Bobaljik, J., and Thrainsson, H., "Two Heads aren't Always Better than One." Msc. Harvard University and the University of Iceland. 1997.

Carston, R., "Implicature, Explicature, and Truth-Theoretic Semantics." In R.M. Kempson (ed.) *Mental Representations*. Cambridge: Cambridge University Press. 1988. Reprinted in Davis, 1991.

Carston, R., "Implicature, Explicature and Truth-Theoretic Semantics." In S. Davis (ed.), *Pragmatics: A Reader*. Oxford: Oxford University Press. 1991: 33–51.

Carston, R., "Explicature and Semantics." *UCLA Working Papers in Linguistics* 12 (2000): 1–44. Forthcoming in S. Davis and B. Gillon (eds.), *Semantics: A Reader*. Oxford: Oxford University Press. 2004.

Chomsky, N., *The Minimalist Program*. Cambridge: MIT Press. 1995.

Cinque, G., *Adverbs and Functional Heads: A Cross-Linguistic Perspective*. Oxford: Oxford University Press. 1998.

Clapp, L. "What Unarticulated Constituents Could Not Be." In J. Campbell, M. O'Rourke and D. Shier (eds.) *Meaning and Truth*. New York: Seven Bridges Press. 2002: 231–256.

Kayne, R., *The Antisymmetry of Syntax*. Cambridge: MIT Press. 1994.

Pollock, J.Y., "Verb Movements, Universal Grammar and the Structure of IP." *Linguistic Inquiry* 20 (1989): 365–424.

Poole, G., "Optional Movement in the Minimalist Program." In W. Abraham et. al. (eds.), *Minimal Ideas: Syntactic Studies in the Minimalist Framework*. Amsterdam: John Benjamins. 1996: 199–216.

Recanati, F., *Direct Reference: From Language to Thought*. Oxford: Blackwell. 1993.

Recanati, F., "Unarticulated Constituents." *Linguistics and Philosophy* 25 (2002): 299–345.

Sperber, D., and Wilson, D., *Relevance: Communication and Cognition*. Oxford: Blackwell. 1986.

Stainton, R., "In Defense of Non-Sentential Assertion." In Z. Szabo (ed.), *Semantics vs. Pragmatics*. Oxford: Oxford University Press. 2004: 383–457.

Stanley, J., "Context and Logical Form." *Linguistics and Philosophy* 23 (2000): 391–434.

LENNY CLAPP

ON THE INTERPRETATION AND PERFORMANCE OF NON-SENTENTIAL ASSERTIONS*

> What is it that we call a sentence? A series of sounds, but only if it has a sense (this is not meant to convey that *any* series of sounds that has a sense is a sentence). And when we call a sentence true we really mean that its sense is true. And hence the only thing that raises the question of truth at all is the sense of sentences.
>
> Gottlob Frege, "Thoughts"

1. INTRODUCTION

The issue that is the focus of this anthology can be characterized in terms of the above citation from Frege. Some, whom I will call *sententialists*, maintain that the view espoused by Frege in the above passage is basically correct; sententialists agree with Frege that utterances of *only* complete declarative sentences can be true (or false), though they would reject Frege's Platonist conception of *sense*. Others, whom I will call *non-sententialists*, maintain that Frege's view is fundamentally flawed, and not merely because he assumes an implausible Platonism regarding sense; against Frege, they maintain that utterances of non-sentential words or phrases can also be true (or false).

It is noteworthy that contemporary sententialists and non-sententialists alike reject Frege's Platonism and claim to be pursuing a much different project than Frege claimed to be pursuing. In the pages preceding the above citation, Frege took pains to distinguish his project from any sort of psychological investigation; Frege very much desired to distinguish the "laws of psychology" and the "laws of logic," and his project was to discover the laws of logic. Thus Frege's attitude is in stark contrast with that shared by both the contemporary sententialists and non-sententialists, who are working within a tradition of empirical linguistics — a discipline proudly defined as a branch of psychology or cognitive science. As a consequence, sententialists and non-sententialists alike must be concerned with sorts of natural language phenomena that Frege could, perhaps with justification, disregard as being mere psychological glitches, or deficiencies of natural language. In particular, if confronted with the at least apparent phenomenon of non-sentential assertion, Frege could say that regardless of whether or not the laws of psychology allow for non-sentential assertions, the laws of logic do not. The contemporary sententialist, however, cannot avail himself of this sort of response. How, then, is the contemporary sententialist, who works within the empirical framework established by Chomsky, to account for the apparent phenomenon of non-sentential assertion? As Stainton and Elugardo note in the introduction to this volume, there are two general strategies of response available to the sententialist, both of which involve denying that

non-sentential assertion is a genuine phenomenon:

> *The Syntactic Strategy*: The sententialist can claim that cases of what appear to be non-sentential utterances expressing truth conditions actually involve some sort of *ellipsis*, and thus what is uttered is really a full sentence — the LF for the utterance has IP as its initial node. This strategy thus grants that the utterances in question express truth conditions, but it denies that they involve sub-sentential syntactic structures.[1]

> *The Pragmatic Strategy*: The sententialist can claim that cases of what appear to be non-sentential utterances expressing truth conditions actually do not express truth conditions at all. This strategy thus grants that the utterances in question involve sub-sentential syntactic structures, but it denies that such utterances really express truth conditions.

As Stanley (2000, 403–4) has pointed out, the sententialist need not commit to only one of these strategies: it is open to him to utilize the syntactic strategy to account for some non-sentential utterances, and to utilize the pragmatic strategy to account for others.[2] My purpose here, however, is to argue that the pragmatic strategy is not a live option for the sententialist. Thus if the sententialist is to succeed in explaining away apparent instances of non-sentential assertions, he must utilize the syntactic strategy, which faces significant difficulties of its own.[3]

The paper proceeds as follows. In section II I describe the general perspective of *truth-conditional semantics* which motivates sententialism, and I distinguish it from *truth-conditional pragmatics*, the general perspective that underlies non-sententialism. I also explicate the model of interpretation that is inherent in truth-conditional semantics.[4] In section III, I illustrate that there are many sorts of *prima facie* counterexamples to truth-conditional semantics in addition to those involving non-sentential utterances. It is important to keep these other sorts of *prima facie* counterexamples in mind because, given the similarities between the various sorts of *prima facie* counterexamples, a sententialist's response to one sort will commit him to a similar response to another sort. And this wider commitment to a strategy of response might have significant consequences. Indeed, this is precisely what I will argue to be the case regarding the pragmatic strategy of response. Thus, in section IV, I explicate Stanley's (2000) utilization of the pragmatic strategy, and I argue that if the sententialist utilizes anything like Stanley's response to account for problematic non-sentential utterances, then he must also utilize this response to account for other sorts of *prima facie* counterexamples. But this wider commitment to the pragmatic strategy is incompatible with the model of interpretation inherent in truth-conditional semantics. In section V, I consider and reject Stanley's appeal to a semantic competence/performance distinction to support his utilization the pragmatic strategy. And finally, in section VI, I briefly consider the consequences of the failure of the pragmatic strategy for the issue of whether or not there are genuine non-sentential assertions, and for the more general debate between truth-conditional semantics and truth-conditional pragmatics.

2. TRUTH-CONDITIONAL SEMANTICS, TRUTH-CONDITIONAL PRAGMATICS, COMPETENCE, PERFORMANCE AND INTERPRETATION

The sententialist claims that only utterances of complete sentences can be assigned truth conditions. Or slightly more precisely, the sententialist claims that only declarative utterances whose LFs are fully sentential — whose initial node is an IP — have truth conditions, and thus only such full sentences (at LF) can legitimately be used to make assertions. What is the motivation for the sententialist's claim?

The sententialist's claim is a consequence of the general principle that "all truth-conditional effects of extra-linguistic context can be traced to logical form" (Stanley, 2000, 391). This general principle is more precisely rendered as follows:

> *Utterance Compositionality*: The truth conditions of an utterance are a *function* of (i) the structure of the LF of the utterance, and (ii) the *semantic values* of the terminal nodes of the LF of the utterance (as determined by the context of utterance).

Genuine non-sentential assertions would constitute counterexamples to this principle. Consider two typical non-sentential utterances, each of which involves an articulation of the adjectival phrase, 'totally useless': (i) Suppose I am giving you an assessment of the computer equipment in my department. Pointing at a particular printer I utter, 'totally useless'. In this context my utterance of this adjectival phrase seems to constitute an assertion and thus seems to express truth conditions — my utterance is true if and only if, roughly, the indicated printer is totally useless. (ii) Suppose I am in a meeting, and all the participants are voicing their opinions about a certain policy that has just been proposed. When it is my turn, I utter 'totally useless'. Again, my utterance seems to constitute an assertion and thus seems to express truth conditions — in this context my utterance is true if and only if, roughly, the recently proposed policy would be totally useless. That two utterances of the very same phrase express distinct truth conditions is problematic for the principle of utterance compositionality. For it at least *seems* that my utterances have the same LF, and, given that my utterances involve the very same (non-context-sensitive) lexical items, my utterances must invoke the very same semantic values. Yet, contrary to what would be predicted by utterance compositionality together with these apparent facts, my utterances express distinct truth conditions. Consequently, the defender of utterance compositionality must claim either that, despite phonological and syntactic appearances, my utterances are associated with distinct LFs, or he must claim that, despite semantic and pragmatic appearances, my utterances do not express truth conditions. That is, the defender of utterance compositionality must utilize close cousins of either the syntactic or the pragmatic strategy.[5]

The above examples illustrate why non-sentential utterances pose a threat to utterance compositionality, but they do not quite make explicit why one who endorses utterance compositionality is also committed to sententialism. The above examples illustrate that if utterance compositionality is to be preserved, then either, despite appearances, my two utterances of 'totally useless' do not have the same LF, or they do not express the truth conditions they seem to express. But *this much* does not commit the defender of utterance compositionality to the sententialist claim that only utterances whose LFs are fully sentential express truth conditions. Why does utterance compositionality imply this additional, stronger, claim?

The reason is that sub-sentential words and phrases often appear within the larger syntactic environment of a full sentence. That is, the adjectival phrase 'totally useless' appears phonetically realized in sentences such as 'Software written before 1990 is now totally useless'. And when it occurs embedded in this broader syntactic environment, the adjectival phrase does not express truth conditions; within this sentential (IP dominated) syntactic environment, the adjectival phrase 'totally useless' is not assigned truth conditions as its semantic value. Rather, as it occurs within this larger fully sentential syntactic environment, the semantic import of the phrase 'totally useless' is merely to contribute to the determination of the truth conditions of the full sentence of which it is a proper part. Moreover, it is typically required that the semantic value assigned to a syntactic structure remain constant regardless of any broader syntactic environment

in which the structure might appear. Following Davidson (1968) this requirement is sometimes referred to as "semantic innocence."[6] Consequently, if 'totally useless' (or its LF) is not assigned truth conditions when it occurs embedded in a sentence (i.e., in an LF whose initial node is IP), then it cannot be assigned truth conditions when it occurs as a non-sentential utterance. To summarize, semantic innocence requires that if *some* occurrences of the adjectival phrase 'totally useless' do not express truth conditions, then (assuming that the phrase is not radically context sensitive) *all* occurrences of the phrase do not express truth conditions. Moreover, the general compositionality requirements of traditional truth-conditional semantic theories dictate that the adjectival phrase 'totally useless' does not express truth conditions when it appears in a sentential (IP dominated) syntactic environment. And hence the commitment to utterance compositionality together with semantic innocence imply the sententialist's claim that only utterances whose LFs are fully sentential express truth conditions.

The sententialist's claim is thus a consequence of utterance compositionality, together with other plausible constraints on an adequate semantic theory. But what is the motivation for utterance compositionality itself? The commitment to this principle derives from a particular construal of the role of compositionality in the process of *interpretation* — the process whereby speaker-hearers actually come to understand one another's utterances. It is now a familiar idea that our interpretative abilities result, at *least in part*, from our implicit knowledge of semantic rules and principles. Some of these rules and principles, the lexical rules, dictate what the meanings (or *semantic values*) of individual words or morphemes, relative to a context of utterance, are. Others, the combinatorial rules, dictate how the meanings of words and morphemes combine, as directed by the LF of the sentence, to determine the meaning, or truth conditions, of an entire sentence. According to this general compositional conception of semantics then, a semantic theory for a language is a specification of implicitly known compositional rules and principles that in combination assign appropriate truth conditions to every sentence of the language. Larson and Segal (1995, 11–12) summarize the attractions of this general compositional conception:

> The hypothesis that we know a set of compositional semantic rules and principles is a highly attractive one having a great deal of explanatory power. In particular, it accounts for three notable and closely related features of linguistic competence. First, it explains why *our understanding of sentences is systematic* — why there are definite, predictable patterns among the sentences we understand. . . . Second, the hypothesis accounts for the obvious but important fact that *we can understand new sentences*, sentences that we have never come across before. . . . Third, the hypothesis accounts for the slightly less obvious but equally important fact that *we have the capacity to understand each of an indefinitely large number of sentences*.

The Frege-inspired truth-theoretic semantic programs defended and developed by Davidson, Montague and their followers are motivated by considerations such as those stated above. It is important to realize, however, that these familiar arguments in support of some sort of semantic compositionality do not entail *utterance compositionality*. The above cited familiar motivations for some sort of compositional semantic theory support utterance compositionality *only if* such a compositional semantic theory is construed as a theory of linguistic *performance*, as opposed to a theory of linguistic *competence*. That is, one might maintain that though implicit knowledge of compositional semantic rules and principles is *necessary* for interpretation, it is nowhere near *sufficient*. One might maintain that implicit knowledge of such compositional semantic

rules and principles is just one part of the knowledge and abilities speakers utilize in making judgments concerning the truth conditions of utterances, and that other — more *pragmatic* — knowledge and abilities are also utilized. Moreover, if one regarded such a compositional semantic theory as a theory of only semantic *competence*, as opposed to semantic *performance*, one could adopt something akin to Frege's attitude toward "aberrant" linguistic phenomena; one could maintain that certain linguistic phenomena are beyond the scope of a theory of semantic competence. In particular one could maintain that our purely *semantic* knowledge, processes, and systems do not allow for non-sentential utterances, and thus to explain how we use non-sentential utterances one would have to invoke *pragmatic* knowledge, processes, and systems. Hence, there is an important distinction between those who regard compositional semantic theories of the sort proposed by Davidson and Montague as theories of *semantic performance*, and those who regard them as theories of *semantic competence*.

Truth-conditional semantics is the view that such traditional compositional semantic theories ought to be regarded as theories of *performance* — as theories explaining how speaker-hearers actually manage to interpret one another's *utterances*. Thus, those who endorse truth-conditional semantics endorse *utterance compositionality*: they maintain that the truth condition of utterances are determined by *semantics* alone, where *semantics* is limited to determination of LFs, the processes of assigning semantic values to terminal nodes of LFs, and finally the computation of the truth conditions thereby determined.[7] In contrast, those who regard traditional semantic theories as theories of only *semantic competence* deny that such theories suffice as theories of *interpretation*. Following Stanley (2000) I will call such theorists *truth-conditional pragmatists*. These theorists maintain that purely *semantic* (where this term is understood as described above) factors do not suffice to determine the truth conditions of utterances, and that additional *pragmatic* factors are also required to determine truth conditions. Hence truth-conditional pragmatists reject utterance compositionality, and therefore they can maintain something akin to Frege's attitude toward "aberrant" linguistic phenomena; in particular, truth-conditional pragmatists can maintain that the "aberrant" phenomenon of non-sentential assertion is accounted for by pragmatic knowledge, processes and systems that go beyond the knowledge, processes and systems described by traditional compositional semantic theories. In summary, both truth-conditional semanticists and truth-conditional pragmatists accept the recently rehearsed reasons supporting a traditional compositional semantic theory, though they disagree as to whether such a semantic theory is to be understood as only a theory of *semantic competence*, or as a more comprehensive theory of *semantic performance*: truth-conditional semantics understands such a compositional semantic theory to be a theory of semantic performance, while truth conditional pragmatics understands it to be only a theory of semantic competence. Sententialism is supported by the broader perspective of *truth-conditional semantics*, while non-sententialism is supported, or at least allowed for, by the contrasting broader perspective of *truth-conditional pragmatics*.

Because truth-conditional semantics is proposing a theory of semantic performance, it entails a particular model of *interpretation* — the process whereby speaker-hearers actually come to understand one another's utterances. This model of interpretation is explicitly described by Stanley and Gendler-Szabo (2000, 11). According to this model, interpretation of "typical assertions" is a two-step process whereby a hearer identifies the proposition the speaker intends to communicate, or equivalently determines the

truth conditions of an assertion. In the first step the hearer uses her syntactic and phonological knowledge, together with whatever clues she can garner from the context of utterance, to determine the LF of the assertion. Stanley and Gendler Szabo (2000, 13) use the equation, "what is articulated + context = what is uttered" to describe this first step, where "what is articulated" is a "phonological sentence," and "what is uttered" is a "grammatical sentence," i.e., an LF. Thus if an interpreter correctly completes the first step, she will have determined the LF of an utterance. In the second step the hearer uses her knowledge of the LF of the utterance, together with her knowledge of the semantic theory for her language and knowledge of the context of utterance, to determine the proposition expressed, or equivalently the truth conditions of the utterance. Stanley and Gendler Szabo (2000, 15) use the equation "what is uttered + linguistic meaning + context = what is said" to describe the second step, where "what is said" is the proposition expressed, or equivalently the truth conditions expressed. I think it is more perspicuous to conceive of this second step as itself proceeding in two sub-steps: In the first sub-step, the interpreter uses her recently acquired knowledge of what is uttered (the LF), and her knowledge of the lexical semantic rules for her language and her knowledge concerning relevant the context of utterance, to determine the semantic values of the semantically significant features of what is uttered (the lowermost nodes of the LF). And in the second sub-step the interpreter uses her knowledge of the thus determined semantic values, together with her knowledge of the combinatorial semantic rules of her language and her knowledge of what is uttered (the LF), to determine "what is said," i.e., the truth conditions of the utterance. In keeping with the principle of utterance compositionality, this two-step model of interpretation requires that the truth conditions interpreters assign to "typical assertions" must be "traced to logical form."

3. OTHER SORTS OF *PRIMA FACIE* COUNTEREXAMPLES TO UTTERANCE COMPOSITIONALITY

An adequate understanding and treatment of the problem non-sentential utterances pose for truth-conditional semantics requires consideration of other similar problems. It is not difficult to generalize from the phenomenon of non-sentential utterances to formulate a general description of *prima facie* counterexamples to utterance compositionality. Let an *expression* be a phonological type — a sound-type constituting a linguistic entity; hence some expressions correspond to complete sentences, others to mere words or phrases. (This is hardly precise, but it will serve my purposes.) Suppose an expression S *at least seemed* to have the following four properties:

(a)   Some occurrences of S express truth conditions.
(b)   S is context sensitive so that the truth conditions it expresses vary from context to context, or perhaps in some contexts S does not express truth conditions at all.
(c)   S is neither lexically nor structurally ambiguous.
(d)   S contains no context-sensitive words and/or features that account for its context sensitivity.

Such an expression would constitute a *prima facie* counterexample to utterance compositionality. For if S possessed (c), every occurrence of S would have the same LF. And if

S possessed (c) and (d), every occurrence of S would invoke the same semantic values. Hence, utterance compositionality would imply that either no occurrence of S expresses truth conditions, or that every occurrence expresses the same truth conditions. For this principle entails that the truth conditions of *every occurrence* of S are a function of (i) the structure of S's LF, and (ii) the *semantic values* of the terminal nodes of S's LF. So if S also possesses (a) and (b), S constitutes a *prima facie* counterexample to utterance compositionality. If utterance compositionality is to be preserved, the defender of truth-conditional semantics must demonstrate that the *prima facie* counterexample does not, despite appearances, actually possess at least one of (a)–(d).

In section I the *pragmatic* and *syntactic* strategies were defined relative to *prima facie* counterexamples involving non-sentential utterances. It is now apparent, however, that non-sentential utterances are just one instance of a general problem for truth-conditional semantics. Consequently, the syntactic and pragmatic strategies generalize in a straightforward way so that they apply to all sorts of *prima facie* counterexamples:

> *The Syntactic Strategy*: Faced with a *prima facie* counterexample S, the defender of truth-conditional semantics claims that the LFs corresponding to occurrences of S contain phonetically unrealized structure, and thus the LFs corresponding to utterances of S are richer than is suggested by the phonological features of S. (In terms of the definition of a *prima facie* counterexample stated above, despite appearances the *prima facie* counterexample does not actually possess property (c) or (d).)[8]

> *The Pragmatic Strategy*: Faced with a *prima facie* counterexample S, the defender of truth-conditional semantics claims that occurrences of S do not actually express the truth conditions they seem to express, or do not actually express truth conditions at all. (Again, in terms of the definition stated above, despite appearances the *prima facie* counterexample does not actually possess property (a) or (b).)

The *prima facie* counterexamples to utterance compositionality are by no means limited to cases of non-sentential utterances. As Travis (1985), Sperber and Wilson (1986), Carston (1991), Bach (1994), Recanati (1996) and other proponents of *truth-conditional pragmatics* have demonstrated, there are many expressions that at least seem to have properties (a)–(d). The list of *prima facie* counterexamples includes almost all quantified sentences, sentences containing comparative adjectives, propositional attitude ascriptions, sentences containing definite descriptions, modal sentences, counterfactuals, and others. The focus of this paper is non-sentential utterances, and thus I will not review all of these sorts of *prima facie* counterexamples. But my argument against the pragmatic strategy does require the premise that if the defender of utterance compositionality utilizes the pragmatic strategy to explain away the *prima facie* counterexamples involving non-sentential utterances, then he must utilize this strategy with regard to other sorts of *prima facie* counterexamples as well. And making a case for this premise requires me to explicate at least some of the other sorts of *prima facie* counterexamples, and to illustrate how the syntactic strategy can be utilized in an attempt to explain them away. So in the remainder of this section I first discuss the general phenomenon of utterances that require, in the terminology of Bach (1994), "completion" or "expansion" — all such utterances constitute *prima facie* counterexamples to utterance compositionality. I also consider a particular sort of utterance that requires "expansion" — quantified sentences — and I briefly explicate Stanley and Gendler Szabo's (2000) utilization of the *syntactic strategy* in an attempt to explain away *prima facie* counterexamples involving quantified sentences.

Bach (1994) introduces the general notions of "completion" and "expansion" in terms of different sorts of *prima facie* counterexamples to utterance compositionality. Consider typical utterances of the following sentences:

(1)   John left (the party/graduate school)
(2)   Spike and Butch got in a fight (with Bruno/with each other)

Interpreters of typical utterances of (1) and (2) must work through a process of *enrichment* to arrive at the intended truth conditions of the utterance. (In the above, I have indicated possible enrichments in parentheses.) Bach (1994) suggests that with regard to utterances of sentences such as (1) interpreters must work through a process of *completion* to arrive at truth conditions. The terminology is meant to reflect that, though (1) is *grammatically* a complete sentence, the phonologically realized elements of (1) are insufficient to determine truth conditions. The idea is that one cannot evaluate the thought that John left *simpliciter* for truth or falsity — one must know what it is that John is being alleged to have left. And what it is that an utterer of (1) is alleging John to have left will vary across contexts. Moreover, (1) contains no relevant, phonologically realized, context-sensitive words or features. Thus sentences such as (1) that require completion constitute *prima facie* counterexamples to utterance compositionality. (Or more precisely, *expressions* corresponding to sentences such as (1) constitute *prima facie* counterexamples.)

According to Bach sentences such as (2) are not in need of semantic completion, because the phonologically realized material is itself sufficient to express a "minimal proposition." The minimal proposition expressed by occurrences of (2) is simply that Spike got in a fight, and Butch got in a fight. This minimal proposition, however, is not what speakers typically express in uttering (2). Usually, but not always, when speakers utter (2) they express the proposition that Spike and Butch got in a fight *with each other*. So though the phonetically realized material in (2) is sufficient to determine a minimal proposition, the truth conditions typically expressed by an occurrence of (2) are more discriminating than this minimal proposition, and thus interpreters must work through a process of *expansion* in order to determine the expressed truth conditions. Consequently expressions corresponding to sentences such as (2) also seem to possess properties (a)–(d): These expressions are used to make assertions, and they are context-sensitive, but they contain no relevant context-sensitive features and involve neither lexical nor structural ambiguity. Consequently such expressions constitute *prima facie* counterexamples to utterance compositionality.[9]

Quantified sentences are one sort of sentence whose corresponding expressions typically require completion, and Stanley and Gendler Szabo (2000) have utilized the syntactic strategy in an attempt to explain away such *prima facie* counterexamples. Consider the quantified sentence discussed by Stanley and Gendler Szabo:

(3)   Every bottle is empty.

A typical utterance of (3) is in need of expansion. For a typical utterance of (3) does not express the minimal proposition that every bottle in the universe is empty; rather a typical utterance of (3) states a weaker, richer, proposition to the effect that every bottle relevant to the people engaged in the discourse is, to some relevant degree, empty of some relevant substance. Thus the domain of quantification does not include *all*

bottles, but is restricted to a proper subset of relevant bottles. But the proper subset of bottles that serves as the restriction varies from context to context. For example, one utterance of (3) is true, roughly, if and only if every wine-bottle on my dinner table at a particular time is empty of wine. But a different utterance of (3) is true if and only if, roughly, every baby-bottle within easy reach of a desperate parent is empty of baby formula. Thus (1) clearly has properties (a) and (b), and it seems to have properties (c) and (d): (3) is neither structurally nor lexically ambiguous. And though (3) contains at least one overt context-sensitive feature, viz. the tensed form of the verb, there is no overt context-sensitive feature that would plausibly account for the variance in the quantifier restriction. Thus sentence (3) seems to possess all of (a)–(d), and consequently it constitutes a *prima facie* counterexample to utterance compositionality.

A closely related phenomenon concerns "incomplete" definite descriptions. Consider the following sentence:

(4)   The bottle is empty.

Assuming that definite descriptions presuppose uniqueness, it seems that an utterance of (4) can express truth conditions only if there exists just one bottle in the universe.[10] But typical utterances of (4) do not seem to presuppose that there exists only one bottle in the universe, thus utterances of (4) also require interpreters to work through a process of expansion. Thus (4) also seems to possess properties (a)–(d) and constitutes a *prima facie* counterexample to utterance compositionality. Clearly (4) can be used to express truth conditions, and moreover, different truth conditions in different contexts: In some contexts (4) is true if and only if, roughly, a particular wine-bottle is empty, while in other contexts (4) is true if and only if, roughly, a particular baby-bottle is empty. But again (4) contains on relevant context-sensitive words or features, and it is neither lexically nor structurally ambiguous. So sentences with "incomplete" definite descriptions, or rather their corresponding expressions, also possess properties (a)–(d), and thus also constitute *prima facie* counterexamples to utterance compositionality.

Stanley and Gendler Szabo (2000) utilize a version of the syntactic strategy in an attempt to explain away the *prima facie* counterexamples of the sort exemplified by (3) and (4). Stanley and Gendler Szabo propose that an expression such as (3) be analyzed as containing at the level of LF a "hidden indexical" which takes on different semantic values in different contexts. More specifically, they propose the LF of an utterance of (3) is something like this
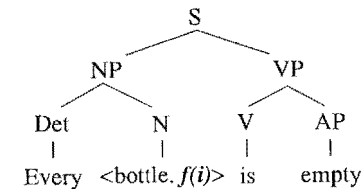
According to Stanley and Gendler Szabo's analysis, the terminal node corresponding to the phonetically realized noun 'bottle' is syntactically complex: It is an ordered pair,

the first member of which is the phonetically realized familiar lexical item 'bottle' and the second member of which is a new sort of phonetically unrealized context-sensitive element $f(i)$. This phonetically unrealized element consists of two parts: $f()$ is a context-sensitive element that has as its semantic value, relative to a context, a function from individuals to sets (or properties), while i is a context sensitive element that has as its semantic value, relative to a context, an individual. The set that is determined by applying the function "provided by context" to the individual provided by context serves to further restrict the domain of quantification: the restricted domain is the intersection of the extension of 'bottle' and the set determined by applying the function "provided by context" to the argument "provided by context." (Stanley and Gendler Szabo, not surprisingly, provide no explanation whatsoever as to how the semantic values of $f()$ and i are "provided by context.") Thus there is a semantic value invoked by an utterance of (3) which — though not the semantic value of any phonetically realized, or articulated, word or feature in (3) — is nonetheless the semantic value of an element of the utterance's LF. Moreover, since $f(i)$ is a context-sensitive element, it is assigned different semantic values in different contexts. In terms of the previous example involving different utterances of (3), in some contexts $f(i)$ is assigned as its semantic value the set of all wine bottles on my dinner table, while in other contexts $f(i)$ is assigned as its semantic value the set of all baby-bottles within reach of a certain desperate parent. Hence on Stanley and Gendler Szabo's utilization of the syntactic strategy, (3) does not actually possess property (d), for on their analysis, despite appearances, occurrences of (3) really do contain relevant context-sensitive features. A similar sort of explanation of course applies to sentences such as (4) containing "incomplete" definite descriptions. In this way Stanley and Gendler Szabo utilize the syntactic strategy to explain away *prima facie* counterexamples involving quantified sentences.

4. THE PRAGMATIC STRATEGY OF RESPONSE AND ITS APPLICATION TO ALL OF THE *PRIMA FACIE* COUNTEREXAMPLES

To utilize the pragmatic strategy in response to a *prima facie* counterexample is to argue that, despite appearances, occurrences of the problematic expression do not actually express truth conditions, or at least not the truth conditions they seem to express. In this section I will argue that the pragmatic strategy is not a live option for the defender of utterance compositionality. My comments will focus on Stanley (2000), because Stanley is the only theorist I am aware of who utilizes the pragmatic strategy with regard to some *prima facie* counterexamples involving non-sentential utterances. It will become apparent, however, that the problems I raise with regard to Stanley's utilization of the pragmatic strategy undermine *any* attempt to utilize the pragmatic strategy to defend utterance compositionality from *prima facie* counterexamples.

Stanley (2000) maintains that, despite the judgments of ordinary speaker-hearers, occurrences of non-sentential utterances do not actually express truth conditions. He presents two arguments in support of this claim. He first suggests that "linguistic speech acts must determinately be made with the relevant sort of force. That is, for an act to count as a speech act of kind k, it must determinately be performed with the force appropriate to acts of kind k" (407). Stanley then uses this criterion of *determinate force* to argue that a particular case of non-sentential utterance is not a "linguistic

assertion, and indeed is not a genuine linguistic speech act" (Stanley, 2000, 407). In the case considered by Stanley, a thirsty man staggers up to a street vendor and utters, 'water'. Stanley argues that in this case it is indeterminate whether or not the thirsty man's utterance has the force of a request, or a command: "It would be equally consistent with the thirsty man's intentions to suppose that the utterance was a request, or a command" (Stanley, 2000, 407). And since the thirsty man lacks intentions that would determine a force appropriate to the act kind *assertion*, the utterance does not count as an assertion, and thus does not count as a "genuine linguistic speech act."

Unfortunately, Stanley's discussion of the thirsty man's utterance of 'water' takes the debate rather far astray. First, even if one grants Stanley's conclusion that the thirsty man's utterance does not count as a "genuine linguistic speech act," it is far from clear how this serves to rescue utterance compositionality. This principle is threatened because it predicts that non-sentential utterances do not express truth conditions, yet many such utterances seem to express truth conditions. Hence, what is relevant is whether or not the thirsty man's utterance expresses truth conditions — whether or not it counts as an *assertion*, or any other "genuine linguistic speech act" is beside the point.[11] Second, though I agree with Stanley that the thirsty man's utterance is not an assertion, and moreover does not express truth conditions, nothing relevant to the debate about non-sentential utterances follows from this.[12] Stanley is allegedly presenting a reason for supposing that "apparent non-sentential assertions" (Stanley, 2000, 407) do not really express truth conditions at all. The problem is that if an utterance lacked the force appropriate to assertion, then it simply would not be an "apparent non-sentential assertion" — it would not be a *prima facie* counterexample. Nobody claims that just any old utterance of a word or phrase constitutes a *prima facie* counterexample to utterance compositionality. Hence the example of the thirsty man *cannot* pose a problem for utterance compositionality; it does not even *seem* to express truth conditions. Only utterances that *seem* to express truth conditions can be *prima facie* counterexamples. In order to explain away such problematic utterances, Stanley must provide a reason for thinking that though such utterances *seem* to express truth conditions, they do not *really* express truth conditions. Stanley's first reason does not even attempt to do this.

Stanley's second reason for supposing that some apparent non-sentential assertions are not really genuine assertions is that linguistic speech acts "must express determinate contents" (407). Stanley then uses this criterion of *determinate content* to argue again that the thirsty man's utterance of 'water' to a street vendor is not a genuine linguistic speech act:

> . . . in the case of the thirsty man's utterance . . . there is no determinate content associated with the speech act. Suppose, for the sake of argument, that the speech act is an assertion. Then, the relevant sort of content is a proposition. But what proposition has thereby been expressed? The point is particularly acute if we assume that propositions are structured. Is the proposition thereby expressed the proposition that the thirsty man wants water? Is it the proposition that the vender should give the thirsty man water? The available facts simply do not determine a determinate propositional content for the alleged assertion. And when a communicative act lacks a determinate content, it is not a linguistic act. (408)

Stanley again mistakenly focuses on whether or not the thirsty man's utterance qualifies as an *assertion*, or any other sort of "linguistic act." Such issues of illocutionary

force are, to repeat, beside the point — the relevant issue is whether or not the thirsty man's utterance really expresses truth conditions. But Stanley's second reason, unlike his first reason, can be interpreted so that it is relevant. That is, Stanley can be interpreted as endorsing a *criterion of determinate content*: An expression expresses truth conditions only if it has "determinate content."

If having "determinate content" is a necessary condition for an utterance's expressing truth conditions, then not only will the thirsty man's utterance not qualify as expressing truth conditions, but *all* non-sentential utterances that are not *obvious* cases of syntactic ellipsis will not qualify as expressing truth conditions. Consider again my seeming assertion of 'totally useless' in the policy meeting. Precisely which determinate proposition have I expressed? That the policy just proposed would be totally useless? That the policy they proposed is totally useless? That the policy being considered in this meeting would be totally useless, if we adopted it? Stanley's insight that speakers' intentions are not rich enough to determine how non-sentential assertions are to be *completed* yields the result that *no prima facie* counterexample involving a non-sentential utterance will qualify as expressing truth conditions. But the consequences of Stanley's insight extend even further: If having "determinate content" is a necessary condition for an utterance's expressing truth conditions, then no utterance requiring any sort of *completion or expansion* will qualify as expressing truth conditions.

Consider again a typical utterance requiring *expansion*:

(4)   The bottle is empty.

Suppose this sentence is uttered by a desperate parent who is bottle-feeding a fussy child. And suppose that this utterance constitutes a *prima facie* counterexample — it is an apparently successful act of communication. To fix intuitions, suppose that as a consequence of uttering (4) the other parent says "OK," and is led to bring the speaker another baby-bottle, sufficiently full of baby formula. According to Stanley the utterance of the desperate parent expresses truth conditions only if it expresses "determinate content." But here again Stanley's insight clearly applies. Which, if any, of the following best expresses the "determinate content" of the speaker's assertion?

(4a)   The bottle (in my hand) is empty.
(4b)   The bottle (I have been using) is empty.
(4c)   The bottle (I am looking at) is empty.
(4d)   The bottle (in little Suzie's mouth) is empty.
(4e)   The bottle (right here) is empty.
(4f)   The bottle (I have been using just now to feed our child) is empty.

I have made the point elsewhere (Clapp, 2001), so I will not belabor it here, but clearly there is no more reason for thinking that, e.g., (4a) is the "determinate content" than there is for thinking that (4c) is the "determinate content." It is of course determinate *which bottle* the parent is *denoting* in uttering (4) — the utterance is successful. What is not determinate is how the "incomplete definite description" is made "complete," or more generally, how the quantifier is further restricted. Moreover, there is nothing extraordinary about the utterance of (4) we have been considering, and thus it is clear that this sort of indeterminacy is present in almost every utterance containing a quantifier term. And neither is this indeterminacy limited to utterances involving quantifiers.

Consider again a typical utterance requiring *completion*:

(1)   John left.

Suppose we are at a painfully dull party, and somebody utters (1) in response to an inquiry as to the whereabouts of John. What is the "determinate content" of this utterance? Again, here are a number of plausible candidates:

(1a)   John left (this dull party)
(1b)   John left (the party)[13]
(1c)   John left (this place)
(1d)   John left (this apartment)
(1e)   John left (the apartment we are now in)
(1f)   John left (the party we are now at)

Again, in a typical utterance of (1) there is no more reason to suppose that, e.g., (1a) represents the "determinate content" or the utterance than there is to suppose that (1d) represents the "determinate content" of the utterance.

It would seem then that "indeterminacy of content" is, as Frege might have put it, a widespread deficiency of natural language. In particular, almost every *prima facie* counterexample to utterance compositionality will suffer from this deficiency; having "indeterminate content" is almost an essential feature of *prima facie* counterexamples.[14] The defender of utterance compositionality might suppose this to be a beneficial result; Stanley presented the criterion of *determinate content* as a means of explaining away only some *prima facie* counterexamples involving non-sentential utterances. But as it turns out the criterion can be used, indeed *must* be used, to explain away (almost) all *prima facie* counterexamples. Not surprisingly, the widespread applicability of the criterion of determinate content has a number of significant consequences for the defender of truth-conditional semantics.

One significant consequence is that the defender of truth-conditional semantics who follows Stanley in endorsing the "determinate content" criterion for expressing truth conditions has no need of the *syntactic strategy*. In other words, one consequence of Stanley's rejection of non-sentential utterances on the grounds that they do not have "determinate content" is that Stanley and Gendler Szabo's utilization of the syntactic strategy to account quantifier domain restriction is rendered otiose, as is Stanley's (2000) appeal to a process of *pragmatic ellipsis*.[15] If utterances of sentences containing quantifiers do not constitute genuine assertions because they lack "determinate content," then there is no need to posit sophisticated hidden syntactic structure to preserve utterance compositionality. It might be claimed, however, that this is good news for the defender of truth-conditional semantics. For such "hidden indexical" theories seem *ad hoc* and problematic for independent reasons.[16] If the *prima facie* counterexamples can be explained away without positing hidden syntactic elements and mysterious processes whereby "context provides" semantic values for these hidden elements, so much the better for truth-conditional semantics.

Another significant consequence of the widespread applicability of Stanley's criterion of "determinate content" is that it commits the defender of truth-conditional semantics to *semantic minimalism*: According to semantic minimalism, utterances that undergo a process of *completion* do not really express truth conditions at all, and

utterances that undergo a process of *expansion* express only the minimal proposition that is semantically encoded in the utterance. Hence according to semantic minimalism the actual truth conditions of utterances are often, perhaps usually, quite different than what ordinary speakers take them to be. For instance, according to semantic minimalism, so long as there is more than one bottle in the universe, (4) cannot really be used to make an assertion, for its presupposition that there is only one bottle is always false.

A good example of semantic minimalism is provided by Salmon's (1986) and Soames' (1987) analysis of attitude ascriptions. Salmon and Soames maintain that, despite interpreters' recalcitrant judgments to the contrary, all utterances of attitude ascriptions report only the "determinate content" encoded in utterances, and consequently attitude ascriptions are *transparent*. Thus, despite speakers' firm judgments to the contrary, occurrences of 'John believes that Twain wrote' and 'John believes that Clemens wrote' express the very same truth conditions. Though, to my knowledge, neither Salmon nor Soames themselves motivates this Russellian analysis of attitude ascriptions by appeal to the criterion of determinate content, it is certainly open for them to do so. That is, against competing Fregean analyses — including "hidden-indexical" analyses — that posit reference to "modes of presentation" of some sort, they could follow Schiffer (1992) in rejecting such theories on the grounds that that they suffer from the "meaning intention problem" (Schiffer, 1992, 512). In brief, Schiffer's meaning intention problem is the fact that speakers do not have intentions that would determine which modes of presentation are referred to in an utterance of an attitude ascription. Therefore, since any content there might be involving modes of presentation is indeterminate, attitude ascriptions cannot involve reference to modes of presentation. The only determinate content expressed by an attitude ascription is the minimal Russellian proposition, and as a result attitude ascriptions must be *transparent*. In this way, following Schiffer's and Stanley's lead, Salmon and Soames could motivate their view that only the minimal, Russellian, propositions encoded in the (mostly) phonetically realized features of an attitude ascription are *really* expressed. As is always the case, any attempt at *expansion* results in a lack of determinate content thereby precluding the expression of truth conditions.

So, if the defender of utterance compositionality follows Stanley in endorsing the determinate content criterion for expressing truth conditions, then he is committed to semantic minimalism, and thus he must claim that fully competent speakers are often, perhaps even usually, incorrect in their judgements of the truth conditions of utterances. Is this consequence acceptable for the defender of truth-conditional semantics? It is not, for it is incompatible with the model of interpretation inherent in truth-conditional semantics. Or rather it is incompatible with that model of interpretation so long as that model is understood as an *empirical* theory attempting to explain how speakers actually interpret utterances, and is not, *a la* Frege, understood as a proposal for reconstructing and improving natural language semantics. Recall that the model of interpretation inherent in truth-conditional semantics proposes that the truth conditions actual interpreters judge actual utterances to have are derived through the following two-step process: In the first step the hearer uses her syntactic and phonological knowledge, together with whatever clues she can garner from the context of utterance, to determine the LF of the assertion. In the second step the hearer uses her knowledge of the LF of the utterance, together with her knowledge of the semantic theory for her language, and knowledge concerning the relevant context of utterance, to determine the proposition expressed,

or equivalently the truth conditions of the utterance. In keeping with utterance compositionality, this two-step model of interpretation requires that the truth conditions interpreters assign to "typical assertions" must be "traced to logical form." We have now seen, however, that if semantic minimalism is true, then this model of interpretation does *not* correctly explain and/or predict speaker-hearers' judgements of truth conditions.

Stanley's utilization of the pragmatic strategy crucially depends upon his endorsement of the criterion of determinate content. This criterion leads to semantic minimalism, and the attendant result that interpreters are often, perhaps usually, incorrect in their judgments of truth conditions. But this result undermines the model of interpretation inherent in truth-conditional semantics, at least in-so-far-as that model is understood as an empirical theory of semantic performance attempting to explain how speakers actually interpret utterances. Like the fabled Viet Nam commander, Stanley has destroyed the village of truth-conditional semantics in attempting to save it.

One might agree that Stanley's way of utilizing the pragmatic strategy to explain away *prima facie* counterexamples involving non-sentential utterances fails, but only because he endorses the criterion of determinate content. Perhaps the problem is with this particular criterion for expressing truth conditions, and not so much with the pragmatic strategy generally. If one could formulate a more discriminating criterion according to which *prima facie* counterexamples involving non-sentential utterances do *not* qualify as expressing truth conditions, but other sorts of *prima facie* counterexample *do* qualify as expressing truth conditions, then perhaps one could apply the pragmatic strategy to non-sentential utterances without having to apply it to all sorts of *prima facie* counterexample. In other words, perhaps a more discriminating criterion for expressing truth conditions would allow the defender of truth-conditional semantics to save the village of truth-conditional semantics without destroying it. As I will now argue, however, the prospects for formulating such a discriminating criterion are bleak.

To argue that problematic non-sentential utterances do not really express truth conditions, one must formulate some sort of *general* criterion for expressing truth conditions that is incompatible with the judgments of actual speaker-hearers. The effect of this criterion will be that all utterances bearing some property, or set of properties, will be deemed to not express truth conditions, even through fully competent speaker-hearers interpret them as expressing truth conditions. But it is difficult to see how the property or set of properties described in the proposed criterion could serve to rule out *only* problematic non-sentential assertions and not other sorts of *prima facie* counterexamples. It would be question begging for the sententialist to maintain that only utterances that constitute complete sentences (at LF) qualify as really expressing truth conditions. The proposed criterion would have to be something along the lines of Stanley's proposed criterion, which does not depend upon the mere fact that the problematic non-sentential utterances are at least apparently non-sentential. But it is difficult to see how any such non-question-begging criterion would not apply to many, perhaps all, sorts of *prima facie* counterexamples. Moreover, once the door is open to this sort of *error theory*, it is difficult to see how it would not trump all utilizations of the syntactic strategy for explaining away *prima facie* counterexamples. Once it is allowed that competent speaker-hearers can be egregiously mistaken concerning the truth conditions of utterances, then, given any *prima facie* counterexample, the pragmatic strategy would be preferable to the syntactic strategy. Given the options of either (i) claiming that a *prima facie* counterexample is just *another* case where

interpreters are wrong, or (ii) positing hidden syntactic machinery to account for the truth of interpreters' judgements concerning the truth conditions of the utterance — where it is wholly mysterious as to how the semantic values for the hidden machinery are "provided by context" — it would seem that (i) would always be preferable.[17]

Consequently it seems that any remotely plausible, non-question-begging, criterion would be applicable to many sorts of *prima facie* counterexample. But, to the extent that the criterion can be used to explain away *prima facie* counterexamples, the two-step model of interpretation inherent in truth-conditional semantics is undermined. If speaker-hearers regularly interpret utterances as having truth conditions that cannot be determined via the two-step model, then obviously speakers regularly do not utilize the two-step model to interpret utterances. So to the extent that the defender of truth-conditional semantics claims that competent speakers make incorrect judgments concerning the truth conditions of utterances, he raises counterexamples to truth-conditional semantics.

5. SEMANTICS AND PRAGMATICS, AND COMPETENCE AND PERFORMANCE

In light of the preceding, it is not obvious that it is even open to a theorist who is attempting to explain how interpreters actually determine the truth conditions of utterances to maintain that often interpreters make mistakes and interpret utterances as expressing truth conditions that they "really" do not express. If one of the tasks of truth-conditional semantics is to explain how interpreters actually do determine the truth conditions of utterances, then it is not clear that it even makes sense to claim that interpreters often *incorrectly* assign truth conditions to utterances. Stanley is aware of this tension in his view, and he attempts to resolve it by invoking a semantic *competence/performance* distinction. After noting that his utilization of the pragmatic strategy is incompatible with how speakers actually do interpret some non-sentential utterances, Stanley writes,

> But this is to be expected. Ordinary discourse often involves the use of complex expressions which would be counted as ungrammatical even by the utterer's own lights. For example, some people regularly start a new sentence halfway through an utterance of another sentence. Such discourse involves few sentences that the utterers themselves would classify as grammatical. It is absurd to suppose that we should count such discourse as grammatical, and thereby modify syntactic theory to account for it, and this despite its (statistically speaking) relative normalcy. It is just as absurd to suppose that our conception of semantics should be modified to account for every communicative action which involves the use of language. (Stanley, 2000, 408).

Stanley is here presenting an argument based upon an analogy between contemporary syntactic theory and truth-conditional semantics. He correctly points out that what speakers often, perhaps usually, utter is according to contemporary syntactic theory ungrammatical. And he correctly maintains that it does not follow from this that contemporary syntactic theory should be rejected. To make such an inference would be to confuse *competence* and *performance*. But, Stanley claims, the case of truth-conditional semantics is analogous. And thus Stanley maintains that although the way speakers often *interpret* utterances is incompatible with the predictions of truth-conditional semantics, it does not follow from this that traditional truth-conditional semantics should be rejected. To make such an inference, Stanley suggests, would be again to confuse *competence* and *performance*.

The analogy, however, does not hold. For, as was explained in section II, the defender of truth-conditional semantics defends a *performance* theory. Consider Chomsky's classic and influential description of the competence/performance distinction and its relevance to linguistics:

> Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of the language in actual performance.... To study actual linguistic performance, we must consider the interaction of a variety of factors, of which the underlying competence of the speaker-hearer is only one.... We thus make a fundamental distinction between *competence* (the speaker-hearer's knowledge of his language) and *performance* (the actual use of language in concrete situations). Only under the idealization set forth in the preceding paragraph is performance a direct reflection of competence. A record of natural speech will show numerous false starts, deviations from rules, changes of plan in mid-course, and so on. (Chomsky, 1965, 3–4.)

And several pages later Chomsky further elaborates on the competence/performance distinction:

> To avoid what has been a continuing misunderstanding, it is perhaps worthwhile to reiterate that a generative grammar is not a model for a speaker or a hearer. It attempts to characterize in the most neutral possible terms the knowledge of the language that provides the basis for actual use of language by a speaker-hearer. When we speak of a grammar as generating a sentence with a certain structural description, we mean simply that the grammar assigns this structural description to the sentence. When we say that a sentence has a certain derivation with respect to a particular generative grammar, we say nothing about how the speaker or hearer might proceed, in some practical or efficient way, to construct such a derivation. These questions belong to the theory of language use — the theory of performance. No doubt, a reasonable model of language use will incorporate, as a basic component, the generative grammar that expresses the speaker-hearer's knowledge of the language; but this generative grammar does not, in itself, prescribe the character or functioning of a perceptual model or a model of speech production. (Chomsky, 1965, 9)

The reason that, as Stanley correctly notes, theories of generative syntax are to some extent insulated from the stops, starts, and muddles of actual speech is that syntactic theory is not a theory of the processes and procedures whereby actual speaker-hearers produce or interpret actual speech in actual situations. Syntax, as a part of generative grammar, is a theory of *competence*, not a theory of *performance*. That is to say the goal of contemporary syntactic theory is to provide an account of the syntactic rules and principles that constitute a speaker-hearer's implicit *grammatical knowledge*. In actual speech this grammatical knowledge interacts "with a variety of factors" to yield actual speech. From the perspective of generative grammar these other factors — which include other knowledge, processes and systems — are "noise" to be factored out. Thus a speaker's muddled actual speech is a result of this complex interaction, only one factor of which is grammatical competence. Thus it is not necessarily a problem for contemporary syntactic theory that it predict that much of what speakers actually say is ungrammatical. Nor, conversely, is it necessarily a problem for contemporary syntactic theory that some utterances speaker-hearers judge to be ungrammatical are predicted to be grammatical.

Truth-conditional semantics, however, is not in this way insulated from what speaker-hearers actually do. This is because the two-step model of interpretation inherent in truth-conditional semantics is a theory of *performance*; it is a theory about the

processes and procedures speaker-hearers actually utilize in interpreting utterances.[18] Consequently, if the truth conditions speaker-hearers actually interpret utterances as having are incompatible with the predictions of the model, then this counts as evidence against the model.[19]

Stanley, on behalf of truth-conditional semantics, might claim that the two-step model of interpretation is not intended to yield predictions about how speakers *actually* interpret utterances; that is, he might claim that the two-step model is a theory of competence, not performance. But this would be simply to abandon *traditional truth-conditional semantics* in favor of *truth-conditional pragmatics* — the conception of semantics proffered by Bach, Sperber and Wilson, Carston, Stainton, Recanati, myself, and others. According to truth-conditional pragmatics, the truth conditions speaker-hearers actually judge utterances to have are *not* determined by LFs and the semantic values of relevant features thereof, and thus truth-conditional pragmatics rejects the principle of utterance compositionality. Rather, according to truth-conditional pragmatics the LF of an utterance and the semantic values of its semantically relevant features are only some of the factors contributing to the interpretation of the utterance. In addition, what have been thought to be mere *pragmatic* processes must also be invoked to yield truth conditions.

6. CONCLUSION: ARE THERE NON-SENTENTIAL ASSERTIONS?

The *pragmatic strategy* is not a live option for the sententialist. The problem, in brief, is that any reason the sententialist might provide in support of the claim that *prima facie* counterexamples involving non-sentential utterances do not express truth conditions will also apply to the many other sorts of *prima facie* counterexample. This widespread applicability has two significant consequences: First, it renders all utilizations of the syntactic strategy otiose. If speaker-hearers habitually treat utterances that do not really express truth conditions as if they did express truth conditions and thus are habitually grossly mistaken concerning the truth conditions of utterances, then there is no need to posit hidden syntactic structure to account for their interpretative judgements. If speaker-hearers are habitually *mistaken* about truth conditions, there is no need to posit hidden syntactic material to render them *correct*. Second, and more importantly, the result that speaker-hearers are habitually grossly mistaken concerning the truth conditions of utterances is incompatible with the two-step model of interpretation inherent in truth-conditional semantics. This model is a theory of *performance*; it alleges to describe, albeit in very general terms, the process whereby speaker-hearers actually determine the truth conditions of utterances. If this model predicts that speaker-hearers are often, perhaps usually, mistaken in their interpretations, then the model, and truth-conditional semantics generally, must be rejected.

It does not follow from the failure of the *pragmatic strategy* that sententialism is false, for perhaps the *syntactic strategy*, in various forms, can be utilized to explain away apparent non-sentential assertions and all the other sorts of *prima facie* counterexamples to utterance compositionality. But I am skeptical. Suppose, as is called for by the syntactic strategy, that the LFs somehow instantiated in the brains of *speakers* are much richer than what is phonologically represented in, or even suggested by, their actual speech. How is this additional phonetically unrealized material discerned by *hearers*? I believe

that the problem posed by this simple question is in many cases insurmountable, and consequently that many non-sentential utterances constitute genuine counterexamples to utterance compositionality. If this is right, then truth-conditional semantics should be rejected in favor of truth-conditional pragmatics. To endorse truth-conditional pragmatics, however, is to take only a small step toward explaining how speaker-hearers actually manage to interpret one another's utterances.

NOTES

[1] Theorists have posited a number of different processes of ellipsis. The standard sort of syntactic ellipsis (e.g., Morgan, 1973) involves a process of "deletion" that results in structure that is not phonetically realized. Another proposal (Williams, 1977) involves a process of "copying" structure already present in the discourse of the elliptical utterance. Both sorts of process require identical syntactic structures to be already present in the discourse environment because only structure that is already phonetically realized in a discourse is allowed to be "copied to," or "deleted from," a subsequent utterance. This condition of identity precludes the standard processes from applying in cases where an apparent non-sentential assertion appears in discourse initial position. To account for these more problematic sorts of cases theorists proposed *pragmatic ellipsis*. (See Sag and Hankamer 1977, and Stanley, 2000). Pragmatic ellipsis does not require a *phonetically realized* linguistic antecedent — rather the material to be "copied" or "deleted" is provided by other features of the discourse context. The problem with *pragmatic ellipsis* is that it violates the condition of recoverability: other features of context often drastically underdetermine what the ellided information might be. (See Barton 1990 and Clapp 2001 for criticisms of *pragmatic ellipsis*.)

[2] Though Stanley is correct that it is open to the sententialist to utilize both strategies, considerations of parsimony suggest that a unified account of all instances of non-sentential utterances is, *ceteris paribus*, to be preferred. That is, a fully adequate unified account is to be preferred over a fully adequate disjoint account.

[3] In a previous paper (Clapp, 2001) I argued against the syntactic strategy, or more precisely against pragmatic ellipses. In this paper criticizing the other general strategy of response available to the sententialist, hence this paper complements the earlier paper.

[4] I borrow the phrases 'truth-conditional semantics' and 'truth-conditional pragmatics' from Stanley (2000). The phrases are defined in section II.

[5] Another strategy of response would be to maintain that at least one of the lexical items involved in my utterances is, despite appearances, context-sensitive, and thus is assigned different semantic values in different contexts. To my knowledge, no defender of truth-conditional compositionality has utilized this 'indexical strategy" in response to the *prima facie* counterexamples involving non-sentential utterances. The indexical strategy has, however, been used in various forms as a response to the *prima facie* counterexamples involving propositional attitude ascriptions (Richard, 1990) and quantifier domain restriction (Szabo and Stanley, 2000).

[6] As Davidson (1968) notes, Frege (1893) himself violated this constraint in order to account for attitude ascriptions. A very similar violation of semantic innocence is proposed in Higginbotham (1991). It is controversial as to whether or not the violations of semantic innocence proposed by Frege and Higginbotham undermine the explanatory power of a truth-conditional semantic theory. But even if these violations are allowed it is, I believe, agreed among truth-conditional semanticists that, e.g., an adjectival phrase such as 'totally useless' cannot in some syntactic environments express truth conditions, and yet in other syntactic environments merely serve to contribute toward the truth conditions of complete sentences in which it appears.

[7] I appropriate this usage of 'semantics' from Stanley (2000), though I think that much of what Stanley would classify as *semantic* would usually be classified as *pragmatic*. For example, the question of how indexical expressions are interpreted as having different semantic values in different contexts would typically be classified as a question for *pragmatics*. But on the usage I have appropriated from Stanley, this is, at least in part, a question for *semantics*. I think nothing important depends upon such terminological issues.

[8] If the defender of truth-conditional compositionality attempts to explain away a *prima facie* counterexample by positing additional, phonologically unrealized, *non context-sensitive* material in the LF of the

expression, where the identity of this material varies across contexts, then he is claiming that the expression does not actually possess property (c). And if he attempts to explain away a *prima facie* counterexample by positing additional, phonologically unrealized, *context sensitive* material in the LF of the expression, where the identity of this material does not vary across contexts, then he is claiming that the expression does not actually possess property (d). It would also of course be possible to utilize a mixed strategy according to which both (c) and (d) are not possessed by a *prima facie* counterexample.

9. Bach thus maintains that occurrences of sentences such as (2) express both a minimal and an expanded proposition, though the expanded proposition is the one the speaker intends to communicate. Bach's views thus differ from Recanati (1984), who denies that a minimal proposition is expressed. This issue is beside the point of my paper. All that matters for my purposes is that because occurrences of (2) can express the expanded proposition, expressions such as (2) constitute *prima facie* counterexamples to utterance compositionality.

10. There are some difficult issues concerning presupposition that I am glossing over. An utterance of a sentence containing a definite description whose presupposition is not satisfied is neither true, nor false. (Or so I shall assume.) But does such an utterance nonetheless express truth conditions? I here assume that it does not, though I am not confident of this. Thankfully, the issue is not directly relevant to my concerns.

11. As Stainton reminded me, many speech acts that seem not to be *assertions* nonetheless express truth conditions. For example, in finishing a joke one might utter, 'Pigs do fly!' Such an utterance expresses truth conditions — the utterance is clearly false — but the speaker is telling a joke and is not making an *assertion*. Hence, issues concerning the *illocutionary force* of non-sentential utterances are for the most part irrelevant to utterance compositionality, and thus also irrelevant to Stanley's thesis that "all truth-conditional effects of extra-linguistic context can be traced to logical form" (2000, 391).

12. Stanley is correct that the thirsty man is neither asserting, nor requesting, nor commanding, but I think the man is performing a genuine sort of speech act with an appropriate sort of force. The thirsty man is *ordering*, where ordering is a sort of speech act that, in appropriate conditions, customers make to various sorts of employees in the service industry.

13. Note that this could not be a complete specification of the "determinate content," as it contains an "incomplete" definite description itself in need of further restriction.

14. I think it is conceivable that there be an apparent assertion in need of either *completion* or *expansion* that nonetheless has "determinate content," but I cannot provide an example of such an utterance. One reason for this is that it is not at all clear to me what it would be for an utterance to have "determinate content."

15. Stanley (2000) presents purely *syntactic* arguments to support utilization of the syntactic strategy for *prima facie* counterexamples involving quantifiers and relational expressions. Stanley argues that certain phenomena involving binding relations require the sort of "hidden indexicals" posited by the syntactic strategy. Consequently, if Stanley's syntactic analyses are correct, there are *syntactic* reasons to utilize the syntactic strategy, despite the widespread applicability of the pragmatic strategy. So, if Stanley's syntactic arguments are sound, then some *prima facie* counterexamples can be explained away by *both* the syntactic strategy and the pragmatic strategy. But this presents the defender of truth conditional semantics with a dilemma, for the pragmatic strategy and the syntactic strategy "pull in different directions" — the former is an "error theory" that rejects the judgments of ordinary speaker-hearers, while the latter renders these judgments compatible with utterance compositionality. Hence something must give; one cannot allow both strategies to apply to a particular *prima facie* counterexample.

16. The most significant problem is of course that there is nothing in most contexts of utterance that might fix the semantic values of the alleged hidden indexicals. In terms of Stanley and Gendler Szabo's hidden indexical account of quantified sentences, there is no mechanism that might fix the semantic value of the posited function indexicals $f()$, and i.

17. Unless of course there is independent, perhaps syntactic, evidence for the existence of the posited hidden machinery, in which case the defender of truth conditional semantics would be faced with the dilemma described in note 15.

18. This is not of course to say that the two-step model is offered by Stanley and Gendler-Szabo as a *complete* performance theory. They are quite clear that they are offering only a sketch of such a theory, but it is nonetheless a sketch of a *performance* theory, intended to describe the processes speakers actually utilize in interpreting actual utterances. In Stanley and Gendler Szabo (2000), they introduce the two-step model of interpretation with the following, "In order to interpret typical assertions of others, we normally need to know what sentence they used. . . . " (228). And in Stanley (2000) he writes, while describing the model, "It is often assumed that the objects of semantic interpretation, that is, syntactic logical forms, are free of lexical and structural ambiguity. However, sometimes the sounds we hear suffer from such ambiguity. One role context plays is in helping us to decide which logical form is the one that has been uttered" (399).

19. Of course, even a theory of linguistic performance must allow for wholly non-linguistic factors — perceived fire alarms, bricks to the head, etc. — to affect actual speech production and interpretation. Even a *performance* theory is a theory in the special sciences, and thus is rife with *ceteris paribus* clauses.

REFERENCES

Bach, K., "Conversational Impliciture." *Mind and Language* 9 (1994): 124–62.

Bach, K., "Do Belief Reports Report Beliefs?" *Pacific Philosophical Quarterly* 78 (1997): 215–41.

Barton, E., *Nonsentential Constituents: A Theory of Grammatical Structure and Pragmatic Interpretation.* Amsterdam: John Benjamins. 1990.

Carston, R., "Implicature, Explicature, and Truth-Theoretic Semantics." In S. David (ed.), *Pragmatics.* Oxford: Oxford University Press. 1991: 33–51.

Chomsky, N. *Aspects of the Theory of Syntax.* Cambridge: MIT Press. 1965.

Clapp, L., "What Unarticulated Constituents Could Not Be." In J. Campbell, M. O'Rourke and D. Shier (eds.) *Meaning and Truth.* New York: Seven Bridges Press. 2002: 231–256.

Davidson, D., "On Saying That." *Synthese* 19 (1968): 158–74.

Frege, G., "Uber Sinn und Bedeutung." *Zeitshrift fur Philosophie und Philosophische Kritik* 100 (1893): 5–50.

Frege, G., "Der Gedanke", 1918. Translated as "Thoughts" in Geach, 1977.

Geach, P. (ed.), *Logical Investigations.* tr. P. Geach and R. Stoothoff. New Haven, Conn.: Yale University Press. 1977.

Higginbotham, J., "Belief and Logical Form." *Mind and Language* 4 (1991): 344–69.

Larson, R. and Segal, G., *Knowledge of Meaning.* Cambridge: MIT Press. 1995.

Morgan, J., "Sentence Fragments and the Notion of 'Sentence'." In *Issues in Linguistics*, B. Kachru et al. (eds.) Urbana: University of Illinois Press. 1973: 719–751.

Recanati, F., "Domains of Discourse." *Linguistics and Philosophy* 19 (1996): 445–75.

Richard, M., *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them.* Cambridge: Cambridge University Press. 1990.

Sag, I., and Hankamer, J., "Syntactically vs. Pragmatically Controlled Anaphora." In *Studies in Language Variation*, R. Gasold and R. Shuy (eds.) Washington D.C.: Georgetown University Press. 1977.

Schiffer, S., "Belief Ascription." *The Journal of Philosophy* 89 (1992): 499–521.

Sperber, D., and Wilson, D., *Relevance.* Cambridge: Harvard University Press. 1986.

Stainton, R., "Using Non-Sentences: An Application of Relevance Theory." *Pragmatics and Cognition* 2 (1994): 269–84.

Stainton, R., "Non-Sentential Assertions and Semantic Ellipsis." *Linguistics and Philosophy* 18 (1995): 281–96.

Stanley, J., "Context and Logical Form." *Linguistics and Philosophy* 23(4) (2000): 391–434.

Stanley, J., and Z. Szabo, "Quantifier Domain Restriction." *Mind and Language* 15(2&3) (2000): 219–261.

Travis, C., "On What is Strictly Speaking True." *Canadian Journal of Philosophy* 15(2) (1985): 187–229.

Williams, E., "Discourse and Logical Form." *Linguistic Inquiry* 8 (1977): 101–139.