ROBERT J. STAINTON

QUANTIFIER PHRASES, MEANINGFULNESS "IN ISOLATION", AND ELLIPSIS[1]

1. Introduction

Semanticists and philosophers of language have traditionally divided expressions into two classes: those that do, and those that do not, have meaning "in isolation". Among expressions which are commonly said to *lack* meaning in isolation are particles in natural language (e.g. 'up' in 'shoot up', 'on' in 'catch on', etc.) and the logical vocabulary of both natural and artificial languages (e.g. the sentential connectives 'and', 'if', etc.). Assuming one can sort out what "having meaning in isolation" amounts to, the question arises: which expression types fall into which category? In particular, and this is the central question of the paper:

(1) *The Question:* Do quantifier phrases have meaning in isolation?

Here's how I plan to address this question. First, I want to get a fix on what "having meaning in isolation" amounts to. After that, I will argue that quantifier phrases, e.g. 'some woman' and 'eight cats', *do* belong in the class of expressions which have meaning in isolation. The central argument for this claim: quantifier phrases can be used and understood outside the context of any sentence. (For example, a man may approach an apple cart and say nothing more than 'Six large apples',thereby requesting six large apples.) But, I aim to show, expressions which lack meaning in isolation cannot be so used. So quantifier phrases have meaning in isolation.

Having said a little about what the Question means, and having argued for an affirmative answer to it, I will consider a possible response to my argument. That response, which has great initial appeal, goes like this: though quantifier phrases *appear* to be used and understood in isolation,

[1] This paper was written during a visit to Umass-Amherst, and revised while visiting Rutgers University. My thanks to my hosts and friends at both institutions, including especially Barbara Partee and Ernie Lepore, who invited me to their respective universities. An earlier draft was presented at McGill University. I'm grateful to everyone who attended, but must single out Stephen Neale, who happened to be at McGill that day. Unsurprisingly, Stephen made numerous, and very useful, comments. Thanks also to Andrew Botterell, Marie-Odile Junker, Barbara Partee (again), and two anonymous *Linguistics and Philosophy* referees, for criticism/commentaries. Finally, I'm grateful to my home institution, Carleton University, for sabbatical time.

Linguistics and Philosophy **21**: 311–340, 1998.
© 1998 *Kluwer Academic Publishers. Printed in the Netherlands.*

this is actually a matter of *ellipsis*. In the end, as I'll explain, I don't think an appeal to ellipsis can really support a negative answer to the Question. So my positive answer stands.

That's the game plan. Before continuing, I want to emphasize as strongly as I can that this paper is not intended as an exercise in the history of philosophy. Specifically, it is not an attempted explication of, and subsequent attack upon, any theses proposed by Bertrand Russell. It is true that, in trying to get a grip on what "meaning in isolation" amounts to, I will appeal to some of Russell's (1905, 1911, 1919) ideas. But, in the end, nothing I say hangs on whether, for example, the notion of "having meaning in isolation" I arrive at really derives from Russell. Or again, in arguing for a positive answer to the Question, I make no claim to have refuted any doctrine held by Russell. Our very different purposes, and our divergent use of terms, makes it hard for me to know whether my result conflicts with Russell's views. (Anyone who does know is cordially asked to pass the word along.)

2. Meaningful in Isolation: A Case Study

The basic problem with explicating the notion "meaningfulness is isolation" is this: what counts as "meaningful in isolation" greatly depends upon what meanings are. So, to move ahead on the Question, I need to figure out what meanings are. Obvious problem: there is, to put it mildly, rather substantial disagreement about what meanings are. Some philosophers think meanings are senses; others think they are Ideas; while still others take them to be (non-conceptualized) objects in the external world. Because of this, rather than trying to provide a general account of meaningfulness in isolation, I will try to illustrate, rather than define, the term 'meaningful in isolation'. If my discussion is sufficiently illuminating, its lessons can be applied to various theories of meaning, thereby yielding a notion of "meaningfulness in isolation" for diverse meaning-theories.

As my illustration, I choose Russell's direct reference theory, essentially because it involves a clear-cut distinction between expressions which are, and are not, meaningful in isolation. Roughly speaking – and recalling that this isn't exegesis – meanings for Russell are external entities. To borrow a happy phrase from Sainsbury (1979), they are *meaning-relata*. These external entities come in two flavours: names refer to *particulars*, while predicates refer to *universals*. Crucially, at least so far as subsentential expression go, these alternatives are meant to be exhaustive.

Given that there are, for Russell, only two kinds of meaning-relata, and that these correspond exclusively to names and predicates, it might seem

that all other sub-sentential expressions should be meaningless. But, of course, it would be absurd to say that every other linguistic item is *gibberish*. The solution, however, is not far to seek. One need only draw a three way distinction between: (a) expressions which have meaning "in the primary way" (in Russell's case, by having a meaning-relatum); (b) expressions which have meaning in some "non-primary" way; and (c) expressions which lack meaning altogether. We may then say that an expression has meaning in isolation only if it gets its meaning in the primary way. Any expression which lacks meaning altogether, or affects the meaning of whole sentences, but not in the primary way, has no meaning in isolation.

Next step: what is it to have meaning in the "non-primary" way? It is, I take it, a matter of making a meaning-difference to whole sentences – though not by the standard means. Applied to Russell's case, then, an expression has meaning in the "non-primary" way if it affects the meaning of whole sentences, but does not have a meaning-relatum. (Put otherwise: though the word/phrase alters the meaning of a sentence *S*, it does not contribute a *constituent* to the proposition expressed by *S*.)

3. The Russellian Approach and Quantifier Phrases

To sum up the previous section: putting aside truly meaningless words and phrases (e.g. 'madatrauts'), an expression is *not* meaningful in isolation if, instead of being paired with a meaning in the usual way, one must give a rule for generating the meaning of *whole sentences*, within which the expression occurs. (As I said, what "pair with a meaning in the usual way" comes to depends very much on what you take meanings to be – e.g. external objects, modes of presentation of objects, Ideas, or what have you.) With this in mind, I now return to (1).

(1) *The Question:* Do quantifier phrases have meaning in isolation?

In Russell's case, (1) has an reasonably straightforward answer. Quantifier phrases are neither names, nor predicates; and only names and predicates have meaning-relata. Hence, quantifier phrases are not meaningful in isolation. Instead, their meaning must be specified by giving a general rule – which determines what sentences containing them mean. An example: to give the meaning of 'Every cat', one could provide a rule like:

(2) 'Every cat' combines with ⌜is G⌝ to yield a sentence; that sentence is true if and only if, for every x, if x is a cat, then x is G.

Abstracting away from the restrictive predicate 'cat', one could give a meaning-contribution rule for ⌜Every F⌝. That rule might be:

(3) ⌜Every F⌝ combines with a predicate ⌜is G⌝ to yield a sentence; that sentence is true if and only if, for every x, if x is F, then x is G.

Generalizing still further, the approach as a whole may be capured by the following, where **Q** is any quantifier word, and the ellipses are completed by the rule appropriate to the particular quantifier word:

(4) *The Not-Meaningful in Isolation Approach:* ⌜**Q** F⌝ combines with a predicate ⌜is G⌝ to yield a sentence; that sentence is true if and only if . . .

Adopting the above as the semantic axiom for quantifier phrases amounts to giving a negative answer to (1). For, what (4) provides for each quantifier phrase is a method for calculating the meaning of whole sentences containing quantifier phrases, rather than giving a meaning-relatum for them. What I now want to ask is whether this is the right approach. My conclusion will be that, whatever its merits vis à vis artificial languages, as a perfectly general approach to quantifier phrases in natural language, (4) does not work. The reason is, (4) is only operative *when there is a predicate* ⌜is G⌝ for the quantifier phrase ⌜**Q** F⌝ to combine with. Lacking such a predicate, the rule simply does not apply. But, as a matter of empirical fact, quantifier phrases can be used and understood in the absence of any such "second predicate". Or so I'll suggest.

4. The Problem: Unembedded Quantifier Phrases

I pause to emphasize: about quantifier phrases outside sentences, (4) says nothing whatever. Given this, rule (4), at least as it stands, can (at best) capture the semantic contribution of quantifier phrases as they occur in sentences. This wouldn't be a problem, but for the fact that quantifier phrases can be used and understood in isolation.

Time to introduce some data.[2] Suppose I'm at a linguistics meeting, talking with Andy Brook. There are some empty seats around a table. I point at one and say, 'An editor of *Natural Language Semantics*'; I then

[2] In case imaginary examples worry you, take a cursory glance at a speech corpus. You'll see that quantifier phrases are very frequently used and understood in isolation. Also, non-sentential expressions of many kinds can be used to perform speech acts: Noun Phrases, Verb Phrases, Prepositional Phrases, etc. Indeed, Barton 1990 argues that any X^{max} can be used on its own. See Stainton 1994, 1995, 1997a, 1997b for examples, and discussion.

indicate another empty seat and say, 'Anyone from *Pragmatics and Cognition*'. Upon hearing these words, Andy forms the belief that the unoccupied seats are reserved for an editor of *Natural Language Semantics*, and for some representative of *Pragmatics and Cognition*, respectively. Another detail. The seats I pointed to are actually reserved for Emmon Bach and M. A. K. Halliday; and, as a matter of fact, they are not involved with these journals. I want to stress two things about this imagined situation. Point one: since, in the imagined situation, the seat I indicated first is not reserved for an editor of *Natural Language Semantics*; and since the second seat is not set aside for someone from *Pragmatics and Cognition*, I spoke falsely in uttering (5) and (6), below. I made a false statement.

(5) An editor of *Natural Language Semantics*
(6) Anyone from *Pragmatics and Cognition*

Point two. What I uttered, in the described situation, were two quantifier phrases. Neither time did I utter a sentence. (You might be tempted to say: "This *isn't* really a use of an unembedded quantifier phrase; it's a use of an elliptical sentence – in particular, an elliptical sentence which *contains* a quantifier phrase". I'll address this shortly.)

To repeat: the no-meaning-in-isolation approach – schematized in (4) – at best says nothing whatever about the unembedded use of (5), (6) and related cases; at worst, it says that a meaningful utterance of (5) or (6) on its own is impossible. The reason, as I said, is that (4) applies only where there is a "second predicate" available to combine with the quantifier phrase. But, crucially, when an *unembedded* quantifier phrase is used, no such predicate appears.

What to do? The obvious solution is to make quantifier phrases meaningful in isolation. This is what I propose. Importantly, however, I want to reject one means of doing this: assimilating quantifier phrases to the category of names, thereby pairing them with *individuals*. The problem with this, as Russell rightly stressed, is that logical puzzles, and bizarre ontological commitments, would thereby arise.[3] Nor is it plausible that quantifier phrases have ("ordinary", i.e. first order) universals as their meaning-relata. So, if quantifier phrases are to be meaningful in isolation, what's needed is a third kind of meaning-relata. So be it.

[3] E.g. what kind of particular could be denoted by 'a man'? And how can it be that 'a man is bald' is true, while 'a man is not bald' is also true? Is it, perhaps, that the bizarre individual denoted by 'a man' is both bald and not bald?

5. The Fix: Generalized Quantifiers

Plausible assumption, part one: quantifier phrases can be used and understood in isolation. Plausible assumption, part two: the no-meaning-in-isolation approach, suggested by (4), does not predict that quantifier phrases can be used and understood in isolation. In which case, an alternative to (4) is called for. Here it is, in a nutshell: quantifier phrases, whether within a sentence or unembedded, correspond to generalized quantifiers.[4]

A generalized quantifier, for the purposes of this paper, will be a function from sets to propositions. (See Lewis 1972, Montague 1974, and Barwise and Cooper 1981 for early work.) Two things deserve to be stressed about my usage. First, a generalized quantifier, as I use the term, is not a kind of expression: though *quantifier phrases* are linguistic items, *generalized quantifiers* are not; they're functions. Second, I'm treating generalized quantifiers in the Russellian spirit: not as functions from sets to *truth values*, but as functions from sets to *propositions*. (The reason? I want to be able to distinguish 'a king of France', 'a unicorn' and similar unsatisfied quantifier phrases from one another.)

The generalized quantifier corresponding to 'some nitwits', for example, is that function f from sets to propositions such that, for any set S, $f(S)$ is a true proposition iff the intersection of the set of nitwits with S is non-empty. And the generalized quantifier corresponding to 'every toadstool' is that function g from sets to propositions such that, for any set S, $g(S)$ is a truth iff the set of toadstools is contained in S. Applied to a sentential example, sentence (7) is true iff the intersection of the nitwits with the smokers is non-empty – i.e. iff something is both a nitwit and a smoker; and sentence (8) is true iff the set of toadstools is contained in the set of broken things.

(7) Some nitwits smoke

(8) Every toadstool is broken

Having the notion of a generalized quantifier at hand, I can now lay out an alternative to (4).

(9) *The Meaningful in Isolation Approach:* ⌜**Q** *F*⌝ denotes the function f from sets G to propositions such that $f(G)$ is a true proposition if and only if . . .

---

[4] It's sometimes supposed that the generalized quantifier approach and the syncategorematic approach are mere notational variants of one another. If I'm right, this cannot be so – because they have distinct empirical consequences *vis-à-vis* the use and comprehension of unembedded quantifier phrases.

Notice: because quantifier phrases do not, on this approach, denote individuals, the aforementioned logical puzzles and weird ontology are avoided; nevertheless, it's worth stressing, because quantifier phrases *are* assigned meaning-relata by (9), they are *meaningful* in isolation. That is, returning to my earlier terminology, they are assigned their meanings "in the primary way". Hence, on this approach, the answer to (1) is: Yes, quantifier phrases are meaningful in isolation. This is in stark contrast with (4). Finally, since this formulation applies with equal naturalness both to quantifier phrases within sentences, *and* to unembedded quantifier phrases, it is *superior* to (4): the advantage of this revised formulation, from the point of view of the use and comprehension of unembedded quantifier phrases, is that it assigns meaning-relata to quantifier phrases whether or not they occur in sentences. To take an example, ⌜Every F⌝ denotes a function from sets to propositions: one which outputs a true proposition when and only when the input set contains the set F.

6. Ellipsis and Other Manoeuvres

I want now to consider two suggestions for defending (4), the "not meaningful in isolation" alternative, while accounting for the use of unembedded quantifier phrases. As will emerge, neither suggestion is satisfactory.

The "No Meaning" Gambit

Here's an initially plausible way of saving the not meaningful in isolation approach; I call it the "no meaning" gambit. Recall the crucial example: saying 'An editor of *Natural Language Semantics*' causes Andy to believe that the indicated seat is reserved for an editor of *NLS*. That this kind of thing can and does occur is indisputable. But, someone might say, this doesn't show that the sounds produced have any kind of semantic content at all, let alone meanings "in isolation". For, it is undeniable that many non-linguistic stimuli can be used to induce beliefs – without having meaning in isolation.[5] To take an obvious case, I may brandish an umbrella, in Anita's direction, thereby inducing in her the belief that it is raining outside. But "umbrella brandishings" have no semantics.

Indeed, to develop the example a bit, even supposing that I, the speaker, induce the belief that a certain chair is set aside for an editor of

[5] My thanks to an anonymous *Linguistics and Philosophy* referee for pressing me on this issue.

NLS by getting Andy to recognize my intention to induce said belief, and so on in familiar Gricean fashion, this *still* doesn't show that 'An editor of *Natural Language Semantics*' has semantic content. For, once again, many kinds of stimuli have *nonnatural meaning*, while nevertheless lacking a semantic value.

If the objection worked, there would indeed be no need for a meaning-in-isolation for quantifier phrases. To see the failings of the "no meaning" gambit, however, it's enough to reflect upon the differences between sentence-based communication on the one hand, and non-linguistic communication on the other; and to ask, of the use of unembedded quantifier phrases, which it resembles.

Sentences can be used to communicate thoughts which are (a) of unlimited complexity and (b) systematically similar/different; what's more, the nature of the thought communicated is (c) comparatively independent of context. Consider just a few examples.

(10) Particle accelerators often cost more than 14 million American dollars each

(11) Particle accelerators often cost less than 12 million American dollars each

(12) Digital clocks seldom cost more than thirty dollars and fifty-five cents a piece, in Canadian currency

The thoughts one would communicate with (10) through (12) are undeniably complex and sophisticated. And more complex thoughts still can be communicated, by adding further qualifier words, or by conjoining these with other sentences, or what-have-you. These examples illustrate that there is no upper bound to the length and complexity of the thoughts one can convey using sentences. Notice also the relatively subtle and systematic similarities and differences between the thought one would typically communicate with (10) as compared to (11). The differences arise from the replacement of 'more' by 'less' and '14' by '12'; the similarities from holding the remainder constant. Finally, (10) can be used in many different circumstances to communicate essentially the same thought: the precise nature of the communicated thought may well change, but there will nevertheless remain something importantly constant across contexts.

In contrast, non-linguistic communication is comparatively simple, course-grained, unsystematic, and context bound. For instance, to communicate something as complex and subtle as the thoughts which (10) or (12) encode, without using language, would be no mean feat. Of course one can, by sticking one's tongue out, express disapproval; one can even be rather more specific than that. Similarly for umbrella brandishing. But,

in general, there's a pretty low threshold to the complexity, precision, and so on of what one can communicate, without using sentence-like symbols.

Now, *why* does sentence-based communication exhibit these three crucial features? Essentially because sentences have semantic content; and that content is determined recursively and compositionally: the meaning of a whole sentence is a function of what its parts mean, and how those parts are put together; and the rules for "putting the parts together" can apply repeatedly, thereby creating an unlimited number of meaningful sentences. In contrast, gestures and other non-linguistic stimuli, *even if* they have conventional content, do not have that content determined by a lexicon and a combinatorix.

Having noted this contrast between communication with and without compositional symbols, consider communication with unembedded quantifier phrases. For example, ask yourself what Andy would have understood, at the linguistics meeting, had I said each of the following – letting all else remain fixed:

(13) A representative of the Uruguayan Linguistics Association
(14) A representative of the Uruguayan Philological Society
(15) The man who coined the phrase 'theta-rôle'
(16) The woman who coined the phrase 'semantic competence'
(17) Someone that you'll really want to meet
(18) A student of Chomsky's that you'll really want to meet

What immediately stands out is: (a) the complexity of the thoughts communicated; (b) the subtle and systematic differences between them; and (c) the limited degree of contextual influence. (I.e. in each case, I tell Andy who will occupy the chair. As before, the chair must be contextually supplied. But practically everything else comes from the phrase uttered.) Indeed, it seems on reflection that the thoughts are, to all intents and purposes, as complex, productive, systematic, subtle, and (roughly) as context-independent as the thoughts communicable via sentences. Let me now suggest why this is: unembedded quantifier phrases, like sentences, have semantic content – which is determined recursively as a function of part meanings, and structure. In which case, it will not do to say that quantifier phrases are *wholly* meaningless.

It's also worth noting, by the way, that unembedded quantifier phrases exhibit such semantic features as ambiguity, anomaly, logical relations, and so on. And that, to understand utterances of such phrases, one must know the language in which they are uttered. Here again, this is nothing at all like umbrella brandishing.

One might reasonably reply that, while unembedded quantifier phrases

must be treated as *linguistically meaningful* stimuli, no less than full sentences, it doesn't follow that they need to be assigned *meanings in isolation*. There is, as I said, an intermediate option between being meaningless, and being meaningful in the primary way. Surely, it might be said, quantifier phrases are meaningful precisely in this intermediate sense: to use slightly different terminology, they have no semantic values – instead, there are syncategorematic rules for how quantifier phrases affect sentence meanings. Thus, this reply might go, in saying an unembedded quantifier phrase ('Three philosophers', for example) the speaker does produce a truly contentful linguistic stimulus. But there's no need to have an account of the *meaning-relatum* of said unembedded quantifier phrase, because it has none.

This line of response raises an obvious question: how can such phrases be *understood* in isolation, if they have no meaning in isolation? Here's a possible answer: the hearer, making use of the context, comes up with some predicate. He then combines this predicate with the heard quantifier phrase, to form a sentence. Only then does he use (4) to interpret the resulting sentence. Outcome? Quantifier phrases are used and understood in isolation, but this in no way impugns the not meaningful in isolation approach, because quantifier phrases, so used, don't need meaning-relata.

My reply. To say that the hearer "comes up with a predicate" merely re-labels the problem of how not-meaningful-in-isolation expressions manage to get interpreted; it does not solve it. Specifically, it remains a mystery how the right predicate is found: the hearer, in his search for the right predicate, cannot rely on the *meaning* of the unembedded quantifier phrase since, by hypothesis, it doesn't come into play until *after* the missing predicate has been found. But, if the quantifier phrase offers *no semantic clue* about where in the context to search, there are going to be far too many "salient predicates" to choose from. Arriving at an interpretation of the speaker would end up being a fabulous stroke of luck. It's just not credible that interpretation works like this. Pretty clearly, the quantifier phrase's content must play a central part in the search for the "right predicate". In which case, the bare phrase cannot be assigned its meaning merely in terms of (4).

Second Defence: Ellipsis

In the foregoing, it was granted that a bare quantifier phrase can be used, but proposed that this didn't entail that the thing used had meaning in isolation. This approach to defending (4) having failed, another suggests itself: deny that what gets used really is a bare quantifier phrase at all.

Here's the idea. When someone says 'An editor of *Natural Language Semantics*' what they produce isn't an indefinite description in isolation; rather, what they produce is an elliptical sentence. This is a natural thought. It's a plausible thought. And it's a thought which might well save (4) – after all, if quantifier phrases *aren't* actually used outside sentences, why make one's theory account for such usage? The thing is, this natural, plausible thought is also (arguably) a *false* thought. I can't go into the evidence in any detail here. But, simply by way of motivating my positive answer to the Question, let me at least sketch a couple of the considerations which mitigate against defending (4), repeated below, by appeal to ellipsis.[6]

(4) *The Not-Meaningful in Isolation Approach:* ⌜**Q** *F*⌝ combines with a predicate ⌜is *G*⌝ to yield a sentence; that sentence is true if and only if . . .

I'll start by laying out a very general account of what ellipsis amounts to. That the following proposal is terrifically general should forestall the hope that my arguments against ellipsis work only for a particular, parochial theory of ellipsis. Crucially, this general account of ellipsis will be such that *if* all apparent uses of unembedded quantifier phrases really are uses of elliptical sentences, *then* something like (4) may still work. Having suggested what ellipsis amounts to, I'll then give several arguments designed to show that this is not what occurs when speakers (appear to) utter unembedded quantifier phrases. I conclude that (4) cannot be defended in this way, and is not the correct account of the semantics of quantifier phrases; and that, indeed, quantifier phrases have meaning in isolation.

Before continuing, however, I want to warn against assuming an ordinary, common-sense notion of ellipsis – according to which, whenever someone leaves something unsaid, they are speaking elliptically. In *this* sense of ellipsis, the use of bare phrases obviously *is* elliptical: the speaker communicates more than what her words mean. But appeal to ellipsis in this weak sense won't rescue (4), since an ellipsis defence of (4) requires that *the expression uttered* be sentential. And this is not established by pointing out that *the speaker* left certain information unspoken.

In particular, so far as I can see, a speaker may convey a proposition without uttering a sentence. (Compare one of Grice's (1975: 156) cases: A professor writes, in a letter of reference, 'Mr. X's command of English is excellent, and his attendance at tutorials has been regular'. Here, the

[6] See Barton 1990, 1991; Brame 1979; Dalrymple 1991; Morgan 1989; Napoli 1982; and Yanofsky 1979 for additional arguments.

proposition communicated is something like *Mr. X is a crappy philosophy student*; but this in no way shows that *these words* were produced by the speaker.) So the mere fact that a proposition is communicated does not establish that the speaker produced any kind of sentence.

In sum, there's a sense of 'elliptical' in which a use of a bare phrase would be elliptical; but it's not a sense which is especially relevant to the issue at hand. For the issue at hand is: What *expression* did the speaker produce, an unembedded quantifier phrase, or a sentence? If the answer is, "A quantifier phrase", then (4) needs patching.

That being said, what is ellipsis – in the sense in which it would preserve (4)? Preliminary remark: linguistic items, both tokens and types, fall into different classes; in particular, some are syntactically sentential while others are syntactically lexical or phrasal. Roughly speaking, a linguistic representation is syntactically sentential if and only if it is headed by an inflectional element at the appropriate level of representation – where inflectional elements include modals, tense and verb-subject agreement. (For further discussion, see Chomsky (1981, 1982, 1986a, 1986b), Haegeman (1991) and references cited there.) An expression is lexical/phrasal otherwise. Now, everyone agrees that speakers *appear to* utter unembedded quantifier phrases. For example, everyone recognizes that a man may approach an apple cart and say 'Three big red apples', or knock on a co-worker's door and say 'The Leibnitz reading group'.[7] In both cases, the speaker performs a speech act: in the first case, the man requests three big red apples; in the second case, he informs his colleague of an upcoming engagement. And, in both cases, the speakers appear to use an unembedded quantifier phrase, in the sense that the expressions they produce sound like phrases. To defend (4), then, its proponent might maintain that, despite the fact that certain utterances do not sound like ordinary sentences, they nevertheless are sentential in the syntactic sense.

How to do this? Here's an "ur-proposal". I will say that an expression r is **shortened** if and only if there exists another expression r' such that r' has a longer phonetic form than r, but r' has the same syntactic structure as r.[8] In a word, a linguistic expression becomes an *ordered pair* of a syntactic structure and a phonetic form. It is the syntactic structure of an expression that determines whether it is sentential or phrasal; and it is the phonetic form that determines how the expression sounds. (In which case,

[7] Here and elsewhere I assume a Russellian theory for definite descriptions, according to which they are quantificational rather than referential.

[8] I leave open the question of how, precisely, the notion of length should be explicated – relying in what follows on an intuitive understanding of this notion.

putting it crudely, ellipsis comes to this: a syntactic structure which is headed by INFL gets paired with an "abbreviated" phonetic form.) Given the notions of syntactically sentential expressions and shortened expressions, I can now introduce the ellipsis hypothesis:

(19) *The Ellipsis Hypothesis:* Whenever a speaker performs a speech act by uttering an (apparently) unembedded quantifier phrase, what that speaker really utters is an elliptical sentence in the sense that her utterance is syntactically sentential, but it is shortened.

It is the "shortening", of course, that explains why the result does not "sound like" an ordinary sentence – even though the utterance is syntactically sentential. (Notice that the ellipsis hypothesis is stated in such a way as to be indifferent to how the abbreviation occurs – syntactic deletion, null elements, phonological deletion, or what-have-you.)

To take one example: a speaker who uses the sound *the head of philosophy* to communicate that the head of philosophy has arrived should be described not as uttering (20), but rather as uttering (21).[9]

(20) ⟨[$_{QP}$ The head of philosophy], *the head of philosophy*⟩

(21) ⟨[$_{IP}$ The head of philosophy has arrived], *the head of philosophy*⟩

Rule (4) would then straightforwardly apply – because the sentence ⟨[$_{IP}$ The head of philosophy has arrived], *the head of philosophy*⟩ is, from the point of view of syntax/semantics, equivalent to the sentence 'The head of philosophy has arrived'; and this latter sentence is easily captured by an instance of (4).

I hope it's now clear what it would amount to, to defend (4) by insisting that when a speaker appears to utter an unembedded quantifier phrase, what she really utters is an elliptical sentence. I now turn to the question of whether it's true. First problem for the ellipsis hypothesis. Witness the fact that, in the situation described above, 'The head of philosophy' are the first words uttered. This strongly suggests that the expression produced

[9] A word about notational conventions. In what follows, I will represent the syntactic structure of utterances by a labelled bracketing. For example, suppose John utters the sentence 'Snow is white'. I use the following notation to give the syntactic structure of John's utterance: [$_{IP}$ snow is white]. I use English orthography in italics to give the phonetic form of utterances. The phonetic form of John's utterance, for example, would be given by *snow is white*. To give the full linguistic representation of an utterance, I use an ordered pair of a syntactic structure and a phonetic form – in that order. Where no confusion will arise, I also use single quotes to talk about expressions – understood as complexes of syntactic structure, phonetic form, etc.

is not elliptical in the desired sense – since, in so far as linguists know anything about ellipsis, they know that a shortened phonetic form demands a *linguistic* context. With certain notable exceptions – e.g., asking permission by saying 'May I?' – you cannot, without awkwardness, begin a conversation with an elliptical sentence: 'He doesn't' or 'I wonder when' sound odd as discourse on-sets, even when they would be understood. (See Hankamer and Sag 1976 for the details.) But – in contrast to cases of authentic ellipsis – you *can* begin a conversation with a quantifier phrase, without (or with much less) awkwardness: recall the man whose first words to the apple vendor are 'Three big red apples'. Nor, incidentally, does the use of unembedded quantifier phrases require a "pragmatic controller", in the sense of Hankamer and Sag (1976): following Yanofsky (1978), I note that 'A tie' can be used to remind someone to wear a tie – this being necessary precisely when no tie is antecedently salient! Here again, there is an important contrast with true ellipsis, which does require such a pragmatic controller when no linguistic antecedent is available.

The above argument has the form: it doesn't quack like ellipsis, and it doesn't walk like ellipsis; so, plausibly, it isn't ellipsis. Arguments like this are not, of course, conclusive: one could, pending a finding of deep similarities, conclude that a non-waddling, non-quacking creature is a previously unfamiliar kind of duck; and one could, after sufficient investigation, conclude that the items produced in discourse initial position are a previously unfamiliar kind of syntactic ellipsis. But, that these fragments lack one of the fundamental properties of elliptical expressions is solid *prima facie* evidence that they are not elliptical.

And too, it would be surprising if elliptical sentences *could* be used in discourse initial position. In discourse initial position there are, of course, no *linguistic* cues in the context for going from the encountered phonetic form to the corresponding syntactic structure. This poses a problem because, as should be obvious, if the ellipsis hypothesis is correct then a given phonetic form does not determine anything like a unique syntactic structure.

Consider an example. It would seem, given the ellipsis hypothesis, that all of (22) through (24) are expressions of English:

(22) ⟨[$_{IP}$ The famous lawyer loves Fred], *the famous lawyer*⟩
(23) ⟨[$_{IP}$ Fred loves the famous lawyer], *the famous lawyer*⟩
(24) ⟨[$_{IP}$ The famous lawyer detests Steve], *the famous lawyer*⟩

These expressions would exist, on the ellipsis hypothesis, because the phonetic form *the famous lawyer* can, given the right context, communicate any of the following propositions: THE FAMOUS LAWYER

LOVES FRED; FRED LOVES THE FAMOUS LAWYER; and THE FAMOUS LAWYER DETESTS STEVE. Given that utterances with the phonetic form *the famous lawyer* can convey any of these, there must be at least three linguistic representations that share this phonetic form – again, if the ellipsis hypothesis is correct.

Next step. For any given phonetic form, there are an *unlimited* number of propositions which it could encode. Therefore, by parity of reasoning, any phonetic form whatever corresponds to an unlimited number of syntactic structures.[10] Evidently, this raises a question about parsing; namely, how does the parser find the single intended syntactic structure? In the standard (i.e. non-elliptical) case, one must assume that it outputs the shortest syntactic structure consistent with the phonetic form. (Otherwise, to take an example, *Mary sleeps*, said in isolation, could get paired with, among other things: [$_{IP}$ Alex mistakenly supposes that on Friday nights Mary sleeps at her friend's house].) In the elliptical case, in contrast, linguistic clues from the context must provide the necessary evidence for getting the right, non-minimal, syntactic structure.

And this shows why *discourse initial* ellipsis would be peculiar: without linguistic clues from the context, it will typically be impossible to find the correct non-minimal syntactic structure of an elliptical utterance, but discourse initial position is precisely a case in which linguistic clues are unavailable. So, it's highly likely that elliptical sentences cannot be used in discourse initial position. The expressions in question, on the other hand, *can* be used in discourse initial position. Hence, it's plausible to suppose, said expressions are not elliptical sentences – at least not in the sense required by an ellipsis defence of (4).

Another argument against the ellipsis hypothesis. Certain constructions (e.g. VP deletion and sluicing) cannot acceptably occur if there is no prior syntactically sentential item in the discourse. This provides a sort of *prima facie* test for syntactically sentential linguistic items in prior discourse. And this test suggests that (apparent) unembedded quantifier phrases *aren't* underlyingly sentential: discourses containing VP deletion constructions and sluicing constructions become significantly less acceptable when (apparent) quantifier phrases are substituted for sentences. Here are examples.

[10] And, obviously, every phonetic form will correspond to infinitely many meanings.

VP Deletion

(25) Jason: The man from Paris is at the door
Mark: And Betty is too

(26) Jason: The man from Paris
Mark: ??And Betty is too

Sluicing

(27) Jason: The man from Paris is at the door
Mark: I wonder why

(28) Jason: The man from Paris
Mark: ??I wonder why

The sentence 'The man from Paris is at the door' differs minimally from the phrase 'The man from Paris', in the sense that both expressions can be used to communicate the proposition that the man from Paris is at the door. Yet, if one substitutes the phrase 'The man from Paris' for the full sentence, in (25) and (27), the result (given in (26) and (28) respectively) is less acceptable – even when the thought that the man from Paris is at the door is successfully communicated.

You might think: this doesn't show that 'The man from Paris' is non-sentential, because 'And Betty is too' requires a *non-elliptical* sentence as its linguistic antecedent. But this doesn't seem right. Notice, for example, that where 'The man from Paris' answers a wh-interrogative (which would plausibly make it an elliptical sentence), 'And Betty is too' becomes quite okay. Thus:

(29) Lenny: Who's at the door?
Jason: The man from Paris
Mark: And Betty is too

(Another point. Though it's unclear why, 'And _ is too' is sensitive to sentential antecedents in a way that other related expressions are not. That's why I use 'And Betty is too', rather than anything else.)

Though far from conclusive, this data suggests that 'The man from Paris' is not syntactically sentential: if it were, the discourses (26) and (28) as a whole should be perfectly acceptable – which they're not. A fortiori, 'The man from Paris' is not syntactically sentential *and shortened*. That is, 'The man from Paris' is not an elliptical sentence.

I repeat: the foregoing considerations are not definitive and decisive. They are not intended to be. But there is much more evidence of this kind, the total weight of which is quite convincing. In any case, my purpose

in this section was merely to make it plausible that quantifier phrases – not elliptical sentences, mind you, but ordinary quantifier phrases – can be used and understood on their own, outside any sentence. Which should be enough to motivate, at least provisionally, hypothesis (9), and the consequent in-isolation meaningfulness of quantifier phrases:

(9) *The Meaningful in Isolation Approach:* ⌜**Q** *F*⌝ denotes the function *f* from sets *G* to propositions such that *f*(*G*) is a true proposition if and only if . . .

7. The Pragmatics of Unembedded Quantifier Phrases

A hypothesis about the semantics of an expression must, at the very least, be consistent with facts about how utterances of that expression are typically understood. I'll suppose it uncontroversial that unembedded quantifier phrases, when used in context, are often understood as communicating quantified *propositions*. But, if treating quantifier phrases as corresponding to generalized quantifiers is on the right track, this is not what unembedded quantifier phrases mean: semantically speaking, an unembedded quantifier phrase corresponds, by hypothesis, to a function from sets to propositions. The failure of fit between the semantics which I have proposed for unembedded quantifier phrases, and how these are understood, might be thought a problem. But, happily, pragmatics can bridge the gap. That, anyway, is what I'll argue.

This isn't the place for laying out the pragmatics of unembedded quantifier phrases in painstaking detail. (Interested readers might look at Stainton (1994), where the interpretation of *other* non-sentences is discussed at length.) So, I will simply sketch a story about how unembedded quantifier phrases can be used in communication. In order to do so, however, I will need to employ numerous ideas from Relevance Theory.

Some definitions. Let **logical forms** be expressions of mentalese. Let **assumptions** be propositional logical forms, and **assumption schemas** be non-propositional logical forms. (In effect, assumptions are sentences of mentalese, while assumption schemas are mentalese predicates, names, quantifier phrases etc.) A logical form is **manifest** to an individual at a time *t* only if she is capable of representing it mentally at *t*. But this is not sufficient for manifestness. Roughly, an *assumption* is manifest at *t* only if the person whose representation it is, is capable of accepting it as true, or probably true, at *t*; an *assumption schema* is manifest to a person at *t* only if what it represents is perceptible at *t*. Let an individual's **assumption-set** be the collection of assumptions currently manifest to her.

(Assumption-sets cannot, by definition, contain assumption schemas.) Finally, call an assumption A **relevant** to an individual at a time to the extent that A positively affects the individual's **assumption-set**, where an assumption-set is positively affected by having assumptions added to it, deleted from it, etc. The more positive effects, the more relevant; but also, the less processing cost (e.g. inferential labour, perceptual effort, memory strain) the more relevant.[11] (Note: I have included an Appendix, at the end of the paper, which lays out some the basic tenets of Relevance Theory; if unfamiliar with Sperber and Wilson's framework, you may want to read the Appendix before going on. Specifically, technical terms such as 'manifest', 'affecting', etc. are discussed there.)

Time to put the definitions to work. To fix ideas, I'll focus on one example: Andy and I hear some loud noises. I peer through the window, and utter the phrase 'Three dogs'. Assuming this phrase corresponds to a function from sets to propositions, as (9) demands, how can it be used communicatively – e.g. to report that there are three dogs outside? Here is the general idea. The utterance of the quantifier phrase makes manifest a number of logical forms. In particular, the utterance makes manifest the following two items:

(30) *The presumption of optimal relevance:* that the utterance is relevant enough for it to be worth the addressee's while to process it; and that the utterance is the most relevant one available.

(31) The logical form corresponding to the quantifier phrase uttered (in the case at hand, the logical form corresponding to the English expression 'Three dogs'), where the content of this logical form is a generalized quantifier.

Here is the key claim: having this much manifest is enough for successful communication, because (30) and (31) can jointly serve as a basis for finding the assumption-set consistent with the presumption of optimal relevance – an assumption-set which contains the assumption that there are three dogs outside. And this, as you'll see, is sufficient for communicating the latter assumption.

The assumption-set, at the outset, does not contain (31) because, though it is a logical form, it is not an assumption. (Remember: assumptions are,

[11] Does the definition of manifestness mean that surprising facts can never be manifest? Well, it does mean that they can't be very manifest, at least until attention is drawn to them. However, this is precisely the reason why surprising facts can be very relevant. And what is communicated is what is most relevant, not what is most manifest. So there is certainly no reason why surprises can't be communicated.

by definition, *propositional* logical forms. (31) is nonpropositional: its meaning is a generalized quantifier, not a proposition.) However, the assumption-set can, with very little effort, be made to contain an assumption got from (31). This is easy enough, assuming some sets (or properties?) are salient. If a set s is salient in the environment, then, by the definition of manifestness, the logical form corresponding to s will be manifest; and a logical form which represents a set, when combined with (31), yields an assumption. (Compare: a quantifier phrase, when combined with a predicate, yields a sentence.) For instance, combining the logical form **be-outside** with the logical form of 'Three dogs' gives (32), which in turn λ-converts to (33):[12]

(32) $\lambda g \in \langle e, t\rangle.\lambda w.[\exists_3 x(dog(x) \;\&\; g(x))]^W$(**be-outside**)

(33) $\lambda w.[\exists_3 x(dog(x) \;\&\; \textbf{be-outside}(x))]^W$

Assumption (33), whose meaning is the proposition that there are three dogs outside, will be *relevant enough* in certain circumstances: for instance, it will be relevant enough in a situation where the hearer wants to know the source of the observed noise. Hence an assumption-set containing (33) satisfies the first part of the presumption of optimal relevance. But – and this is crucial – **be-outside** will be manifest in a subset of these circumstances.[13] Indeed, in some cases where (33) is relevant, **be-outside** will be the *most* manifest logical form, beyond (30) itself. In such circumstances, the assumption-set containing (33) will be the most accessible assumption-set: the first one the hearer considers in her interpretive task. And, as Sperber and Wilson argue, the *most* accessible assumption-set which is relevant is the *only* assumption-set consistent with the presumption of optimal relevance; in turn, the only assumption-set consistent with the presumption of optimal relevance is the assumption-set communicated by the speaker. This set includes (33). And thus you see how an utterance of 'Three dogs', taken as encoding a generalized quantifier, succeeds in communicating a proposition. Presto.

Except for two worries. First worry. It might seem that an utterance of 'Three dogs' cannot really succeed in communicating a proposition because it will never be the most relevant utterance available. Surely, one

[12] Not knowing how to write in mentalese, I use the notation of intensional logic.

[13] To say that **be-outside** is manifest seems to suggest that the set of things $\{x\colon x$ is outside$\}$ is perceptually salient. This may sound odd to some ears. Frankly, it sounds odd to mine. However, I believe the oddness derives not from the supposition that **be-outside** can be manifest, but rather from an overly simple notion of what this assumption schema might represent. Since what is crucial for my purposes is that **be-outside** should be manifest, I will simply abstract away from this issue.

might say, a fully sentential utterance will always be more relevant, since spelling everything out for the hearer would require less inferential work (hence less processing effort) on her part. The point is well taken. But, as I've argued in Stainton (1994), it is far from clear that supplying more linguistic material *must* make an utterance easier to process. Quite the contrary, where an appropriate assumption schema is already very manifest (as **be-outside** is assumed to be in the example above), forcing the hearer to decode a predicate (e.g. [$_{1'}$ are outside]), and then develop the predicate – only to arrive at an (already manifest) logical form – would involve more effort than simply decoding the non-sentence, and conjoining it with the manifest assumption schema. Put more flat-footedly: sometimes it's more efficient to leave the already-obvious unspoken.

What's more, a point also elaborated upon in Stainton (1994), the use of a bare quantifier phrase might have stylistic effects which make the recovered assumption-set richer than it would otherwise be. Such a use might indicate urgency, or informality, etc. Given this, though it remains an open empirical question, I think it's fair to say that an utterance of an unembedded quantifier phrase could well be the most relevant utterance available – given the right circumstances.

Second worry about this Relevance Theoretic story. I insisted that the use of unembedded quantifier phrases was *not* a matter of ellipsis. But – you might wonder – haven't I just given an ellipsis-based account? I don't think so. Let me tell you why.

First of all, the interpretive process, as I see it, goes like this. The hearer of an unembedded quantifier phrase decodes the bare quantifier phrase, and recovers its logical form; since logical forms, being formulae in mentalese, determine meanings, this entails that the hearer recovers the meaning – *before* she recovers the assumption communicated. On the ellipsis story, on the other hand, the hearer cannot assign a meaning to the expression uttered until a whole sentence has been recovered. (Remember, the whole point of introducing ellipsis would be to preserve the idea that only *within sentences* are quantifier phrases meaningful.) So, to speak loosely, I'm in disagreement with the ellipsis proponent about *the stage* at which the quantifier phrase is decoded and understood.

A closely related point. As I see it, in forming interpretive hypotheses the hearer employs the assumption schema got from decoding – i.e. the logical form of the unembedded quantifier phrase – as an important clue in finding the apppopriate assumption schema to combine with it. (In terms of the example, having the logical form of [$_{QP}$ Three dogs] already manifest will surely make it easier to determine that **be-outside**, among all the manifest assumption schemas, is the desired one.) Taking the

ellipsis story seriously, however, the quantifier phrase itself provides no semantic clue to the interpreter – precisely because it is assumed to have no meaning in isolation.

Which takes me to the final important difference between my account and an ellipsis story. (I get a bit fuzzy-headed when I think about this last contrast; but I think the point is worth making nonetheless.) In an ellipsis story, the hearer has to recover a *linguistic* unit – specifically, a natural language predicate – in order to continue the interpretive process. That is, she must recover whatever linguistic material was elided from the fully sentential source. This looks hard. The difficulty arises for two reasons. First, the quantifier phrase itself can provide no semantic help – I'll say it again: the ellipsis story is supposed to preserve the idea that quantifier phrases have no meaning in isolation. Second, generally speaking the context won't allow the hearer to single out the sought-for linguistic item – even if she *did* receive a semantic clue from the quantifier phrase. I don't know how to elaborate this point, except metaphorically. Here it goes. To re-work a passage from Davidson (1978: 263), a non-linguistic context is not worth a thousand words, or any other number. Natural language words are the wrong currency to exchange for a non-linguistic context. Put otherwise, there is no single, correct mapping from situations to public language descriptions of them; so the non-linguistic context cannot determine a linguistic item. In which case, there will generally be no (unique) "salient predicate" for the hearer to concatenate with the quantifier phrase. In contrast, logical forms, unlike words, are *precisely* the right "currency" to exchange for a non-linguistic context – since what is manifest is determined, in very large part, by the actual environment. Remember, anything *perceptually salient* is manifest; and what is perceptible in an environment obviously depends closely on the contents of the environment. So, whereas a non-linguistic something cannot generally be mapped onto a public language structure (elliptical or otherwise), a perceptible (but not necessarily perceived) non linguistic something *will* always make a logical form manifest. This provides a final contrast between (a) selecting a natural language unit – got from the quantifier phrase and the context – and interpreting it (i.e. ellipsis); and (b) combining the logical form of a quantifier phrase – a mentalese expression – with another manifest logical form (anti-ellipsis). Hence, my account is not a mere variation on ellipsis.

To further bring out this contrast, it may help to introduce a view which is not my own. The ellipsis story has the uttered quantifier phrase combining with a contextually salient predicate, to form a sentence; this sentence is then interpreted using (4). An alternative anti-ellipsis story

(again: not my own) has the *denotation* of the quantifier phrase – as given by (9) – combining with some contextually salient *non-linguistic entity*, to yield a proposition. This proposition is communicated by the speaker, but it is not the content of the words which the speaker uttered; for, what she uttered denotes not a proposition, but a propositional function.

Thinking about it this way, there is a clear difference between the two accounts. Crudely, one describes the process in the formal mode (the ellipsis account), while the other uses the material mode (the anti-ellipsis account). Furthermore, this picture highlights the sense in which a context can determine salient *entities* in a way that it need not determine salient *labels* for these entities.

Returning to my own view, this contrast is captured via explicit reference to mentalese versus natural language. That is, I cash "a non-linguistic entity is salient" as: "the mentalese predicate corresponding to said entity is manifest"; and I construe "the denotation of the quantifier phrase combines with this entity" as: "the mentalese assumption-schema which translates the public language quantifier phrase combines with the contextually manifest mentalese predicate".

Putting it my way, one can no longer draw the stark contrast between, roughly speaking, the formal mode approach (i.e. ellipsis) and the material mode approach (i.e. anti-ellipsis). For, seen Relevance Theoretically, *both* approaches are committed to the recovery of some kind of predicate, which is then combined with some kind of quantificational *expression*. But the contrast between a natural language predicate being salient, and being combined with a natural language quantifier phrase; and a mentalese predicate being manifest, and being concatenated with a mentalese correlate of a quantifier phrase, though subtle, is real enough. And it distinguishes my view, according to which quantifier phrases really are used in isolation, from the ellipsis view, according to which they are not so used.

8. Summary and Conclusion

Time to sum up. I have suggested a positive answer to (1), repeated below:

(1) *The Question:* Do quantifier phrases have meaning in isolation?

By asserting that quantifier phrases are meaningful in isolation, I mean that they are assigned meaning "in the primary way" (i.e. by receiving a denotation), rather than being allotted their meaning solely in terms of how they affect the meaning of sentences in which they occur. The grounds

for this conclusion are, admittedly, a bit odd: quantifier phrases can be used and understood even when no *natural language predicate* is available, to combine with the quantifier phrase. Odd or not. I believe this fact about actual usage favours (9) over (4).

By way of defending my positive answer to (1), I considered and rejected two alternative views about how unembedded quantifier phrases apparently manage to be used and understood, consistent with (4). According to one view, they are *not* so used – appearances to the contrary are engendered by misconstruing elliptical sentences as truly unembedded phrases. This alternative was discarded on empirical (especially syntactic) grounds. According to a second view, quantifier phrases really are used and understood in isolation; but, so used, they exhibit no semantic content whatever. This alternative was rejected on the grounds that it failed to account for the complexity, subtlety and relative context-independence of phrasal speech acts. Finally, I attempted to render pragmatically plausible the idea that quantifier phrases are meaningful in isolation – and that what they "mean in isolation" are generalized quantifiers. Here, I used Relevance Theory to show that assigning this kind of meaning to quantifier phrases is at least consistent with the fact that quantifier phrases are used and understood in isolation.

9. Appendix: A Tutorial on Relevance Theory

Interpretation, as Sperber and Wilson (1987, 1995) see it, consists of two steps. On the one hand, the hearer must *decode* the linguistic signal; on the other hand, she must infer utterance meaning – on the basis of what is decoded, plus any other available evidence. Decoding furnishes only the linguistic representation of the utterance, including both its syntactic structure and its *logical form* – where the latter is the mentalese symbol that gives the meaning of the expression. (I'll abstract away from how the decoder achieves this task.)

Inference **develops** this logical form, to arrive at the proposition(s) expressed. Development is almost always required because the logical form output by the decoding process will not, in general, be fully propositional – simply because, put in more familiar terms, expression meaning often falls short of a complete proposition. (Think of 'He bought that'. This expression, the type that is, does not correspond to a proposition; to get a proposition, the meaning of the expression must be supplemented by a context.) Put in Relevance Theoretic terms, the logical form of the expression uttered is ordinarily not an **assumption** – where, by definition, an assumption is a logical form that *does* express a proposition. So, to

understand the proposition communicated, the hearer must "complete" the logical form, until he arrives at a logical form that *does* express a proposition. The question is, how does the hearer do this? To tell that tale I need to introduce still more background.

Sperber and Wilson (1995: 39) say that an assumption is **manifest** to an individual at a given time if and only if she is capable of representing that assumption mentally and accepting that representation as true or probably true at that time. There are several ways that an assumption may be manifest to an individual. It may be perceptible in the physical environment; it may be inferable from assumptions which are already manifest; or it may be retrievable from memory. It is important to stress the modality at work in this definition. To be manifest, an assumption need not have been already perceived, remembered or inferred; rather, what is required for manifestness is the mere possibility that the assumption be perceived, inferred or remembered. Notice too, it is not propositions or states of affairs which are manifest. Rather, assumptions – formulae of mentalese are manifest.

Manifestness, according to Sperber and Wilson, admits of degrees. Assumptions which are more likely to be held true are more manifest. Consider an example. It may be manifest to Watson that Holmes is holding a pipe, but more manifest to him that Holmes is speaking – because Watson is more likely to hold this latter assumption true. In all likelihood, it will be less manifest to Dr. Watson that Holmes has never been to the moon. Not because Watson harbours any doubts; only because he is unlikely to even entertain the possibility that Holmes has been to the moon. Nevertheless, Watson is capable of considering this assumption, however odd. Hence it is manifest to him, albeit very slightly.

Final bit of terminology. (Here I'm simplifying like mad.) Call an assumption A relevant to an individual at a time to the extent that A positively affects the individual's **assumption-set**: the individual's *collection* of currently manifest assumptions. (E.g. A may be relevant by *adding* new and useful assumptions to the individual's assumption-set; or by making *more manifest* assumptions which are already a little manifest; or by *removing* false assumptions from the assumption-set.) Next step. The very assumption A can be more or less relevant, to an individual at a time, depending on how much effort is required to process A. A has greater relevance to the extent that the processing effort it requires is small; and it has less relevance to the extent that the processing effort it demands is large. Given the notions of manifestness and relevance, I can now employ Sperber and Wilson's (1995: 270) **principle of relevance**.

(34) *Principle of Relevance:* Every communicative act communicates the presumption of its own optimal relevance.

(35) *Presumption of Optimal Relevance:*
a. The utterance is relevant enough for it to be worth the addressee's while to process it, and
b. The utterance is the most relevant one available.

If communication is to succeed, the speaker must communicate the presumption of optimal relevance. He must communicate (a) that he is communicating assumptions which are relevant enough, and (b) that he has chosen the most efficient means available for communicating these assumptions. Why is this so, according to Sperber and Wilson? Here's the idea. If the speaker is to succeed in communicating, he must persuade his audience to interpret his utterance. This requires convincing that audience to expend the necessary interpretive effort. Sperber and Wilson claim that speakers convince the audience by making it manifest that the speaker intends to communicate assumptions which are relevant to the audience. Sperber and Wilson are aware, of course, that speakers do not always communicate in good faith. A speaker may claim his audience's attention without having anything truly relevant to communicate. But, they maintain, unless a speaker at least pretends to be aiming for relevance, he will fail to communicate anything. It is in this sense that speakers inevitably communicate, about their very own utterances, that the latter are relevant enough.

So much for part (a) of the presumption of optimal relevance. What about part (b)? According to Sperber and Wilson, if communication is to succeed, a speaker must communicate that his utterance is the *most* relevant utterance available for communicating the set of assumptions in question – call it {I}. Why is this? Well, as they point out, the most effective signal for communicating some set of assumptions {I} is the one which makes it as easy as possible for the addressee to understand {I} (Sperber and Wilson 1995: 157). And the signal which makes it easy as possible for the addressee to understand {I} is precisely the one which requires the least processing effort. Finally, the signal which requires the least processing effort to recover {I} is the most relevant signal capable of making {I} manifest – because the most relevant signal is the one which yields the most positive changes in the interpreter's assumption-set, at the least cognitive cost. In a word, the following identities hold:

(36) The most effective signal for communicating {I} = the signal

which makes it as easy as possible for the addressee to understand {I}

(37) The signal which makes it as easy as possible for the addressee to understand {I} = the signal which requires the least processing effort to recover {I}

(38) The signal which requires the least processing effort to recover {I} = the most relevant signal for making {I} manifest

By transitivity of identity,

(39) The most effective signal for communicating {I} = the most relevant signal for making {I} manifest

Now, if the speaker wishes to successfully communicate a set of assumptions {I}, she will undoubtedly select the most effective signal available for communicating {I}. And, as just explained, the most effective signal for communicating {I} is the most relevant signal for making {I} manifest. So, if the speaker wishes to successfully communicate, she will select the most relevant signal for making {I} manifest.[14]

In speaking utterers communicate both that their utterance is relevant enough, and that it is the most relevant one available. By assuming that this promise of optimal relevance was made in good faith, the hearer can eliminate very many hypotheses about what a speaker might have meant: she can reject any hypothesis which would have the speaker violating the presumption of optimal relevance. But, it might be thought, this criterion leaves a multitude of possible interpretations, all of which are consistent with the presumption of optimal relevance. If many interpretations satisfy this demand, how does the hearer select a single interpretation? In re-

[14] This establishes that a speaker *must choose* the most relevant stimulus available – if he wishes to communicate successfully. But Sperber and Wilson make a stronger claim. They maintain that speakers inevitably *communicate* that they are using the most relevant stimulus available. Why this extra step? Sperber and Wilson (1995: 157) answer as follows:

> An addressee who doubts that the communicator has chosen the most relevant stimulus [available] – a hearer, say, who believes that he is being addressed with deliberate and unnecessary obscurity – might doubt that genuine communication was intended, and might justifiably refuse to make the processing effort required. All of this is mutually manifest; it is therefore mutually manifest that the communicator intends it to be manifest to the addressee that she has chosen the most relevant stimulus capable of fulfilling her intentions.

That is, by communicating that she has chosen the most relevant stimulus, the speaker helps to insure that the hearer will interpret her. For, if she fails to communicate this – if, for example, the hearer takes her to be using a less than optimally relevant stimulus – the hearer may not make the necessary interpretive effort. So, speakers not only inevitably select the most relevant stimulus available; they inevitably communicate that they have selected the most relevant stimulus available.

sponse to this question, Sperber and Wilson (1995: 167) argue that there is only ever one set of assumptions which is truly consistent with the presumption of optimal relevance: the only set {I} consistent with the presumption of optimal relevance is the *first* set of assumptions {I} which the hearer considers, and which is relevant enough.

This talk of "the first" presupposes some ordering of sets of assumptions. The ordering is in terms of **accessibility**. Sperber and Wilson (1995: 77) write that, "A more accessible assumption is one that is easier to recall". They add, ". . . the more a representation is processed, the more accessible it becomes". It is not wholly clear what accessibility comes to, but the intuitive idea can be brought out as follows. Some assumptions are more easily brought to consciousness than others; furthermore, some assumptions can be retrieved from long term memory with ease, while others require significant effort. Similarly, some assumptions can easily be introduced into an individual's assumption-set; other assumptions could become part of the individual's assumption-set only with a good deal of effort. Those assumptions which require less effort to become part of an individual's assumption-set at a given time are more accessible for that individual at that time.

Sperber and Wilson maintain that, in assessing interpretive hypotheses, hearers begin by testing the most accessible set of assumptions – in this sense of "accessible". If this set of assumptions is not consistent with the presumption of optimal relevance, the hearer goes to the next most accessible set, and tests it. This continues until a set of assumptions is found which is consistent with the presumption of optimal relevance. The most accessible set of assumptions which is consistent with the presumption of optimal relevance is the only one consistent with the presumption of optimal relevance. And it is the set of assumptions being communicated.

Why, according to Sperber and Wilson, is the most accessible set of assumptions the only set of assumptions consistent with the presumption of optimal relevance? They write:

> An addressee . . . who wants to maximize cognitive efficiency, will test hypotheses in order of accessibility. Suppose he arrives at a hypothesis which is consistent with the principle of relevance. Should he stop there, or go on and test the next hypothesis on the grounds that it might be consistent with the principle of relevance too? It is easy to show that he should stop there. Suppose he does go on, and finds another hypothesis which verifies the first part of the presumption of relevance: the putative set {I} is relevant enough. In these circumstances, the second part of the presumption of relevance is almost invariably falsified. If it was at all possible, the communicator should have used a stimulus which would have saved the addressee the effort of first accessing two hypotheses consistent with the principle of relevance, and then having to choose between them (Sperber and Wilson 1995: 167–168).

This argument goes by rather fast, and it establishes a rather important

conclusion. So let me unpack it. Sperber and Wilson want to establish the conclusion below:

Conclusion. The first interpretation of an utterance u which is consistent with the presumption of optimal relevance is the only interpretation consistent with the presumption of optimal relevance. Therefore, for any u, there is only one interpretation of u consistent with the presumption of optimal relevance.

To establish this conclusion, assume that there is some utterance u which has two interpretations consistent with the presumption of optimal relevance. From this assumption, a contradiction will be derived.

Premise 1. There is at least one utterance u such that u has two interpretations consistent with the presumption of optimal relevance: {I1} and {I2}.

Sperber and Wilson then observe that, "almost inevitably":[15]

Premise 2. There exists some other utterance u' such that {I2} is the first interpretation of u' consistent with the presumption of optimal relevance.

They then point out that finding the first interpretation of u consistent with the presumption of optimal relevance (i.e. finding {I1}), rejecting it, and finally finding {I2} involves more processing effort than finding the first interpretation of u' consistent with the presumption of optimal relevance (i.e. finding {I2}). In a word:

Premise 3. Interpreting u as communicating {I2} requires more processing effort than interpreting u' as communicating {I2}.

Now, recall the second extent condition on relevance. An assumption is relevant to the extent that the effort required to process it is small (Sperber and Wilson 1995: 125). So, by premise 3, u' is more relevant than u – when both are taken as communicating {I2}. But then it is not true that {I2} is an interpretation of u which meets the presumption of optimal relevance, for there exists a more relevant means of communicating {I2}, viz. u'. This contradicts Premise 1.

This argument establishes that, for any utterance u, there cannot be two interpretations of u consistent with the presumption of optimal relevance. There can be only one. That one is the first interpretation that passes the

[15] Sperber and Wilson include the hedge because of situations in which the communicator has at his disposal a very limited range of stimuli with which to communicate. When this happens, there may be no stimulus that has {I2} as its most accessible interpretation. They maintain, however, that natural languages are not limited in this way.

test; and that single interpretation is the content of the communicative act.[16]

REFERENCES

Barton, Ellen: 1991, 'Nonsentential Constituents and Theories of Phrase Structure', in K. Leffel and D. Bouchard (eds.), *Views on Phrase Structure*, Dordrecht, Kluwer Academic Publishers.

Barton, Ellen: 1990, *Nonsentential Constituents*. Philadelphia, John Benjamins Publishing Company.

Barwise, Jon and Robin Cooper: 1981, 'Generalized Quantifiers and Natural Language', *Linguistics and Philosophy* **4**, 159–219.

Brame, Michael K.: 1979, 'A Note on COMP S Grammar vs. Sentence Grammar', *Linguistic Analysis* **5**, 383–386.

Chomsky, Noam: 1986a, *Barriers*, Cambridge, The MIT Press.

Chomsky, Noam: 1986b, *Knowledge of Language: Its Nature, Origin and Use*, New York, Praeger.

Chomsky, Noam: 1982, *Some Concepts and Consequences of the Theory of Government and Binding*, Cambridge, The MIT Press.

Chomsky, Noam: 1981, *Lectures on Government and Binding*, Dordrecht, Foris Publications.

Dalrymple, Mary: 1991, 'Against Reconstruction in Ellipsis', CSLI Manuscript, Stanford University, Stanford CA.

Davidson, Donald: 1978, 'What Metaphors Mean', *Critical Inquiry* **5**, 31–47. Reprinted (1984) in *Inquiries into Truth and Interpretation*, Oxford, Oxford University Press.

Grice, H. Paul: 1975, 'Logic and Conversation', in P. Cole and J. L. Morgan (eds.), *Speech Acts, Syntax and Semantics*, vol. 3, New York, Academic Press.

Haegeman, L: 1991, *Introduction to Government and Binding Theory*, Oxford, Blackwell.

Hankamer, Jorge and Ivan Sag: 1976, 'Deep and Surface Anaphora', *Linguistic Inquiry* **7**, 391–426.

Lewis, David: 1972, 'General Semantics', in D. Davidson and G. Harman (eds.), *Semantics of Natural Language*, Dordrecht, Reidel.

Montague, Richard: 1974, *Formal Philosophy: Selected Papers of Richard Montague*, edited by R. Thomason, New Haven, Yale University Press.

Morgan, J. L.: 1989, 'Sentence fragments revisited', *Chicago Linguistics Society: Papers from the parasession on language in context* **25**, 228–241.

Napoli, Donna Jo: 1982, 'Initial Material Deletion in English', *Glossa* **16**, 85–111.

[16] Here's an example. Suppose Joe wishes to communicate the set of assumptions {A} by using the sentence *S* in circumstances *C*. Joe realizes that, given his audience's initial assumption-set, {A} is not very accessible in *C*. Indeed, let us assume that, in *C*, {A} is the third most accessible set of assumptions which is relevant enough to warrant processing the utterance. That is, before getting to {A}, the hearer will recover two other (more accessible) sets of assumptions, both of which are relevant enough to warrant the effort expended. Sperber and Wilson maintain that if Joe uses *S* in *C* to make {A} manifest, he will violate the second part of the presumption of optimal relevance. His utterance will not be the most relevant stimulus for making {A} manifest in *C*, because, "almost inevitably", there exists some sentence *S'* such that, given *C*, {A} would be the first interpretation of an utterance of *S'* consistent with the presumption of optimal relevance. Hence, in *C*, an utterance of *S'* is a more relevant stimulus than an utterance of *S*, because an utterance of *S'* does not require the hearer to access two prior sets of assumptions that are relevant enough.

Russell, Bertrand: 1919, *Introduction to Mathematical Philosophy*, London, George Allen & Unwin.

Russell, Bertrand: 1911, 'Knowledge by Acquaintance and Knowledge by Description', *Proceedings of the Aristotelian Society* **11**, 108–128.

Russell, Bertrand: 1905, 'On Denoting', *Mind* **14**, 479–493.

Sainsbury, R. M.: 1979, *Russell*, London, Routledge and Kegan Paul.

Sperber, Dan and Deirdre Wilson: 1995, *Relevance: Communication and Cognition*, second edition, Oxford, Blackwell.

Sperber, Dan and Deirdre Wilson: 1987, 'Précis of *Relevance*', *The Behavioral and Brain Sciences* **10**, 697–710.

Stainton, Robert J.: 1997a, 'Utterance Meaning and Syntactic Ellipsis', *Pragmatics & Cognition* **5**(1), 51–78.

Stainton, Robert J.: 1997b, 'What Assertion Is Not', *Philosophical Studies* **85**, 57–73.

Stainton, Robert J.: 1995, 'Non-Sentential Assertions and Semantic Ellipsis', *Linguistics & Philosophy* **18**, 281–296.

Stainton, Robert J.: 1994, 'Using Non-sentences: An application of Relevance Theory', *Pragmatics & Cognition* **2**, 269–284.

Yanofsky, Nancy: 1978, 'NP Utterances', *Chicago Linguistics Society: Papers from the Regional Meeting* **14**, 491–502.

Carleton University
Department of Philosophy
1125 Colonel By Drive
Ottawa, Ontario
Canada K1S 5B6