

The Handbook of Pragmatics

Edited by

*Laurence R. Horn and
Gregory Ward*

12 The Pragmatics of Non-sentences

ROBERT J. STAINTON

1 The Appearances

I want to begin by describing some appearances. The word "appearances" is important – as will emerge at length below, these appearances might be misleading. That warning being issued, it appears that ordinary speakers routinely utter non-sentences, and in so doing perform full-blown speech acts.

In saying this, I don't mean that they appear to produce non-linguistic gestures and such, thereby performing speech acts. (How could something be a speech act, and be non-linguistic?) Nor do I mean that they utter ungrammatical sentences and yet succeed in asserting, or asking, or ordering. It seems plausible that speakers do the latter, but that isn't the issue I will be discussing. Instead, what I mean is that speakers appear to utter, consciously and by design, fully grammatical expressions which happen to be less-than-sentential: nouns and NPs, adjectives and AdjPs, as well as PPs, VPs, and so on. That is, speakers routinely utter bare words and phrases not syntactically embedded in any sentence, and they thereby perform speech acts like asserting, asking, commanding, and so on. Again, so it appears.

That bare words/phrases can be so used might seem obvious. Many would grant, for example, that a hearer may answer a question with a mere word or a phrase. For instance, Tracy says, "Where do you live?", and Isaac replies, "London." There are also examples of correction (repairs): Tracy says, "I think we met in London," and Isaac responds, "Paris, actually." Not everyone would grant that such cases are, in fact, subsentential; some will insist that in such cases the answers are actually elliptical sentences. But, even if these are subsentential, they aren't the sort of cases I want to emphasize here.¹ Rather, I want to focus on examples in which it is not prior linguistic context but non-linguistic context that somehow "completes" what is asserted, asked, commanded, etc.

Here are some examples of the kind of thing I want to discuss. A woman could enter a room, and Leah could say to Anita, looking at the woman in the

doorway: "Sam's mom." Here, Leah says about the woman entering the room that she is Sam's mom, but what she utters is not a sentence. Instead, she utters an NP.² What's more, it's not an NP that answers an interrogative, nor is it an NP that corrects a previously spoken sentence. Similarly, Anita and Sheryl could be looking at a tote board, watching the progress of shares in Acme Internet. As the stock rises, Anita could say, "Moving pretty fast!" In this example, Anita appears to utter a bare VP, not a sentence. And, here again, it's not a VP that answers an interrogative, nor is it a VP that corrects a previously spoken sentence (I'll exclude this qualification in what follows), yet Anita still succeeds in making a statement. Other examples abound: pick up any magazine, leaf through the ads, and you will find carefully edited (and hence, surely grammatical) copy of the following sort:

- (1) America's most frequent service to Asia
- (2) Fast relief for arthritic pain
- (3) From the sun-soaked mountains of Colombia

Let me describe the appearances a bit more formally. In doing so, I'll make use of some non-obvious theoretical machinery, so in that sense what follows goes beyond "obvious appearances." But it will, I hope, clarify the phenomenon being discussed. The generalization at the level of syntax, first formulated by Barton (1990 and elsewhere), seems to be that speakers can utter not just sentences, but any MAXIMAL PROJECTION. This emphatically does include maximal projections whose GRAMMATICAL HEAD is a lexical item: a noun, an adjective, a preposition, etc. It is uses of these LEXICAL PROJECTIONS that are precisely the cases of interest here.³ As for the semantics of the things used in such cases, it does not seem to matter what content is assigned to the expression in the language. It can stand for an object, a property, or even a function from a property to a truth value, and yet still be used to make a statement, or ask a question, or issue an order. In the notation of Montague grammar, the point can be put this way: an expression need not be of SEMANTIC TYPE $\langle t \rangle$ to be used to perform a speech act; it can be of type $\langle e \rangle$, or $\langle e, t \rangle$, or $\langle \langle e, t \rangle, t \rangle$, etc.⁴ Thus, returning to an earlier example, the phrase *moving pretty fast* does not express a proposition. It doesn't even do so after reference is assigned to indexicals and such, there being none. Put differently, the phrase type *moving pretty fast* is not synonymous with the sentence type *That is moving pretty fast*, as it would have to be if the contextualized meaning of the former were to be a proposition – and this lack of synonymy of the types obtains even if one can make an assertion by tokening either type. (Notational aside: the word "type" unfortunately refers to two quite different things in semantic theorizing. There is the TYPE VERSUS TOKEN distinction, which is in play in the preceding sentence, and there is the distinction between various semantic categories in Montague grammar: semantic types $\langle e \rangle$, $\langle e, t \rangle$, $\langle t \rangle$, etc. Where there is a risk of confusion,

I will use “expression type/phrase type/sentence type” for the former, and “Montagovian semantic type” for the latter.) The complete sentence *That is moving pretty fast*, despite containing context-sensitive elements, is indeed of Montagovian semantic type (t). But the phrase *moving pretty fast* is not of type (t): its semantic type is (e, t). This becomes evident when the phrasal expression is embedded. What *moving pretty fast* contributes to the complete sentence *That stock is moving pretty fast*, once reference has been assigned to indexicals and such, is not a proposition, but a property: that property shared by things which are moving pretty fast. Thus, when a speaker utters *Moving pretty fast* on its own, it appears that she utters an expression that, even after it is contextualized, means a property, not a proposition. Similarly, the phrase *Sam’s mom*, even contextualized, does not express a proposition. The same can be said of the expressions in (1)–(3). And yet, it appears that these expression types, which purportedly have both the syntax and the semantics of ordinary phrases, can be used to make statements. Indeed, it appears that they can be used to perform speech acts of many kinds. For instance, one could ask about a displayed letter, “From Colombia?” Or one could issue a command to one’s child by saying, “To your bedroom. Right now.”

Notice too that, whereas it’s sometimes supposed that complete sentence meanings contain FORCE INDICATORS that account for the kind of act the sentences are typically used to perform, bare word and phrase meanings are assumed not to contain such things. This is still another (apparent) difference in content. Let me explain. One view of the difference in meaning between the sentences (4a–c) is that, though they share the same propositional content, viz. that-John-is-running, (4a–c) exhibit distinct force indicators.

- (4)a. John is running
- b. Is John running?
- c. Run, John!

That is, their content is bipartite: part of it is a proposition, the other part is a force indicator. The first sentence, the syntactic type, has as its non-propositional content an assertoric force indicator; the second sentence type has an interrogative force indicator as part of its context-insensitive content; and the third has an imperatival force indicator. Notice that force indicators are considered here to be part of content of the expression type: they are (at least in English) syntactically carried by the MOOD of the sentence. Mood is a constant feature of syntax, and thus force indication is a matter of context-invariant content. It is the presence of force indicators, encoded by mood, that helps explain why one makes an assertion by uttering (4a) but one asks a question by uttering (4b).

Under these assumptions, then, the phrase type *moving pretty fast* surely does not contain a force indicator as part of its context-invariant content. Clearly it has no mood as part of its syntax, neither declarative nor any other. Its constant content, then, is just a property of things, not a property/force pair.

We thus have another apparent difference between the sentence type *That is moving pretty fast* and the phrase type *moving pretty fast* – a difference that extends beyond the fact that they are of distinct Montagovian semantic types.

To clarify this difference, it will help to say a word about expression types, speech acts, and force. I don't deny that speech acts performed using subsentences exhibit force; indeed, it's part of my description of the appearances that they do so. It may even be true that there are sound pattern types – consisting of the sound of a lexical phrase, though modulated by a special intonation pattern – that have illocutionary force. For instance, the type *From Colombia?* said with rising intonation might be claimed to have interrogative force. Be that as it may, the point I'm making here is that the force of such things – the speech acts or the phrase-intonation pairs – is not inherited from the syntax and semantics of the phrase type itself, since that phrase type does not syntactically encode a force indicator. Suffice it, then, to say that the syntax of complete sentences apparently (often) encodes a force indicator, via mood, but the syntax of lexical projections does not appear to encode this. That is another apparent difference.

There is then – or at least, there appears to be – a significant mismatch in non-sentential speech between what the expression type means in the language and what the speaker of it means. Now, cases of speaker meaning that outpace (contextualized) expression meaning are very familiar in pragmatics. For instance, in CONVERSATIONAL IMPLICATURE (see Horn, this volume), the speaker means something different than (or in addition to) what his words mean, even once reference has been assigned to context-sensitive elements. As Grice would say, "what is said" in such cases does not (wholly) capture what is meant. Recall, for instance, the delightful sort of case imagined by Grice (1989: Chap. 2). Professor Koorb writes a letter of reference for a student that says only:

- (5) Mr. Tonstain has neat handwriting, and he usually arrives on time for class. Yours, J. A. Koorb.

Here what the speaker means goes well beyond what his words mean. What he means is something like: This student is appallingly bad; don't even dream of hiring him. But that is not what his words, even in context, mean. Similarly in non-sentence cases, it appears that what the speaker means, which is a proposition, is quite different from what her words mean, which is not a proposition but an object, or property, or something along those lines.

Interestingly, however, though there is this mismatch between what the expression uttered means in the context (i.e. an object, or property, etc.), and what the speaker of it meant (i.e. a complete proposition), this does not appear similar to cases of non-literal communication. Of course, there are special cases in which one can speak metaphorically or ironically while using a subsentence: Richard could utter "The next Nobel Laureate" while pointing at a notoriously brainless politician, thereby **saying** that the politico is the next

Nobel Laureate – but **meaning** that he is a buffoon.⁵ But not all uses of subsentences are non-literal. For instance, recalling Leah and Anita, if Leah knew perfectly well that the person coming through the doorway was not Sam's mother, but she wanted to convince Anita otherwise, she would have lied in uttering "Sam's mother." Leah could not later say, "Oh I didn't tell Anita that she was Sam's mother. In fact, I made no literal statement at all about the woman. Anita just drew her own conclusions." To the contrary, Leah did make a statement: she strictly and literally said, about the woman at the door, that she was Sam's mother. Unlike in the Gricean case of Professor Koorb described above, in speaking non-sententially it doesn't look like Leah merely suggested, or implicated, a proposition: what Leah does looks very much like assertion, and very much unlike non-literal speech – despite the mismatch between expression meaning and speaker meaning.

To sum up so far: speakers utter ordinary words and phrases, with the syntax and semantics of ordinary words and phrases, and thereby perform speech acts. More formally put, they produce projections of lexical items – which, seen semantically, are not of semantic type $\langle t \rangle$ and contain no force indicator – and yet they thereby make assertions, ask questions, etc. Since there is an assertion of something of semantic type $\langle t \rangle$, what the speaker means in these cases extends beyond what her words mean. And yet, this mismatch is not strikingly similar to metaphor, or conversational implicature, or speaker's reference, or other clearly non-literal speech acts. Rather, one seems to have perfectly literal communication in these cases. As I stressed at the outset, this is how things appear. (Or anyway, how they appear to me.)

There are, of course, two possible responses to such appearances. One is to say that what I've just described are only appearances. That is, one possibility is to deny that there are genuine cases of non-sentential speech acts of the sort just introduced. One must then go on to explain away the appearances. The other possibility is to say that the reason people appear to use subsentential things to perform speech acts is because that's what they really do. (Compare: "The reason the car over there looks purple is because it is purple." This certainly accounts for the appearances.) Taking this second route, the burden is not to "explain away," but to "explain how." Specifically, to explain how speakers manage to do this – a non-trivial task since, as just noted, if the phenomenon is genuine there is an important gap between the meaning of the things used and the nature and content of the act itself. Much of this chapter will be dedicated to exploring these two responses to the just-described appearances.⁶

2 Rejecting the Appearances: Introducing the Options

2.1 Not a genuine speech act

Notice that what appears to be the case is a conjunction: it appears that, in some cases anyway, speakers produce non-sentences and in so doing they perform a speech act. One obvious maneuver in resisting the appearances is to go after the second conjunct: to deny that a genuine speech act is ever performed when something less-than-sentential is produced. Let's explore this option. (Doing so will also help clarify what the second conjunct actually commits one to.)

It would be difficult to maintain that speakers and writers only utter complete sentences, even allowing for performance errors, slips of the tongue, and outright grammatical mistakes. To give just a few obvious examples, book titles are often single words or phrases: *Symbolic Logic*, *Language*, etc. And, of course, there are signs that simply say *Exit* and *Fire extinguisher*. In addition, borrowing an example from Shopen (1973), we affix phrasal labels to objects, e.g. *strawberry jam*. There are spoken cases too: if someone really wanted to, she surely could stand on a street corner and just repeat the word *cymbidium*. (According to my dictionary, this word refers to some kind of orchid.) It's hopeless to insist that in all these cases there is a sentence employed: people surely can, and they surely do, produce plain old words and phrases. The conjunction above requires more than this, however. It requires not only that words and phrases be uttered, but also that such utterances sometimes result in the speaker having performed a speech act: a conventional linguistic act like naming a ship or making a promise, a request, or an assertion, etc. But notice that when the *Exit* sign is posted there is no assertion made. No question is asked by the person who constantly repeats the word *cymbidium*. No promise is issued by the book title. It seems reasonable to say that in these cases there is simply no ILOCUTIONARY ACT at all: there might be language use, but there is no speech act properly so called. Having noticed cases of non-sentential speech in which no speech act is made, the next step would be to argue that whenever a mere word or phrase is uttered, no speech act results (putting aside, as noted at the outset, answers to questions, repairs, and the like). True enough, goes this line of thought, people utter mere words in isolation – but they do not thereby “make a move in the language game,” to use Wittgenstein's famous phrase. Put otherwise, the path of less resistance for the theorist who wishes to reject the appearances is to simply deny that a speech act is performed in subsentential speech of the kind here under consideration.

Unfortunately for those who want to explain away the appearances in this way, this line of argument seems unlikely to succeed. There may be cases that initially look like speech acts but in which one can make the case that appearances mislead. Stanley (2000) gives the example of a thirsty man, emerging parched and sunburned from the desert, who scratchily utters, “water.” Stanley says that this isn't determinately a request, or an assertion, or an order; as such, it might be written off as subsentential but not genuinely a speech act. Maybe that's right. On the other hand, there are lots of cases, like Leah's

utterance to Anita, that are lie-prone, and hence really are assertions. And there are lots of non-sentential utterances that are clearly questions. And others which are promises. And so on.⁷ So, while this strategy might allow one to explain away some apparent cases of non-sentential speech acts – e.g. an utterance of *water* by the thirsty man – it cannot on its own explain away the appearances described in section 1.

2.2 Three senses of “ellipsis”

The next obvious means of explaining away the appearances is to suppose that whenever a genuine speech act is performed, the speaker actually uses a sentence. To account for the appearances, specifically the sound produced by the speaker, one would then have to claim that the thing used was an ELLIPTICAL SENTENCE. This despite the fact that the utterance is not an answer to an explicit question, or a repair. This questions not the second conjunct, but the first: there are genuine speech acts being performed, in the cases at hand, but they aren't in fact performed using non-sentences – because elliptical sentences are a kind of sentence. So, if this maneuver works, what appears to be non-sentential speech is really sentential after all.

Now, there are lots of things that might be meant by the word “ellipsis.” It's not possible to canvass them exhaustively here (but see Huang, this volume). Instead, I'll introduce two promising possibilities.⁸ Before I do so, however, I want to issue a warning. One might say “That's ellipsis” about an apparently non-sentential speech act, and simply mean: The speaker produced words that mean less than the complete thought that she communicated. Thus, put pre-theoretically, she “spoke elliptically.” To anticipate: I think that this is true, but that it doesn't help to reject the appearances. To explain why, consider a comparison. If I say “Postal, Ross, Lakoff, et al.,” one might describe me as “speaking elliptically” about early Generative Semanticists. Or, if I say, “John has finished,” meaning that he has finished eating dinner, I do not explicitly say what it is that he has finished – hence one might again describe this, pre-theoretically, as “speaking elliptically.” But in neither of these cases is it imagined that somehow, under my breath as it were, I really uttered more linguistic material than what one hears on the surface. I did not utter the phrase *early Generative Semanticists*, nor did I utter the phrase *eating dinner*. The only sense in which I spoke elliptically is that I let the hearer fill in contextually available information for herself – I did not produce a special “elliptical” expression.

But now, if this sort of thing is all that is meant by the term *ellipsis*, then one cannot explain away the appearances described in section 1 by appealing to ellipsis. Instead, in saying “That's ellipsis,” one is at best redescribing those appearances, using alternative vocabulary. No doubt the agent “spoke elliptically” in this extremely weak sense: the very description of the appearances highlights (1) the content “mismatch” between the expression used and the complete thought that the speaker of that expression meant; and (2) the key role of non-linguistic context in filling the resulting gap between expression-

meaning and speaker-meaning. But to capture this fact by saying that the agent "spoke elliptically" is not to grant that the appearances are mere appearances. In particular, saying this is not to deny that the agent really did produce a word/phrase and really did not produce a sentence; on the contrary, it's to presuppose that the appearances reflect what is genuinely going on, in that a word/phrase was produced. Here is the warning, then: In what follows I will never use the term *ellipsis* in the extremely weak sense of a speaker meaning more than her words mean.⁹ Thus, both varieties of ellipsis described below are designed to explain away non-sentential speech acts by not granting the existence of genuinely non-sentential speech acts.

Having issued my warning, now consider what I'll call "the fundamental feature of ellipsis." In ellipsis, the sound produced by the speaker is abbreviated *vis-à-vis* the message encoded, but the hearer can recover the complete message because the abbreviated sound somehow linguistically encodes that message. Thus, when a speaker means more than what his words mean that isn't a case of ellipsis, as intended here – essentially because the "more" that he meant is not linguistically encoded; instead, it is supplied by other means. Given the fundamental feature – i.e. abbreviated sound produced with complete meaning encoded – an issue that immediately arises is: how can this occur? There are at least two ways, as I'll now explain. (Readers already familiar with linguistic theories of ellipsis may wish to merely skim what follows.)

It is a truism of linguistic theory that sound patterns don't directly correspond to meanings. Rather, there is an intermediate level, syntax, the level at which words combine to make phrases, and phrases combine to make sentences.¹⁰ What a sound pattern immediately corresponds to, on this view, is a syntactic structure; that syntactic structure, in its turn, corresponds to a semantic content. That is:

(6) Sound pattern → Syntactic structure → Semantic content

Given this tri-level picture, it's easy to see the two ways that a sound pattern can end up seeming "abbreviated" *vis-à-vis* the complete message linguistically encoded. There can be something irregular about the sound/syntax correlation, giving rise to the apparent shortening, or there can be something irregular about the syntax/semantics correlation. This yields two quite different notions of ellipsis, namely **SYNTACTIC ELLIPSIS** and **SEMANTIC ELLIPSIS**, respectively. Put in terms of (6), on one view ellipsis has to do with abbreviation occurring between level one (sound pattern) and level two (syntactic structure): this is syntactic ellipsis. (Some might equally call it "phonological ellipsis," but I'll use the more traditional term.) On the other view, ellipsis has to do with abbreviation occurring between level two (syntactic structure) and level three (semantic content): this is semantic ellipsis.

On the first kind of ellipsis, the correspondence between the sound produced and the syntactic structure of the utterance is not perfectly ordinary, in the sense that a comparatively "short" sound gets mapped onto a comparatively

“long” syntactic structure (but the subsequent mapping from syntax to semantics is perfectly normal). VP-ellipsis is typically considered to be an example of this. Thus when a speaker pronounces the sound pattern /*John wants a dog, but Jane doesn't*/, this is short when compared with the corresponding syntactic structure, which (as a first pass) one might represent as in (7):

(7) [_{S1}[_{S2} John [_I [_{VP} wants [_{NP} a dog]]]]] but [_{S3} Jane [_I doesn't [_{VP} want [a dog]]]]]

Specifically, (7) contains two VPs, whereas the corresponding sound contains a sound pattern for only one of them; the second VP is simply not pronounced.¹¹

According to the second notion of ellipsis, the correspondence between the sound and the associated syntactic structure is perfectly ordinary, but the mapping from syntax to semantics is non-standard in that the comparatively short syntactic structure gets mapped onto a comparatively long semantic content. An example of such semantic ellipsis might be the word *cheers*. It is simply a word – specifically the noun [_N cheer] in the plural form – pronounced in the usual way. This ordinary word could be used in a sentence as follows: *I heard cheers coming from the stadium, but I didn't know why*. But, someone might say, this perfectly ordinary plural noun, with its perfectly standard pronunciation, sometimes exhibits a special meaning – for instance, when it is said in a pub as glasses are raised. (What is produced in that circumstance might be called the “one-word sentence” *Cheers!*) Still, they could add, what the sound maps onto syntactically speaking is a plain old word. So, goes this line of thought, the connection between level one and level two is perfectly ordinary: what's odd is the mapping from level two, the plain old word, to level three – the special meaning.

One might wonder: if the speaker has simply used a word, with its ordinary pronunciation, what is meant by calling it elliptical? Just that, semantically speaking, this ordinary syntactic item “has a special meaning,” here in the sense of being conventionally assigned a special use. (One might be tempted to claim similar things about *Out!* as uttered by a baseball umpire.)

How is semantic ellipsis different from syntactic ellipsis? Here is a heuristic for seeing the difference. As Sylvain Bromberger pointed out to me (in conversation), it doesn't make sense to ask about utterances of *Cheers!*: “What verb was uttered? What was the subject of the sentence?” There patently was no verb – and there was no grammatical subject. The sound that the drinkers produced corresponded to an ordinary word, *cheers*. Compare, however, the previous example of VP-ellipsis: “What was the second VP in the utterance of *John wants a dog, but Jane doesn't*?” is a perfectly reasonable question. The answer is fairly straightforward as well: the second VP was [_{VP} want a dog].

Let me sum up this section. I began by considering whether a genuine speech act is performed when someone speaks subsententially. It seemed likely that, at least in many instances, that is the case. Given this, one cannot reject the appearances by saying that speakers who speak subsententially never perform speech acts. I then considered various senses of “ellipsis.” I put aside an

extremely weak notion, in which the agent "speaks elliptically." It is of course true that the appearances presented in section 1 involve the agent "speaking elliptically." But to say this is not to reject the appearances: it is not being denied that the agent produced only a subsentence; nor is it denied that a genuine speech act resulted. (For example, as I described the case, Leah made an assertion of a complete proposition, even though her words encoded neither a force indicator nor a semantic content of type $\langle t \rangle$. Thus, she meant more than what her words, even contextualized, meant.) If the appearances are to be rejected by appeal to ellipsis, a much stronger notion of ellipsis is required, one according to which the thing produced is actually, in some sense, genuinely a sentence after all. There are, I noted, two possibilities. Either the thing produced is syntactically a sentence (with the associated meaning of a sentence), though it doesn't sound like a sentence, or it is not syntactically a sentence, but the subsentential expression produced nevertheless has the content of a sentence. These two may be summarized as follows:

- (8) **SYNTACTIC ELLIPSIS:** Using an abbreviated sound pattern that corresponds to what is syntactically a more complete structure, which structure then linguistically encodes the complete message recovered.
- (9) **SEMANTIC ELLIPSIS:** Using an abbreviated sound pattern that corresponds to an equally abbreviated syntactic structure, but where that syntactic structure somehow linguistically encodes the complete message recovered.

Notice that both options respect what I called the "fundamental feature" of ellipsis. According to both, the sound pattern produced by the speaker is abbreviated *vis-à-vis* the complete message linguistically encoded. Where they differ has to do with how (i.e. where) the abbreviation occurs: the former has a non-standard mapping from sound to syntactic structure (but a standard mapping from syntax to meaning), the latter has a non-standard mapping from syntactic structure to meaning (but a standard mapping from sound to syntax). Having these two options on the table, we must now consider whether they successfully "reject the appearances" described in section 1 by explaining them away.

3 Challenges to the Ellipsis-based Rejection of the Appearances

Recall two of the examples introduced in section 1. Leah used *Sam's mom* on its own to say about the woman entering the room that she is Sam's mom. And Anita used *Moving pretty fast* in isolation to assert that a certain stock was moving quickly, in terms of its rising price. The syntactic ellipsis story, applied to these examples, would proceed as follows: Leah and Anita did indeed perform genuine speech acts, specifically, each made an assertion, but they did

not perform genuine non-sentential speech acts because in both cases what they produced were sentences, reduced via syntactic ellipsis. What the sounds corresponded to, the proponent of syntactic ellipsis could say, were the full syntactic structures:

(10) [_S That [_I is [_{NP} Sam's mom]]]

(11) [_S That stock [_I is [_{VP} moving pretty fast]]]

And, of course, these structures linguistically encode the complete propositions asserted by Leah and Anita, respectively. So this explains both why the message recovered was fully propositional, and why it appeared that the speakers had produced mere phrases. (It appeared that way because the utterances did indeed **sound** just like phrases.)

The syntactic ellipsis gambit looks initially promising, but it is an open empirical question both whether it is a correct view of the cases at hand and whether it can be extended to handle a broader range of (apparently) non-sentential speech. My own view – argued for in Stainton (1997b, 1998a, in press), as well as in Elugardo and Stainton (2001a, in press) – is that it cannot. My aim in this paper, however, is simply to introduce the issue of non-sentential speech, and not to settle all the outstanding disputes. So I will rest content with explaining the syntactic ellipsis approach and noting a couple of the obstacles that it faces.

There are clear disanalogies between, say, VP-ellipsis and the sort of speech described in section 1. In particular, VP-ellipsis requires the presence of explicitly spoken linguistic material in prior discourse.¹² Indeed, this prior material must be such as to allow a purely grammatical reconstruction of the elided material.¹³ Notice, for instance, that the elliptical sentence *Jane doesn't* typically cannot be used in DISCOURSE INITIAL POSITION. Nor can it be used in the middle of a discourse but without appropriate prior material. (Thus, imagine that Mary-Liz shows Paul a photograph of her daughter, in which the daughter, Karen, is smoking. Mary-Liz says, "Karen at school," and sighs. Paul cannot grammatically reply with *My daughter doesn't*, though he might well get his message across by so speaking. The reason is that there is some prior linguistic material, namely *Karen at school*, but there isn't the right sort of material – the sort that would allow a straightforward grammar-driven reconstruction of *My daughter doesn't smoke*.) VP-ellipsis is not a matter of the agent "guesstimating," on an all-things-considered basis, which material was omitted. Rather, it is a highly constrained algorithmic process that uses the "fragment" spoken and appropriate prior linguistic material to recover the source from which the fragment was derived. In contrast, as Barton (1990) first stressed, prior linguistic material is not required for the use of non-sentences. (Notice, in this regard, that both Leah and Anita's speech acts initiated a conversation in the examples described at the outset.) And understanding what is meant is very much a matter of the agent, drawing on all her relevant resources, "figuring out" from

the non-linguistic context what the speaker might be talking about. It is not a wholly grammatical process, and it is very flexible about the sorts of linguistic material (if any) that precede it.

Another disanalogy is this: In VP-type ellipsis, the elided element is always a SYNTACTIC CONSTITUENT. One cannot choose to omit any old sequence, up to recoverability of the intended meaning. But look again at (10) and (11), the purported sources for Leah and Anita's utterances. What would have to be left out of them are, in each case, non-constituents – specifically *That is* and *That stock is*. Similarly, in familiar kinds of syntactic ellipsis, the remnant expression reflects the GRAMMATICAL CASE that the words would have had in the full sentence. This is not, however, a true generalization for non-sentential speech: e.g. in English one says "Me, me!" to get across that one wants a free T-shirt, though the corresponding sentence is not "Me want a free T-shirt" but rather "I want a free T-shirt." (Similar points apply to more richly case-marked languages such as German and Korean. See Stainton, in press, for extended discussion.)

There is, at present, no consensus about whether syntactic ellipsis can explain away the appearances. It does seem that if syntactic ellipsis is occurring, the variety of ellipsis in play is not one that can be assimilated to other, better understood, varieties of ellipsis. For example, it would be quite unlike VP-ellipsis (*Juan doesn't __*), sluicing (*I wonder why __*), or gapping (*Lucia lives in Spain and Alain __ in France*). Still, there are many possible moves that a proponent of syntactic ellipsis could make.¹⁴ Rather than pursue the question further here, however, I will turn to how a semantic ellipsis account might work.

The idea, recall, is that the sound–syntax correspondence is perfectly normal, so that when it sounds like the agent has produced what is syntactically a word/phrase, that is precisely what she has produced. However, the syntax–semantics mapping is non-standard, since the word/phrase has a special meaning. This is why the message encoded is a complete proposition, even though the sound produced is abbreviated *vis-à-vis* that encoded proposition. Applied to the (apparent) non-sentence cases noted above, the idea would be that when Leah produces the sound /*sam's mother*/, she really did produce an NP. No subject, no auxiliary verb, no tense here: syntactically, what Leah uttered was just a single phrase. But this NP has a non-standard meaning. It is of semantic type ⟨t⟩, not of type ⟨e⟩.¹⁵ Hence, in this crucial sense, the thing uttered wasn't really non-sentential; semantically at least, it was sentential. (Optionally, it might be added that the NP had, as part of its content, a force indicator.) Similarly, the proponent of semantic ellipsis will say that Anita really did utter the bare VP [_{VP} moving pretty fast] but that this VP was assigned a non-standard meaning, namely a sentential meaning. This is what explains both what was asserted, namely the propositional content of the one-phrase sentence, and why it sounded like a mere phrase – because it was a mere phrase, though only syntactically speaking. This explanation shows, then, that what appeared to be non-sentential speech acts, though they are speech acts, aren't really non-sentential.

How promising is this approach? Well, appealing to semantic ellipsis may allow one to explain away certain isolated cases of apparently non-sentential speech. For instance, when the baseball umpire yells "Out," he clearly performs a speech act. But one can, it seems, reasonably reject this as a case of a non-sentential speech act – by noting that *out* here has a special meaning. What the umpire produced was syntactically a particle, the story would go, but it doesn't have the same meaning as when it appears embedded in complete sentences. Similarly for *Exit* signs, or cries of "Fire!" in crowded theaters: in each of these cases it's at least plausible to suppose that there is a special convention that establishes a special use, and hence a special meaning (including even a force indicator). However, it surely isn't plausible that every phrase in the language has two meanings, one that it exhibits embedded and one that it exhibits when used on its own, unembedded. As will emerge below, this would be far more meanings than one really needs to explain the speech observed. But such a duplication of meanings is precisely what the semantic ellipsis approach requires, if it is to be a general solution. After all, performance limitations aside – e.g. eventually speakers run out of breath – any NP can, in the right circumstances, be used to perform a speech act. Ditto for every VP, PP, and so on. So each of them would have (at least) two meanings.

Let me stress: what would have to be postulated is a genuine ambiguity, rather than mere context sensitivity. Everyone grants that *She is a professor* can be used to make quite distinct statements, but it would be naïve indeed to conclude from this that the sentence type is ambiguous. What happens in this latter case is just that the meaning of the type – its character, to use David Kaplan's phrase – gets completed by having reference assigned to the indexical *she*. Notice, however, that the sentence-type *She is a professor* is such that the meaning of each token is always a proposition. No token ever stands for an object, or for a property. What varies is simply who the resulting proposition is about. This is not a difference in logical structure. Similarly then, if the phrase-type *moving pretty fast* were univocal, all tokens of it would share the same logical structure: they would all express properties. So if some token of *moving pretty fast* expresses a proposition, this must be because the expression type is ambiguous.

Here is a comparison. There is a rock band called *Better than Ezra*, and so one can say things like, "Better than Ezra will be playing the hockey stadium on Friday." Given this, some tokens of the (sound) type */better than ezra/* actually denote an object, namely the rock band, while others express the property shared by all things that are better than Ezra. As a result of this naming, the expression itself became not just context sensitive, in the sense in which *She is a professor* is context sensitive, but ambiguous: it became both a complex name and a one-place predicate, exhibiting both semantic type $\langle e \rangle$ and semantic type $\langle e, t \rangle$. Now, in the same vein, suppose that the sound */better than ezra/* also corresponded to what is semantically a sentence, as per the semantic ellipsis hypothesis now being considered. (This would account for why, for example, one may point at a surfer, exclaim */better than ezra/*, and

thereby say of that surfer that she is a better surfer than Ezra.) To capture this in the same way that we capture the use of */better than ezra/* as a name, the sound type */better than ezra/* would have to exhibit a third kind of meaning, a propositional meaning. Here is why. The proposition about the surfer could not result merely from filling in any indexical slots in the character of the one-place predicate [_{AdjP} better than Ezra], since that would always yield a property, not a proposition; nor could slot-filling applied to the band-name version, [_{NP} Better than Ezra], result in this proposition. Such slot-filling would (vacuously) yield the rock band as referent. Thus, if slot-filling is to yield a proposition, then the sound must have as one of its meanings something of type $\langle t \rangle$. That is, it must share the semantics of the sentence type *He/she/it is better than Ezra*. Thus the sound type would need to have three conventionalized meanings on this proposal.

Returning to the original cases, it really isn't plausible that *moving pretty fast* and *Sam's mom* have two standardized meanings, one of type $\langle t \rangle$ and the other **not** of type $\langle t \rangle$. (Where it is the latter, subpropositional meaning that the word contributes to the truth conditions of larger wholes, e.g. in *Sam's mom is in her car, and it is moving pretty fast*. Clearly **here** neither *Sam's mom* nor *moving pretty fast* contribute a proposition, even once contextualized.) So this method of rejecting the appearances is quite unpromising. Though, to repeat, it may allow one to explain away some apparent cases: *Out* as said by the umpire, for example.

Some theorists remain hopeful that the above strategies, possibly taken in combination, will allow one to "reject the appearances." Stanley (2000), for instance, proposes (very roughly) to treat some cases as not genuinely speech acts, some as syntactically elliptical, and some as semantically elliptical. He expresses the hope that such a divide-and-conquer approach will ultimately cover all cases of apparently subsentential speech acts. Others think this hope is in vain (see for example Clapp 2001). Unfortunately, the issue is too complex, and the relevant empirical questions too unsettled, to know for sure who is right, so I leave the issue here.

4 Accepting the Appearances

Suppose, for the sake of argument, that neither special syntax nor special semantics is enough to account for the appearances. It would, then, surely fall to pragmatics to explain how subsentential communication can succeed. (Notice that, for the pragmaticist, the aim would be to explain how, not to explain away.) The general strategy should be obvious: The person who speaks subsententially cannot be trying to assert an object, or a property, or a property-of-properties, or anything else non-propositional. Even if we could make sense of what that would be, doing it would not carry forward the talk exchange in an acceptable way. Thus, the hearer would inevitably recognize that the speaker must mean more than what his words mean (and the speaker

will presumably intend for the hearer to recognize this, and so on). The hearer will, therefore, use all the evidence at her disposal to find out what the speaker likely did mean.

This general story seems plausible, at least to me. The question is, what exactly happens during this process? In what follows, I sketch one possible story.¹⁶ It is uncomfortably speculative, but it gives a sense of the sort of tale a cognitively oriented pragmaticist might tell. To tell the story, however, I need to sketch by way of background two quite different conceptions of linguistic interpretation.

According to one way of thinking about language, a language is a system of shared rules – a complex algorithm – that, in advance of occasions of interpretation, determines the meaning of utterances. True, the algorithm needs to take as inputs not just the form of the thing uttered, but also a (quite limited) set of CONTEXTUAL PARAMETERS. But once speaker, addressee, time, and place are specified, employing a language is simply applying the algorithm in question: the speaker employs the algorithm, and nothing else, to “encode” her thought; the hearer employs the self-same algorithm, applying it to the form of the thing uttered and the contextual parameters, to decode the thought. End of story. Because the algorithm assigns meanings compositionally, and because the composition is recursive, knowing it explains the ability of speaker and hearer to interpret an (in principle) unlimited number of novel utterances in a systematic way. Taking a leaf from Davidson (1986: 437–8), who later came to have serious reservations about this approach, the idea is this:

You might think of this system as a machine which, when fed an arbitrary utterance (and certain parameters provided by the circumstances of the utterance), produces an interpretation. One model for such a machine is a theory of truth, more or less along the lines of a Tarski truth definition. It provides a recursive characterization of the truth conditions of all possible utterances of the speaker, and it does this through an analysis of utterances in terms of sentences made up from the finite vocabulary and the finite stock of modes of composition. I have frequently argued that command of such a theory would suffice for interpretation.

It's worth stressing that “an interpretation” here means, at a minimum, an assignment of truth-evaluable content, at least in the case of declarative speech. So what the “machine” does is to take in sounds (plus some highly restricted set of contextual parameters for assigning reference to explicit indexical elements – i.e. speaker, addressee, time of utterance, place of utterance) and output something propositional. For ease of reference, let's call this the ALGORITHM IS SUFFICIENT (AIS) conception. On the AIS conception, employing a language emphatically is not a process of “guesstimating” what a specific utterance literally means. Guessing may play a role in understanding metaphor and conversational implicature, but it has no place in literal comprehension. True enough, it is creative – in the sense of being generative – but it is not creative in the sense of requiring cleverness and imagination.

A rather different way of thinking about language, a way that I personally find much more plausible, goes like this: A language is indeed a system of shared rules, a complex algorithm. Moreover, here too the algorithm is considered to be compositional and recursive. But the algorithm, though it is necessary, is not anything like sufficient for interpretation: the algorithm does not, in advance of occasions of interpretation, determine the meaning of literal utterances all on its own. Specifically, the algorithm often does not – even given the aforementioned contextual parameters – inevitably assign something propositional to the utterance. Rather, it (often) assigns something that must be COMPLETED or ENRICHED to arrive at something truth-evaluable. Thus, although knowing English is required for understanding English speech, it isn't enough – not even when supplemented by knowledge of the highly constrained contextual factors like addressee, time of utterance, etc. It is this second non-AIS conception of interpretation that will play a key role in the positive, pragmatics-based story about how subsentential communication works that I will present at the end.¹⁷

Employing a language, so conceived, involves not only the process of applying the algorithm, but some other process as well. It is the second process that does the enriching. Of particular interest recently is the idea that the additional process is one of drawing general-purpose inferences to arrive at all-things-considered judgments about what the utterance meant. (The idea being not that all things have been considered, which is impossible given the time constraints, but that anything that the person knows is relevant in principle.) Employing a language would thus involve two quite distinct processes, neither of which is individually sufficient for discovering the meaning of the utterance. The first process is algorithmic, but the second is not: It is non-deterministic inference. (Both processes are "creative," but in quite different senses: the former is creative in the sense of involving a generative procedure; the second is creative in the sense in which an artistic creation is.)

One way of thinking about the different conceptions is psychologically. (I expect, however, that proponents of the AIS conception would not themselves think of the difference this way.) The second conception is MODULAR, in something like the sense of Fodor (1983). In particular, the psychological concomitant of the non-AIS-conception goes like this: There is a language-specific subfaculty, in which the algorithm is stored; and there is another component of the mind-brain, call it the "Central System," where inferences get drawn; and both of these subfaculties play a role in both the production and comprehension of speech. (The second subfaculty plays a role in much else besides, of course, including the interpretation of non-linguistic communication. It is, I repeat, not language-specific.) In contrast, the psychological concomitant of the first conception would have to be either that there are not mental subfaculties at all or that there are such subfaculties, but that exactly one is employed in speech and interpretation. This would be the way of cashing out the idea that there is only one process at play – be it in "the mind as a whole" (the no-subfaculties version) or in exactly one subfaculty thereof. Put

in COMPETENCE-PERFORMANCE talk, the two conceptions would see things this way: For the non-AIS-conception, more than one mental competence plays a causal role in yielding observed performance; for the AIS-conception, exactly one competence yields performance. (This may, in fact, lead one to prefer doing away with the competence-performance distinction altogether.)¹⁸ In a nutshell, on the non-AIS-conception the psychology looks like this: there is a specifically linguistic process of decoding the signal to arrive at a representation of its meaning. This process has to do with what the expression spoken – the expression type, that is – means. That decoding is carried out by a task-specific module called the language module. On the other hand, there is a process of general-purpose inference, drawing in principle on all the information at the disposal of the agent (not, to repeat, in the sense of the agent actually accessing all of her information; rather, in the sense of all of it being in principle relevant). This second process is crucial not only for determining what the speaker meant, above and beyond what he said, but also for determining what the speaker literally asserted/stated/said.

Having drawn the contrast, I now want to apply the non-AIS conception to the understanding of non-sentential speech. Let's start with the complete sentence case, by way of introduction. On the picture I am assuming here, when a complete sentence is spoken, the linguistic decoder does not typically yield a truth-evaluable mental representation. Two processes, decoding and inference, are required even here. Still, when given a complete sentence the decoder does yield something – let's call them PROPOSITIONAL FORM SCHEMAS – which, once "fleshed out," are truth-evaluable. What "fleshing out" means here is: all indexical slots are filled, all vague terms are sorted out, and all ambiguities are resolved – but nothing more. That is, the process in the complete sentence case is as follows: the decoder outputs a propositional form schema, the Central System then fleshes out this schema (i.e. assigns reference to indexical slots, disambiguates, and sorts out vague terms), and these two processes together yield a propositional form.

Crucially, this is not what occurs in subsentential speech, or so I have argued elsewhere. In non-sentence cases, what the decoder yields is not even a propositional form schema, let alone a propositional form. Instead, it outputs a representation that, even after fleshing out, represents an object, or property, or property-of-properties. (For instance, in the *Sam's mother* case, we might suppose that what decoding-plus-fleshing-out provides is a mental representation that applies uniquely to the person who is the mother of the contextually salient Sam.) The mental representation produced by the decoder in subsentential speech cases must be altered to arrive at something which is truth-evaluable – i.e. to arrive at a propositional form; this much is the same. But the kind of "alteration" is importantly different: the decoded representation is not altered merely by fleshing it out, i.e. by assigning reference to existing slots, disambiguating, and sorting out vagueness. Rather, it is altered by **conjoining the bit got from decoding-plus-fleshing-out with another representation entirely**. This is what is distinctive about subsentential communication.

Where does that other representation come from? The sources don't seem any different than in the case in which a referent must be assigned to an indexical. Thus, the representation may come from long-term memory, from short-term memory, from the imagination, from a perceptual module, etc. What's different is how the sources are employed: they don't just help to assign a referent to a pre-existing natural language expression; rather they provide a whole new, not-in-spoken-language piece of the whole mental representation. Thus, returning to the original case, Leah and Anita could both see the woman in the doorway. Anita could, therefore, combine a mental representation from the visual faculty, which denotes the woman, with the mental representation resulting from the decoding/fleshing out. This would yield a sentential mental representation – not a sentence of English, mind you, but a propositional form in a mental language. In different examples, the "other representation" could come from memory or many other sources. The key point is that it does not come from the language module. This isn't so surprising, when all is said and done – at least once it's agreed that linguistic interpretation involves not one process (e.g. an algorithm that blindly takes structure-plus-contextual-factors to interpretation), but two processes, where the second process involves drawing inferences using all information available to the agent. If linguistic interpretation works like this, it is understandable that the inferential process could bridge the gap to pragmatically yield a complete proposition.

Notice, in conclusion, that this story does indeed treat the appearances as reflecting what is really going on. Speakers do indeed routinely produce subsentential expressions – nouns and NPs, verbs and VPs, PPs, etc. – and that is precisely what the hearer decodes. Moreover, the expression types produced are not just syntactically subsentential, but semantically subsentential as well: they are not of Montagovian semantic type $\langle t \rangle$, and they do not have force indicators as part of their content. And yet, in speaking subsententially, speakers really do perform fully propositional speech acts. This is possible because the hearer can understand the proposition meant by employing not only information got from linguistic decoding, but also information from a host of other possible sources.

Let me end this section with some questions. They are, I think, among the most pressing issues facing the pragmaticist who wishes to explain, in this sort of way, how subsentential speech works. The questions are:

- 1 How precisely does the representation of the object/property/etc., which doesn't come from linguistic decoding, get combined with the part that does? For example, how does the hearer know which object/property/etc. representation to use from memory, vision, etc.?
- 2 What is the representational format like, such that it is conducive to the combination of (i) something from linguistic decoding and (ii) a non-linguistic representation – retrieved from memory, or created in the Central System, or deriving from a perceptual module?

- 3 How do perceptual and other non-linguistic representations get into this integration-conducive format?

These questions are pressing, not just because they are inherently interesting, but because if they cannot be answered, then one is pushed back toward the "reject the appearances" approach. Depending on how one takes them, these questions either indicate directions for exciting further research, or they pose possibly devastating objections. It will be noticed, for instance, these sorts of questions either do not arise, or are easily answered, on the ellipsis approaches. Some attempt to address these questions within a non-AIS conception can be found in Carberry (1989), Stainton (1994), and Elugardo and Stainton (in press). I leave them open here.

5 Summary

I began the paper with a couple of examples of non-sentential speech, also providing a more general description of what appeared to be going on in such examples. What appeared to be going on was this: Speakers utter things that are, both syntactically and semantically, subsentential, but they nevertheless manage to perform genuine speech acts (e.g. asserting) in so speaking. (And this isn't a matter of replying to an interrogative, or repairing a prior utterance.)

Having described the appearances, I presented two broad strategies for rejecting them. The first was to deny that a genuine speech act was being performed. This might work for some isolated cases, but it seems unpromising as a general approach. The second broad strategy was to deny that the speech episodes in question really were non-sentential. To make this latter idea plausible, it was necessary to introduce the notion of "ellipsis." I thus described three senses of ellipsis. Only two of them were robust enough to actually support the rejection of the appearances, however; the weak sense of ellipsis, it will be recalled, was essentially a redescription of the appearances, rather than a rejection of the reality of non-sentential speech acts.

The first "strong" variant, syntactic ellipsis, involved treating (apparently) subsentential speech as rather like VP-ellipsis. The second variant, semantic ellipsis, involved positing the existence of "one-word sentences" and the like with special propositional meanings (and, optionally, force indicators). Both variants of the ellipsis strategy, it emerged, face important obstacles, and only further research will tell us whether some version of the ellipsis story can succeed.

After exploring a few avenues for rejecting the appearances, I briefly sketched a pragmatics-based approach that accepts that the phenomenon is genuine. The key background presupposition for this cognitive-pragmatic approach was that there are two processes involved in speech-understanding, namely decoding and unencapsulated inference. Because of the latter, mental representations from many different sources can be brought together with mental representations derived from linguistic decoding. On this story, then, what

occurs in subsentential speech cases is this: The linguistic decoder works on the (genuinely) subsentential utterance; it outputs a subsentential mental representation; and this gets combined with another (non-decoder-derived) mental representation to yield a sentential mental representation, which is not in any natural language. This latter encodes the complete message meant by the speaker.¹⁹

NOTES

- 1 For discussion of "completing fragments" of this sort, see Morgan (1973c) and Shapley (1983).
- 2 Actually, according to some recent syntactic theories *Sam's mom* is not an NP. Rather, it is a Determiner Phrase (DP), whose grammatical head is not a noun but a determiner. I'll ignore such issues for present purposes. For an early discussion of the use of NPs in isolation, see Yanofsky (1978).
- 3 If the jargon is unfamiliar, the following gloss will do: Speakers may utter any phrase, including in particular phrases whose "core" is an ordinary word. These latter are the lexical projections. Such phrases contrast with items whose "core" is a tense marker or an agreement feature, the latter two being the "core" of sentences. Examples of lexical projections include [_{NP} a dog], [_{PP} from Brazil], [_{AdjP} red], etc.
- 4 This notation is introduced and explained in detail in Dowty et al. (1981: Chap. 4). For present purposes, however, the following will do: Once reference is assigned to all indexicals, an expression of semantic type $\langle e \rangle$ denotes an individual; an expression of semantic type $\langle e, t \rangle$ denotes a function from an individual to a truth value (equivalently, a set of individuals); an expression of semantic type $\langle t \rangle$ denotes a truth value; and an expression of semantic type $\langle \langle e, t \rangle, t \rangle$ denotes a function from a set to a truth value (equivalently, a set of sets of individuals). Note that, for the sake of simplicity, I am ignoring intensions throughout.
- 5 My thanks to Rebecca Kukla for the point, and for the example.
- 6 There are many possible implications of subsentential speech, implications about Frege's (1884) context principle, the domain of logical form, external determinants of content, the relationship between thought and "inner speech," the pragmatic determinants of what is said, etc. For extensive discussion of these possible implications, see: Dummett (1973, 1981, 1993); Stainton (1995, 1997a, 1997b, 1998a, 1998b, in press); Carstairs-McCarthy (1999); Kenyon (1999); Stanley (2000); Clapp (2001); and Elugardo and Stainton (in press).
- 7 For instance, imagine Leah is often late for dinner. Randy is giving Leah a familiar nasty look, as the latter goes out the door. She responds: "Seven o'clock. Without fail." Here Leah seemingly makes a promise to Randy to be home by seven o'clock.
- 8 Work on ellipsis exists in a host of distinct frameworks. For examples, see: Shopen (1973); Halliday and Hasan (1976); Sag (1976); Williams (1977); Chao (1988); Carberry (1989);

- Dalrymple et al. (1991); Stainton (1995, 1997b, in press); Fiengo and May (1996); Lappin (1996); Lappin and Benmamoun (1999).
- 9 It's worth noting that some contributors to the literature do use "ellipsis" in this very weak way. One example is Carberry (1989). She thinks of the issues in terms of "understanding elliptical fragments," but she clearly means "elliptical" in the weak sense noted here, since her solution to the problem is not to explain it away (by appeal to elliptical sentences), but rather to explain how speakers/hearers, using pragmatic information about discourse goals and the speaker's plans, manage to communicate using genuine subsentences.
- 10 It might be more accurate to speak of MORPHOSYNTAX here, since morphology is part of the "intermediate" level too. However, to simplify the exposition I will continue to speak of syntactic structure. Actually, there are several quite different accounts of how the mapping between "short" sound and "long" syntactic structure is achieved. In particular, some accounts introduce not an unorthodox mapping, but rather an unorthodox resulting structure – one which contains, say, elements that have no pronunciation. (That is, the posited elements are in this respect like big PRO, or trace, in Chomsky's Government and Binding Theory.) Thus, regarding the present example, such a theorist might suggest that the corresponding syntactic structure is not (7) but:
- (i) $[_{S1}[_{S2}$ John $[_r$ $[_{VP}$ wants $[_{NP}$ a dog of his own]]]] but $[_{S3}$ Jane $[_r$ doesn't $[_{VP}$ Δ $[_{NP}$ Δ]]]]
- where Δ has no pronunciation at all. Williams (1977) offers a theory of this sort. For a useful survey of theories of ellipsis, see Chao (1988) and Lappin (1996). Good collections on ellipsis include Berman and Hestvik (1992) and Lappin and Benmamoun (1999).
- 12 Some apparent counterexamples can be found in Hankamer and Sag (1976) and Sag and Hankamer (1977). They discuss the sort of case in which, for example, a woman walks up to a cliff, and someone else says, "She won't." This looks like genuine ellipsis without appropriate prior material. For discussion of the implications of such cases for subsentential speech, see Stanley (2000) and Stainton (in press). The latter paper also goes into much greater depth about the nature of syntactic ellipsis and how it differs from non-sentence use.
- 13 In fact, the constraint is usually taken to be that deletion of material may only occur when the **identical** material occurs in the immediately preceding discourse. Put otherwise, only identical material counts as "appropriate prior discourse." For an early statement of the identity constraint, see Shopen (1973). Arguments for it, and refinements on the notion of "identical material," are provided in Sag (1976). But see note 12 for some possible hedges to this constraint.
- 14 For instance, it would be worth examining closely fragments apparently created by phonological deletion of initial material – especially of determiners, subjects, and auxiliaries. Examples include *Paper boy's here*, in which the determiner is left out, *Seems suspicious*, and *Guess I should be more careful*, in which subjects are omitted, and *Find what you were*

- looking for?* in which both the subject and the auxiliary are left out. For discussion of this interesting phenomenon, see Schmerling (1973) and Napoli (1982). They take it to be a quite superficial process, occurring at the interface of syntax and phonology, of the same sort that derives *'sgusting* from *disgusting*. Clearly not all subsentence cases could be assimilated to deletion of this sort, since what would need to be omitted would very often not be **initial** material, but some apparent cases of subsentential speech might be explained away in this fashion.
- 15 Or, if one wishes to be more of a Montaguean, the expression is not of type $\langle\langle e, t \rangle, t \rangle$. (Montague himself thought that all NPs, including proper names, actually stood for generalized quantifiers.)
- 16 A positive, pragmatics-based story is offered in rather more detail, and with supporting empirical evidence, in Elugardo and Stainton (in press); see also Carberry (1989) and Stainton (1994).
- 17 The AIS and non-AIS conceptions are not the only possibilities, of course: there are intermediate positions that assign the language-specific algorithm more or less weight in linguistic interpretation. The motivation for introducing the contrast this starkly is that it highlights how the accept-the-appearances strategy fits into a comprehensive view about language interpretation.
- 18 I realize that some will balk at reading psychological commitments into these conceptions, especially the first one. But, as I've repeatedly noted, I am not interested in exegesis here. My aim, rather, is to highlight a contrast by drawing it starkly. Hence, I will not try to pin the two conceptions, or their psychological concomitants, on particular philosophers. On the other hand, it's only fair to give credit where credit is due, so let me note that I have gleaned the non-AIS-conception of interpretation quite directly from conversations with, and readings of, Sperber and Wilson and other Relevance Theorists; see Sperber and Wilson (1986a) and Carston (1988). See also Recanati (1989) for related discussion. Speaking of giving credit where it is due, it is also worth noting that Strawson anticipated the possibility of stating something with a subsentential expression some 50 years ago. He writes, in "On Referring":
- There is nothing sacrosanct about the employment of separable expressions for these two tasks [i.e., the tasks of forestalling the question "What (who, which one) are you talking about?" and the task of forestalling the question "What are you saying about it (him, her)?"]. Other methods could be, and are, employed. There is, for instance, the method of uttering a single word or attributive phrase in the conspicuous presence of the object referred to; or that analogous method exemplified by, e.g., painting of the words "unsafe for lorries" on a bridge, or the tying of a label reading "first prize" on a vegetable marrow. (Strawson 1950: 303)
- 19 I am very grateful to the following for comments on an earlier draft: Ash Asudeh, Andrew Carstairs-McCarthy, Ray Elugardo, Rebecca Kukla, Zoltan Gendler Szabo, and Gregory Ward.