The story says that the idiot broke into Clements' home and attacked and robbed him. Obviously, particularly since the guy only got $27, if Clements doesn't do something, the guy comes back and does it again a day or two later."[4]

---

[4] <http://209.157.64.200/focus/news/780053/posts>. The news story is on this page as well as the reaction to it in (9).

Once again, we now know it is *possible* to interpret epithet content with widest-scope. Pairs like (10) permit us to strengthen this, assigning epithets to widest-scope status as well.

(10)   a.   Ellen: *Rambo* is a stupid movie.

      b.   Frank: I liked *Rambo*, but Ellen says the stupid thing isn't worth seeing.

This mini-discourse is strange. Though the context works to support a reading of the epithet in (10b) on which its propositional content is attributed to Ellen, we interpret it as an emotive contribution of Frank's (the speaker's). Aoun et al. (2001) make this observation about epithets in Lebanese Arabic, specifying that they should have a "main clause" interpretation (p. 386). As with EAs, apparently embedded readings reveal themselves to involve quotation. In (11b), the quotative "*total snooze*" is oriented neither to the speaker nor the matrix subject, but rather to Ellen, who is not directly mentioned in the sentence.

(11)   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.

We can gather together the above observations at a slightly more technical level. The widest-scope effects for EAs and epithets are obtained if we ensure that no at-issue functor (negation, quantifier, modal, etc.) ever takes a function containing their meanings as an argument. If we achieve this, then this expressive content will never end up in the scope of anything. It will thus attain speaker-oriented status in the same manner in which main clause assertions attain this status. With an important qualification to allow, for example, **damn** to apply to **republican** in *damn Republicans*, we want at-issue and CI content to be impermeable to each other. Achieving this would have benefits not only in terms of scope, but also in terms of our ability to model the independence of the at-issue and CI dimensions (clause (2d)). At the level of expressive content, this would model speakers' intuitions that they can agree to either of the examples in (12) without committing themselves to the emotive baggage engendered by the modifiers.

(12)   a.   The damn Republicans are aggressively cutting taxes.

      b.   We saw that bastard Charlie at the pool hall.

The other clauses are also easy to match to facts about EAs and epithets. It is quite clear, for instance, that we are dealing with a specific group of lexical items; the content does not flow from the maxims and general considerations about how conversation works. (In the usual terminology, expressive content is 'nondetachable'.)

4.   An undistinguished syntax

A dominant theme of the above is the apparent widest scope of expressive content. This would easily be achieved in the absence of the CI hypothesis if it could be shown that

the items in question had a syntax involving root-level adjunction. The main purpose of this interlude in the semantic discussion is to head-off this alternative, by showing that neither EAs nor epithets display syntactic properties that suggest a nonstandard syntax. A more concrete result of this attention to the syntax is that it suggests an ideal shape for the interpreted structures, and in turn highlights inevitable deviations from that ideal (see (25)).

English has a rich array of EAs (*bleeding*, *(gol)darn*, *mother loving*, and so forth), and new ones are coined fairly regularly, often by popular media as substitutes for swear words. It is useful to use *damn*, *darn*, and *frigging* in examples, because they are unambiguously EAs (the regular adjective is *damned*, and *frigging* is a tame alternative to *fucking* but without the literal meaning for most speakers).

Huddleston and Pullum (2002:553) show that EAs are syntactically much like other strictly attributive adjectives (e.g., *former*, *premier*). They are restricted to prenominal position, but freely intermingle with other adjectives:

(13) a. Sheila said that we must look after her (biggest) friggin' brown dog.

     b. "What's the Big Friggin' Deal About Sony PlayStation 2?"[5]

In German, where attributive adjectives are marked for case, EAs are not distinguished from other adjectives in this sense. For example, one has *dass die verdammte Industrie zu geizig ist* ('that the damn.NOM-FEM industry.NOM-FEM too miserly is'), with the same obligatory case-marking found on regular prenominal adjectives.

There are apparently no restrictions on the kind of determiner that can head a nominal containing an EA:

(14) a. The company says that every damn piece of software we use has to be made by them!

     b. Ed claims that no damn idea of his should be ignored!

This is an important point. As discussed in the next section, EAs, though nominal internal, often modify the entire proposition expressed by the immediate clause. Similar properties are found with adjectives like *occasional* in *An occasional native strolled by* (Stump 1981), which means the same thing as *Occasionally, a native strolled by*. One might seek to extend to EAs Zimmermann's (2000) syntactic movement analysis, on which the infrequency adjective raises to form a quantifier $INFREQ$, denoting a family of sets of event–individual pairs. At least three arguments suggest that this is not a fruitful direction. First, as Zimmermann shows (p. 295), infrequency adjectives permit adverbial readings only with articles and possessives. The adverbial reading of EAs is not limited by the determiner. Second, infrequency adjectives are required to appear adjacent to the determiner for their adverbial readings; in contrast, both examples in (13) can involve the EAs as clause-level modifiers, but neither is determiner-adjacent. Finally, EAs express no notion of (in)frequency; $INFREQ$ is quite obviously not the proper denotation for these expressions.

---

[5] <http://www.zdnet.com/anchordesk/stories/0,10738,2645040,00.html>

The upshot of the above syntactic discussion is that an EA plays no special role in the syntax of a nominal it appears in, beyond simply adjoining as any modifier would. That is, EAs determine routine structures of the form in (15), in which they are simply left-adjoined modifiers (node labels highly negotiable).

(15)   the $[_{N^1}$ $[_{A^2}$ damn] $[_{N^1}$ Republicans]]

It seems safe to conclude that the contrasts between EAs and other attributive adjectives don't follow from properties of the structures they determine.

We can say the same for epithets. The work of Jackendoff (1972:§4.1), Lasnik (1976), Aoun and Choueiri (2000), and Aoun et al. (2001) has shown that epithets are cross linguistically much like full nominal expressions in their syntactic distribution. For instance, they are sensitive to c-commanding antecedents, as in (16).

(16) $^{\#}${Paul$_1$/No musician$_1$/He$_1$} thinks [the vain snob]$_1$ is tiresome.

Where discourse considerations remove this c-command effect, epithets are allowed:

(17)   a.   "In 1654 a friend had written him$_1$ to ask if Pascal$_1$ could solve the *problème des parties*, or problem of points."[6]

         b.   The professor wrote every student$_1$ to ask if the lazy bum$_1$ could solve the problem of points.

Aoun and Choueiri (2000:2–3) report that some Lebanese Arabic epithets are internally distinguished form other nominals, by the presence of an extra definite marker. But, outwardly, it appears cross-linguistically true that epithets are treated by the syntax as regular definite nominals.

5.   A description logic for CIs

The heart of my formal reconstruction of the at-issue/CI divide is the set of types, defined in (18).

(18)   a. e^a and t^a are basic at-issue types.

         b. e^c and t^c are basic CI types.

         c.   If τ and σ are at-issue types, then $\langle \tau, \sigma \rangle$ is an at-issue type.

         d.   If τ is an at-issue type and σ is a CI type, then $\langle \tau, \sigma \rangle$ is a CI type.

         e.   The full set of types is the union of the at-issue and CI types.

---

[6] Thomas A. Bass. 1985. *The Eudaemonic Pie*, p. 117. Boston: Houghton Mifflin Company.

I adopt a syntactic view of these types: rather than acting merely to index sets of denotations (Montague 1970; Halvorsen and Ladusaw 1979), they serve as categories for lambda terms (Barendregt 1992; Reynolds 1983; Shan 2002). In essence, the types regulate semantic composition in the same way that natural language syntactic categories regulate the projection of category labels in syntactic structures. Since typing information is essential to my analysis, I always provide terms along with their types: where α is a term and τ is a type, the expression '$\alpha : \tau$' is glossed 'α is of type τ' or 'α is in τ', in the same way that '*dog* : N' would naturally be read '*dog* is of category N'. I remain informal about the nature of the terms themselves, though my notation is standard.

Even at the level of the type definition, we begin to make inroads in the task of understanding CIs. The definition in (2) specifies that the two dimensions are independent. This must, though, be qualified: CIs are characterizable as comments upon the at-issue core. In order to function in this capacity, they take at-issue content and use it to form CI content. Importantly, the reverse is unattested: we do not find at-issue content borrowing from the CI dimension. This fundamental asymmetry is matched by an asymmetry in the type definition: we have types taking at-issue types into CI types, but not the reverse: for no σ, τ is $\langle \sigma^c, \tau^a \rangle$ in the set defined by (18). Since these types provide, in essence, a space of meaning categories to work in, no meaning can violate the asymmetry we see at the descriptive level.

I state the semantic combinatoric rules as tree-admissibility conditions. I assume that a single-node subtree labelled only with a lexical meaning is well formed. In addition, I provide a set of local-tree conditions. A parsetree $\mathcal{T}$ is well-formed if and only if every local tree of $\mathcal{T}$ instantiates one of the tree conditions. I leave open the interpretation of these trees. They can be viewed in various ways (as proof rules, tree-generation procedures, etc.). I specify only that the ordering of terminal elements is irrelevant.

The most basic and familiar of the rules is at-issue application:

(19) **At-issue application**

$$
\begin{array}{cc}
\alpha(\beta) : \tau^a & \alpha(\beta) : \tau^a \\
\diagup \quad \diagdown & \diagup \qquad \diagdown \\
\alpha : \langle \sigma^a, \tau^a \rangle \quad \beta : \sigma^a \qquad \alpha : \langle \sigma^a, \tau^a \rangle \;\bullet\; \gamma : \rho^c \quad \beta : \sigma^a
\end{array}
$$

As noted, the terminal nodes should be regarded as unordered. The bullet '$\bullet$' is a metalogical symbol used to separate two independent lambda terms. At its core, this is just the standard rule of functional application (Klein and Sag 1985:171; Heim and Kratzer 1998:44). It has the same interpretation as the clause of the lambda calculus that specifies that if α is a term of type $\langle \sigma, \tau \rangle$, and β is a term of type σ, then $(\alpha(\beta))$ is a term of type τ. The two conditions are basically identical; the one on the right is necessary to allow for structures in which application happens in the presence of a term decorated with a CI term.

The second composition rule is the heart of the CI logic. The rules of Karttunen and Peters (1979) follow a pattern: with only a few exceptions, lexical items have one dimension that is an identity function. The exceptions to this either involve presupposition

triggers (e.g., existence presuppositions in quantifier restrictions) or cases of conversational implicatures (e.g., the name *Bill* suggests maleness). Rather than posit lots of these identity functions, I build them into the system, in the form of the following rule of CI functional application:

(20) **CI application**

$$\beta : \sigma^a \bullet \alpha(\beta) : \tau^c \qquad\qquad \beta : \sigma^a \bullet \alpha(\beta) : \tau^c$$
$$\alpha : \langle \sigma^a, \tau^c \rangle \qquad \beta : \sigma^a \qquad\qquad \alpha : \langle \sigma^a, \tau^c \rangle \qquad \beta : \sigma^a \bullet \gamma : \rho^c$$

This formalizes the intuition that CIs comment upon, but do not intrude upon, the at-issue meaning of a sentence. We can describe its action as follows: apply the CI functor to the at-issue meaning to produce a CI meaning. But, in addition, pass along the at-issue argument unmodified. Here, for instance, is the scheme for a modified common noun like *damn Republicans* (I henceforth stack independent lambda terms, for typographic reasons):

(21)
$$\textbf{republican} : \langle e^a, t^a \rangle$$
$$\bullet$$
$$\textbf{damn}(\textbf{republican}) : t^c$$
$$\textbf{republican} : \langle e^a, t^a \rangle \quad \textbf{damn} : \langle \langle e^a, t^a \rangle, t^c \rangle$$

We duplicate **republican**. The overall result is that the at-issue value for a parsetree $\mathcal{T}$ is always identical to the tree $\mathcal{T}'$ obtained from $\mathcal{T}$ by trimming all nodes labelled with terms of type τ^c, for any τ.

In order to model the content of EAs completely, it seems best to adopt one more composition scheme:

(22) **Isolated CIs**

$$\beta : \langle e^a, t^a \rangle \qquad\qquad \beta : \langle e^a, t^a \rangle$$
$$\alpha : t^c \qquad \beta : \langle e^a, t^a \rangle \qquad\qquad \alpha : t^c \qquad \beta : \langle e^a, t^a \rangle \bullet \gamma : \rho^c$$

The purpose of this rule is to account for readings of EAs on which they seem not to take any arguments at all. Rather, they function to express the speaker's disposition. In (23), for example, the EA need not express disapprobation of expressive modifiers. The most prominent reading conveys a measure of frustration or uneasiness in general.

(23)   What are we going to do with these damn expressive modifiers?

Using rule (22), we can have the following parsetree for *damn expressive modifiers*:

(24)
$$\textbf{expressive-modifiers} : \langle e^a, t^a \rangle$$
$$\textbf{frustrated}(\textbf{the-speaker}) : t^c \quad \textbf{expressive-modifiers} : \langle e^a, t^a \rangle$$

I add (22) only for perspicuity's sake. We could get by with (20), by allowing some EAs to involve vacuous abstraction; *damn* could translate as $\lambda f.\ \textbf{frustrated}(\textbf{the-speaker})$. But this seems only to obscure the true nature of what happens in these cases. For some tentative independent motivation for a rule like this, one that does not involve only CI propositions, I refer to Potts 2003b.

The only remaining combinatoric question is what to do with CI propositions as they collect. Various options are discussed in Potts 2003b. The most appealing fits in with the view implicit in the above rule notation that parsetrees are first-class semantic objects: we specify that the interpretation of any parsetree $\mathcal{T}$ is the interpretation of the at-issue term on $\mathcal{T}$'s root node plus the interpretation of any term in $\mathcal{T}$ of type t^c.
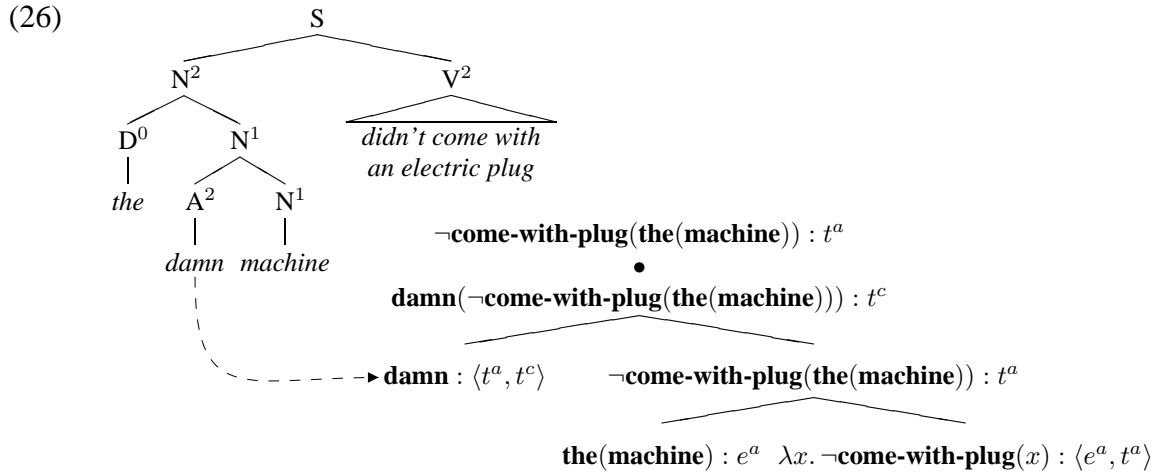
6.   Lexical meanings

The technical apparatus is in place, and I have suggested how it works for EAs. The task of formulating lexical entries raises new issues.

I begin the lexicography with EAs. They are a rich, open class. But, semantically, it is unclear how the members are distinguished from each other. There is a kind of scale, with *darn* at the tame end and *fucking* at the obscene end, but the details are heavily discourse-conditioned: a playful use of *fucking* might sound less angry or disapproving than a stern *damn*. I simplify by giving only the narrow semantics of EAs, mapping them all to the same lambda term.

But which term? Another issue that I have not addressed thus far is the fact that EAs need not be interpreted as taking their common noun sisters as arguments. The immediately containing full noun phrase or the entire clause are both fair targets:

(25)   a.   We have to look after Sheila's damn dog.

b.   Nowhere did the instructions say that the damn machine didn't come with an electric plug!

With (25a), the speaker probably does not express disapprobation of all dogs, but rather just Sheila's; (25b) arguably expresses the speaker's frustration with the fact that the machine in question arrived plugless. Since all known syntactic evidence mitigates against movement of attributive adjectives, but the existence of these readings indicates that some can act as clause-level functors, the simplest overall picture involves viewing the parsetrees in this work as semantic; the interpreted object is not a natural language syntactic object, but rather a designated semantic structure, related to the syntax by mapping principles (as in Lexical Functional Grammar, or interpretations of Minimalism that derive logical forms using rules that are different from those of overt syntax). This view of logical forms frees us to have sentence representations such as (26) without worrying that the attributive adjective violates constraints on movement.

(26)

$\neg$**come-with-plug**(**the**(**machine**)) : t^a

damn($\neg$**come-with-plug**(**the**(**machine**))) : t^c

damn : $\langle t^a, t^c \rangle$ $\neg$**come-with-plug**(**the**(**machine**)) : t^a

the(**machine**) : e^a $\lambda x. \neg$**come-with-plug**(x) : $\langle e^a, t^a \rangle$
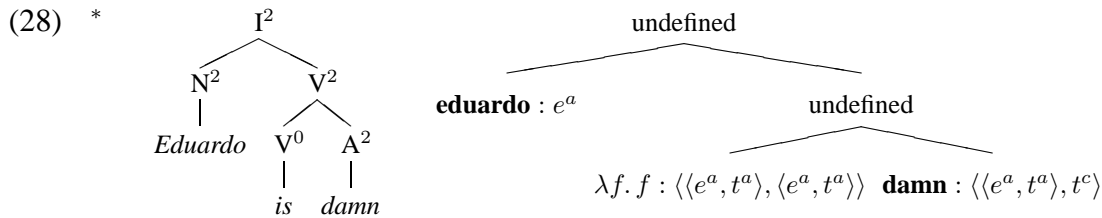
At the level of denotations, the variability of the arguments to an EA indicates polymorphism in the domain of the EA meaning. I offer a general lexical entry, on which an EA can take any argument in $\langle \tau^a, t^a \rangle$ to produce a term of type t^c:

$$
(27) \quad \left\{
\begin{array}{c}
\textit{damn} \\
\textit{bloody} \\
\vdots \\
\textit{fucking}
\end{array}
\right\} \rightsquigarrow \lambda X. \, \textbf{bad}(^\cap X) : \langle \langle \tau^a, t^a \rangle, t^c \rangle
$$

The nominalizing type-shifter $^\cap$ is that of Chierchia (1984). When defined extensionally, it takes any function and returns the plural individual composed of all members of the input set. (In symbols, $^\cap = \lambda f. \, \iota x : \forall y : f(y) \to x \leqslant y$, where $\leqslant$ is the 'part of' relation and ι is the definite operator.)

The translation in (27) contains a function called **bad**. To be more precise would imply a degree of understanding of the semantics of EAs that I do not possess. My interest is in managing the content, whatever it is. (Section 9 briefly addresses the difficult question of what the models for expressive modifiers look like.) The important thing for present purposes is the typing of the EA denotation: it takes at-issue sets and returns CI truth values. This suffices to account for EAs' inability to appear in predicative position; I illustrate using a *be* that takes properties into same, but the result holds across theories of the copula.

(28)   *

undefined

eduardo : e^a undefined

$\lambda f. \, f : \langle \langle e^a, t^a \rangle, \langle e^a, t^a \rangle \rangle$ **damn** : $\langle \langle e^a, t^a \rangle, t^c \rangle$

In (28), *be* denotes an identity function on properties. Since the problem is one of typing, the result holds equally if *be* translates as $\lambda f \lambda x. \, f(x) : \langle \langle e^a, t^a \rangle, \langle e^a, t^a \rangle \rangle$. Other theories

of *be* also fail to produce a meaningful meaning; though we might find cases in which *be* is a function on generalized quantifier types ($\langle\langle e^a, t^a\rangle, t^a\rangle$), these could not be functions on generalized quantifier *conventional implicature* types, for the simple reason that we have no types whose first members are CI types ($\langle\langle e^a, t^c\rangle, t^a\rangle \notin$ (18)). This same logic affords an account of why EAs, despite forming a loose scale of strength, are not gradable. Since gradable modifiers like *very* are functions from at-issue meanings into same, they cannot apply to the CI dimension of a word like *damn*.

The lexical entries facilitate an explanation for why CIs are unembeddable, but the general result is largely independent of individual meanings. In order to embed a CI meaning under an at-issue operator A, it would have to be the case that A had a type of the form $\langle t^c, \tau^a\rangle$. But we have no such types. As a result, there is no provision in the set of local tree conditions for taking CI content into at-issue content. An example helps bring out this property of the logic: if we take the parsetree in (26) and embed it under *Ed says*, we have the following semantics:

(29) $\textbf{say}(\textbf{ed}, \neg\textbf{come-with-plug}(\textbf{the}(\textbf{machine}))) : t^a$

         $\textbf{ed} : e^a$ $\lambda x.\,\textbf{say}(x, \neg\textbf{come-with-plug}(\textbf{the}(\textbf{machine}))) : \langle e^a, t^a\rangle$

                 $\textbf{say} : \langle t^a, \langle e^a, t^a\rangle\rangle$ $\neg\textbf{come-with-plug}(\textbf{the}(\textbf{machine})) : t^a$
                                                  $\bullet$
                         $\textbf{damn}(\neg\textbf{come-with-plug}(\textbf{the}(\textbf{machine}))) : t^c$

By our provision for interpreting trees, this is defined as having the value $(1, 1)$ iff Ed says the machine doesn't come with a plug, and it is bad that the machine doesn't come with a plug. The speaker-orientation of the latter, CI proposition follows from the same interpretive specifications that make the former proposition speaker-oriented.

It is somewhat more challenging to find lexical denotations for epithets, since they have the same rich internal structure as regular noun phrases. I propose that all epithets have the structure of those that place an appositive modifier on a name (Huddleston and Pullum 2002:447–448); where the name is absent, a free variable fills its spot:

(30)   a.   that/the stupid jerk Eddie

       b.   that/the stupid jerk x_{25}

A fuller picture is given in (31).

(31)
              N^2 $\textbf{chuck} : e^a$

       N^2 N^1 $\bullet$
                             $\textbf{bastard}(\textbf{chuck}) : t^c$

   D^0 N^2 *Chuck*

                    $\lambda y.\,\textbf{bastard}(y) : \langle e^a, t^c\rangle$ $\textbf{chuck} : e^a$
   the *bastard*

I note that this analysis ignores the definite article in the semantics. This seems in line with the fact that, unlike regular definites, epithets do not presuppose that a unique entity meets the conditions specified by their descriptive content.

7. Quantifiers and a variable environment dimension

So far I have considered only epithets that have a referential semantics. However, we have already seen cases in which an epithet appears to be (in some sense) dependent upon a higher quantifier:

(32) a. The professor wrote every student$_1$ to ask if the lazy bum$_1$ could solve the problem of points.

b. The judge told every dead-beat dad$_1$ that the bum$_1$ must help out.

Aoun et al. (2001) also report instances of quantifier–epithet connection:

(33) kəll muttahame saʔalto ʔəza   ha-l-maʒduube nħabasit
      *each suspect.*SF *asked.*2P *whether this-the-idiot.*SF *imprisoned.*3SF
      'Each suspect, you asked whether this idiot was imprisoned.'

<div align="right">(Aoun et al. 2001:373, (6))</div>

A large part of Aoun et al. 2001 is devoted to explaining why such relationships are apparently not possible in Lebanese Arabic unless an island boundary intervenes. Their explanation for why (34a) is impossible is that it would have to have the logical form in (34b), in which the first occurrence of x is unbound.

(34) a. *kəll muttahame ʕrəfto   ʔənno ha-l-ma-ʒduube nħabasit
      *each suspect.*SF *know.*2P *that   this-the-idiot.*SF *imprisoned.*3SF
      'Each suspect, you know that this idiot was imprisoned.'

<div align="right">(Aoun et al. 2001:(5a))</div>

b. x is an idiot and each suspect x is such that you know x was imprisoned

But, given only what they say, it is mysterious why (33) is not bad because its logical form contains an unbound variable:

(35) x is an idiot and each suspect x is such that you asked whether x was imprisoned

It would seem that the present account is in the same bind. Based on the referential cases, a natural interpretation of the components of meaning for (32b) would seem to be (36).

(36) a. at-issue: the judge told every dead-beat dad x that x must help out

b. CI: x is a bum

If this is the correct representation, then we need to adjust the logic somehow. In virtue of the fact that (36a) is a distinct term from (36b), the variable x in the CI meaning is free; it gets its interpretation from the assignment, thereby denoting an individual.

But attention to the semantics of the epithet's expressive content in these cases reveals that the connection between the two dimensions of meaning is not nearly so tight as binding would imply. The analysis we seek is actually one on which the propositional content of the epithet is a generic quantification over the restriction on the at-issue determiner. Let's see why this is so.

A factual argument against binding across the dimensions comes from the insensitivity of the epithet's content to changes in the at-issue quantifier it connects with. The expressive content of (37) is identical to that of (32b) despite changes in the at-issue meaning resulting from variation in the quantified object of *tell*.

(37) a. The judge told almost {every/no} dead-beat dad that the bum must help out.

 b. The judge told {few/many/most} dead-beat dads that the bums must help out.

For all examples in (37), the CI is a generic quantification of roughly the form 'generally, dead-beat dads are bums'. We would wrongly predict variation in the nature of the CI if the CI contained a variable bound from the at-issue dimension. (This is another point of contrast with presuppositions, which display a rather complex mix of properties when in the nuclear scope of a quantifier (Heim 1983; Cooper 1983:152–154; Krahmer 1998:§4).)

These facts are admittedly subtle. It would be helpful to supplement them with factual considerations of a less delicate nature. One argument of this form concerns the well-known generalization that only referring expressions can associate with nonrestrictive modifiers (Karttunen 1976; McCawley 1998:451, Potts 2002:83; and others), a generalization that extends even to instances in which the associate is a bound variable (hence locally referential).

(38) a. Every student$_1$ spoke with a psychiatrist of hers$_1$ that welcomes calls at home.

 b. *Every student$_1$ spoke with a psychiatrist of hers$_1$, who welcomes calls at home.

 c. Sally$_1$ spoke with a psychiatrist of hers$_1$, who welcomes calls at home.

Epithets are nonrestrictive; (32b) is not equivalent to (39).

(39) The judge told every dead-beat dad who is a bum that he must help out. $\neq$ (32b)

This descriptive fact alone provides independent reason to doubt that the CI aspect of epithets is dependent upon an at-issue quantifier, as this would involve nonrestrictive modification of a dependent element. (On the present analysis, the nonrestrictiveness of epithets follows directly from the fact that their descriptive content cannot possibly influence the at-issue proposition expressed.)

But, as noted, there is an undeniable link between the at-issue quantifier and the epithet's content. A factually accurate analysis of (32b) is (40). (G is a generic quantifier.)

(40)   The judge told [every dead-beat dad]$_1$ that [the bum]$_1$ must help out.

      a.   at-issue: $\forall x : \textbf{db-dad}(x) \rightarrow \textbf{tell}(\textbf{the-judge}, x, \textbf{must}(\textbf{help}(x)))$

      b.   CI: $\mathsf{G}x : \textbf{db-dad}(x) \rightarrow \textbf{bum}(x)$

The technical question is how to get **db-dad**, the restriction to dead-beat dads, into the CI dimension in a systematic way. The restriction logic of Gawron (1996) provides exactly the needed mechanism. Space considerations prevent me from giving the full details of how to incorporate the necessary features of restriction logic into the present logical system; I refer to Potts 2003b for them. Here, I provide only a suitable meaning-language representation, with a brief description of its semantics:

(41)   a.   environment: $x \mid \textbf{db-dad}(x)$

      b.   meaning: $\forall x : \textbf{tell}(\textbf{the-judge}, x, \textbf{must}(\textbf{help}(x))) : t^a$
                                  $\mathsf{G}x : \textbf{bum}(x) : t^c$

The environment serves to restrict possible values of the variable x to those that have the property named by **db-dad**. An innovation of restriction logic is making this restriction hold even for bound variables; the semantics for quantifiers like $\forall$ and G appeal to the information in the environment for their restrictions. The result is an identity for x beyond its binding quantifier. The environment is itself a new dimension of meaning.

    I close this section by noting a stubborn bit of context sensitivity: referential epithets involve predication of their descriptive content of some individual, whereas dependent epithets place this descriptive content in the nuclear scope of a generic quantification over an independently provided restriction. Thus, whereas a referential *the bum* x_{25} has its CI dimension given by $\textbf{bum}(x_{25}) : t^c$, a dependent use has the CI dimension $\mathsf{G}x : \textbf{bum}(x) : t^c$. It would be good to remove this context dependency, but I do not see how. I note, though, that it is not the only contrast of this type: (38) indicates that pronouns behave differently depending on how they receive their values.

8.   A scope-based alternative

It is worth heading off an alternative analysis that attempts to locate the usual properties of expressive content modifiers in their scopal properties, thereby reducing expressive content to regular at-issue meanings. A reduction of this sort is assumed to be feasible in Kaplan 1989:55, fn. 71, a short footnote in which Kaplan characterizes epithet examples like *John says that the lying S.O.B. who took my car is honest* in terms of what he calls 'pseudo de re'. He suggests that quantifying-in is a suitable mechanism for handling such examples.

    Since I do not advocate this approach, it would be foolhardy to develop such an analysis. My strategy is more general: I enumerate the properties an at-issue account would have to obtain, along with critical comments and links with the present work.

Lexical marking   Something must set epithets and EAs apart from other modifiers. This will require lexical marking. In the CI analysis, the marking is type-theoretic, but other techniques are feasible. It should be noted, though, that while the type-theoretic account yields easy answers to the question of why EAs cannot be modified by gradable adjectives or appear in predicative position, this seems likely to involve a series of unconnected stipulations on the at-issue account of their meanings.

Obligatory wide-scope   Epithets and EAs do not display the kind of rich scope variability that makes a scope-shifting account so appealing for regular quantified expressions. There seem only to be wide-scope, speaker-oriented readings.

Nonrestrictiveness   As noted above, epithets and EAs are never restrictive. But if we simply interpret their content (in the actual world) as we would a regular modifier's, we will allow them to function restrictively; *the damn Republicans* could pick out a proper subset of the Republicans (those that the speakers looks on with disapproval). Similarly, an epithet in the scope of a quantifier would wrongly be able to contribute to the restriction on that quantifier.

Multidimensionality   In order to model speakers' intuitions that expressive content is a comment upon the at-issue core, the at-issue account needs to shift the domain of sentences from $\{0, 1\}$ to $\{0, 1\}^n$, the set of all ordered tuples with elements drawn from $\{0, 1\}$. (Parallel comments hold for an intensional semantics.) This multidimensionality is the defining feature of the CI account.

In sum   Given the above comments, it seems that the at-issue account must duplicate all the important aspects of the CI theory. In the end, disputes might all be located in the terminology. While we should welcome alternative formalizations of expressive content, it seems likely that all such accounts will converge on the above findings.

9.   On to the question of what the models are like

I mentioned in section 8 that Kaplan (1989) uses the term 'pseudo de re' for, roughly speaking, epithets. His description suggests that he would welcome EAs under this heading as well. So it is worth providing the following snippet from his footnote:

(42)   "I do not see that the existence of the *pseudo de re* form of report poses any issues of theoretical interest to make it worth pursuing." (Kaplan 1989:555–556, fn. 71)

I hope to have made a convincing case that this dismissal is too hasty. EAs and epithets help validate Grice's (1975) definition of conventional implicatures as speaker-oriented comments upon the at-issue core of utterances, potentially reinvigorating this class of meanings. The account suggests a method for managing expressive content. In this setting, previously intractable-looking problems start to make sense; promising directions to take the present work include obviative marking, the complex and elusive content of honorifics in languages like Japanese, and emotive, nonrestrictive uses of other adjectives (*Ed said I could take one of his lovely vases — he thinks they are ugly*). Though it is useful to link

these phenomena with Grice's definition of CIs, the substantive abstract insight, the one that Grice hinted at and Karttunen and Peters (1979) made explicit, is the *multidimensionality* of these meanings.

Greater depth of coverage might yield insights into a problem I have left unaddressed: does expressive content require a revision of our usual conception of the models for semantic theory? I have assumed that EAs and epithets contribute propositions in the usual sense; model-theoretically, I have not differentiated *damn* from *red*. But speakers perceive differences. For instance, falsity seems inappropriate for *damn* in much the same way it seems inappropriate for *ouch*. One might in turn prefer a semantics in which *damn* could not, even in principle, make a false contribution. Other differences are also apparent. A single EA or epithet can be used repeatedly to indicate a speaker's hostility without the sense of redundancy that would accompany repeated use of *I feel hostile*. Cruse (1986:271ff) touches upon these points and others, but only at a descriptive level. The task of formulating a model-theory for expressive content will surely prove a challenging task. I refer to Potts 2003b for an initial attempt.

Significantly, the task of managing expressive content is separable from the task of describing precisely what that content is. This paper provides a description logic that is flexible enough to work for a wide range of model theoretic interpretations. I've shown that familiar assumptions about semantic composition extend quite readily to the domain of expressive content. It is much less clear, though, that our usual assumptions about the models for semantics will fare so well here.

References

Aoun, Joseph and Lina Choueiri. 2000. Epithets. *Natural Language and Linguistic Theory* 18(1):1–39.

Aoun, Joseph, Lina Choueiri, and Norbert Hornstein. 2001. Resumption, movement, and derivational economy. *Linguistic Inquiry* 32(3):371–403.

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

Barendregt, Henk P. 1992. Lambda calculi with types. In Sampson Abramsky, Dov M. Gabbay, and T. S. E. Maibaum, eds., *Handbook of Logic in Computer Science*, Volume II, 117–309. New York: Oxford University Press.

Beaver, David Ian. 1997. Presupposition. In Johan van Benthem and Alice ter Meulen, eds., *Handbook of Logic and Language*, 939–1008. Cambridge, MA and Amsterdam: MIT Press and North-Holland.

Beaver, David Ian. 2001. *Presupposition and Assertion in Dynamic Semantics*. Stanford: CSLI.

Chierchia, Gennaro. 1984. *Topics in the Syntax and Semantics of Infinitives and Gerunds*. Ph.D. thesis, University of Massachusetts, Amherst. [Distributed by GLSA].

Cooper, Robin. 1983. *Quantification and Syntactic Theory*. Dordrecht: D. Reidel.

Cruse, D. A. 1986. *Lexical Semantics*. Cambridge, UK: Cambridge University Press.

Dong, Quang-Phuc. 1971. English sentences without overt grammatical subject. In Arnold M. Zwicky, Peter H. Salus, Robert I. Binnick, and Anthony L. Vanek, eds., *Studies out in Left Field: Defamatory Essays Presented to James D. McCawley on the Occasion of his 33rd or 34th birthday*, 3–10. Amsterdam: John Benjamins.

Gamut, L. T. F. 1991. *Logic, Language, and Meaning*, Volume 2. Chicago: University of Chicago Press.

Gawron, Jean Mark. 1996. Quantification, quantificational domains, and dynamic logic. In Shalom Lappin, ed., *The Handbook of Contemporary Semantic Theory*, 247–267. Oxford: Blackwell Publishers.

Green, Mitchell S. 2000. Illocutionary force and semantic content. *Linguistics and Philosophy* 23(5):435–473.

Grice, H. Paul. 1975. Logic and conversation. In Peter Cole and Jerry Morgan, eds., *Syntax and Semantics*, Volume 3: Speech Acts, 43–58. New York: Academic Press.

Halvorsen, Per-Kristian and William A. Ladusaw. 1979. Montague's 'Universal grammar': An introduction for the linguist. *Linguistics and Philosophy* 3(2):185–223.

Heim, Irene. 1983. On the projection problem for presuppositions. In Michael Barlow, Daniel P. Flickinger, and Michael T. Wescoat, eds., *Proceedings of WCCFL 2*, 114–125. Stanford: Stanford Linguistics Association.

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

Huddleston, Rodney and Geoffrey K. Pullum. 2002. *The Cambridge Grammar of the English Language*. Cambridge, UK: Cambridge University Press.

Jackendoff, Ray S. 1972. *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press.

Kaplan, David. 1989. Demonstratives. In Joseph Almog, John Perry, and Howard Wettstein, eds., *Themes from Kaplan*, 481–614. New York: Oxford University Press. [Versions of this paper began circulating in 1971].

Karttunen, Lauri. 1976. Discourse referents. In James D. McCawley, ed., *Syntax and Semantics*, Volume 7: Notes from the Linguistic Underground, 363–385. New York: Academic Press.

Karttunen, Lauri and Stanley Peters. 1979. Conventional implicature. In Choon-Kyu Oh and David A. Dinneen, eds., *Syntax and Semantics*, Volume 11: Presupposition, 1–56. New York: Academic Press.

Klein, Ewan and Ivan A. Sag. 1985. Type-driven translation. *Linguistics and Philosophy* 8(2):163–201.

Krahmer, Emiel. 1998. *Presupposition and Anaphora*. Stanford: CSLI.

Lasnik, Howard. 1976. Remarks on coreference. *Linguistic Analysis* 2(1):1–22.

McCawley, James D. 1998. *The Syntactic Phenomena of English*. Chicago: Chicago University Press, 2 ed.

Montague, Richard. 1970. Universal grammar. *Theoria* 36:373–398. Reprinted in Montague, Richard. 1974. *Formal Philosophy: Selected Papers of Richard Montague*, 222–246. Edited by Richmond H. Thomason. New Haven, CT: Yale University Press.

Potts, Christopher. 2002. The lexical semantics of parenthetical-*as* and appositive-*which*. *Syntax* 5(1):55–88.

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. *The Logic of Conventional Implicatures*. Ph.D. thesis, UC Santa Cruz.

Reynolds, John C. 1983. Types, abstraction and parametric polymorphism. In R. E. A. Mason, ed., *Information Processing 83: Proceedings of the IFIP 9th World Computer Congress*, 513–523. Amsterdam: Elsevier Science. Available at <ftp://ftp.cs.cmu.edu/user/jcr/typesabpara.pdf>.

Shan, Chung-chien. 2002. A continuation semantics of interrogatives that accounts for Baker's ambiguity. In Brendan Jackson, ed., *Proceedings of SALT 12*, 246–265. Ithaca, NY: CLC Publications.

Steedman, Mark. 2000. Information structure and the syntax–phonology interface. *Linguistic Inquiry* 31(4):649–689.

Stump, Gregory. 1981. The interpretation of frequency adjectives. *Linguistics and Philosophy* 5(2):221–256.

van der Sandt, Rob. 1988. *Context and Presupposition*. London: Croom Helm.

Zimmermann, Malte. 2000. Pluractional quantifiers: the *occasional*-construction in English and German. In Brendan Jackson and Tanya Matthews, eds., *Proceedings of SALT X*, 290–306. Ithaca, NY: Cornell University.

Department of Linguistics
University of California, Santa Cruz
Santa Cruz, CA 95064

potts@ling.ucsc.edu