

Thurs 30 Mar 06

Empty Representations in Linguistic Perception
Draft of Material for Talk for Early Content Conference, April 06

Abstract

I argue that, *pace* Chomsky (2000, 2003), standard theories of linguistic competence are committed to taking talk of representations seriously, in particular, to recognizing that the “of x” clause that invariably follows “representation” is a way of specifying that representation’s intentional content. One reason to insist upon intentional content in such cases is that the “x” in “of x” may not exist (as in “of Zeus”). This issue is especially relevant to linguistics since, recapitulating considerations raised by many linguists, I go on to argue that most of the SLEs themselves seldom, if ever exist: it is doubtful there are many, if any, tokens of them in space and time; indeed, their existence is by and large inessential to the needs of either communication or serious linguistic theory. All that linguistic theory requires to be real in this regard are the representations, presumably entokened in people’s brains, understood, however, in terms of their intentional contents.

Introduction

In my talk today I want to summarize work I’ve been pursuing for several years on the role of representations in linguistic theory. Although I am concerned with the *intentionality* of these representations –what they are representations *of*– I am *not* concerned with this topic as it arises for the meaning of *expressions in natural languages themselves*, what has been called a “*linguo*-semantics. Rather, I’m concerned only with a *psycho*-semantics, or the meaning of the representations that, according to most linguistic theories, *the mind* employs for the syntactic and phonological structure of expressions in natural language, i.e., for such things as NPs, VPs, morphemes, phonemes, syntactic and phonological features (I refer to all such “things” as “Standard Linguistic Entities,” or “SLE”s).

In earlier papers (Rey 2003a and b), I have argued that standard theories of linguistic competence are committed to taking such talk of representations seriously, in particular, to recognizing that the “of x” clause that invariably follows “representation” is a way of specifying that representation’s intentional content: a representation of the word “dog” is a representation with {“dog”} as its content.¹ One reason to insist upon intentional content in such cases is that the “x” in “of x” may not exist: one can have a representation of Zeus even though there is no Zeus; so we need to be able to talk instead about the “idea” of Zeus –i.e. the content {Zeus}. Indeed, a provisional characterization of intentional content

1. I shall refer to intentional contents by enclosing words that express them in curly-brackets (I usually use square ones, but in linguistics –and so in some places in this paper– expressions in square brackets refer to phones or phonetic features).

might be: *however x is to be understood in the expression "representation of x, but there is no x."*² This issue is especially relevant to linguistics since, recapitulating considerations raised by many linguists, I go on to argue (Rey 2005, 2006) that most of the SLEs themselves seldom, if ever exist: it is doubtful there are many, if any, tokens of them in space and time –certainly not as many as we ordinarily think. More importantly, even if some SLEs do turn out to exist, their existence is by and large inessential to the needs of either communication or serious linguistic theory. All that linguistic theory requires to be real in this regard are the representations, presumably entokened in people's brains, understood, however, in terms of their intentional contents.

There are two issues motivating my discussion. One is to call attention to the significant role of intentionality at the "meta-" level of linguistic theory itself. This, I think, is important for a number of reasons having to do with a theory of mind, not least of which is Chomsky's recent (2000) surprisingly Quine-like denials of the role of intentionality in any serious science. But another reason, perhaps closer to Chomsky's heart, is my dissatisfaction with the "externalist" semantics of philosophers like Dretske, Millikan, Neander, Fodor, Taylor and Tye, that has become virtually the received view in philosophy since the work of Kripke, Putnam, Evans and Burge in the 1970s. Indeed, it is, I think, this pervasive externalism that leads "internalist" linguists like Chomsky and Jackendoff to distance linguistics from intentionality entirely: given what they reasonably see as the inadequate notion of "reference" that underlies so much externalism, they wrongly conclude that the notion of intentionality has no place in their theories at all. What they and their opponents here agree about is that intentionality is committed to some strong externalist thesis, and that there is little place for the kind of more *internalist intentionality* that was once a commonplace in theories about the mind, for example, in the work of Franz Brentano, who regarded reference to the non-existent as one of the marks of the mental.

For the record, I should emphasize that I shall not be defending a *purely* internalist view of intentionality. I think the familiar examples of Kripke and Putnam do support what I call a "Weak" (what Francois Recanati once called a "holistic") Externalism: *some* intentional ascription depends in *some way or other* upon *some* facts outside the skin of a thinker. What I call "Strong" Externalisms try to cash out these mere existentials and insist on some other specific causal relation between every primitive of thought and some actual external phenomena that provides that thought with its content. It's these Strong Externalisms that I join Chomsky in rejecting. Unlike Chomsky and these other philosophers, however, I see no reason for a theory of intentionality to be burdened with them. Indeed, the main point of my discussion is to make this apparent at some of the basic levels of psychological processing. (And since this is my main point, I shall not be concerned with establishing the non-existence of every SLE, but merely to show that it's plausible about enough of them to give the Strong Externalist serious pause.)

In §1 of my talk I'll briefly summarize the reasons for thinking SLEs do not exist, essentially summarizing the considerations linguists since Sapir (1933/63) have marshalled for claiming that they have only a "psychological reality." This latter phrase has I think been responsible for a number of confusions that I will explore in §2, not least Chomsky's and other's surprising denial of the role of

2. I say "provisional" since I do not want to present here (what in any case is not yet available) a *theory* of intentional content, and want to remain as close to the usage of "representation" linguists seem to have in mind, as, for example, when Chomsky writes (in passage I'll discuss below):

The representations are postulated entities, to be understood in the manner of the mental image of a rotating cube, whether it be the result of tachistoscopic presentations of a real rotating cube or stimulation of the retina in some other way; or imagined, for that matter. --(Chomsky 2000:158??)

intentionality in serious linguistic and psychological theory. I'll argue that, while Chomsky is perfectly right to insist on the "psychological reality" of various *representations* of natural language in the brain, this does not entitle him to insist on a similar –or any!– reality of the "things," the SLEs, those representations represent. Talk of SLEs is better understood as talk about the *intentional content* of those representations. More specifically, owing to a notorious oddity in the way in which we talk about intentional content, talk of SLEs is talk about the *intentional inexistents* projected from that intentional content: i.e. it is a way of characterizing intentional content by talking about the things the content is about, even when those things do not exist –neither in the world, the brain, or anywhere else. In the concluding §3 I'll briefly explain how this (what I call "*folieist*") conception of language is not quite as mad as it might seem, how it doesn't entail a general idealism often inferred from it, and how it is in fact precisely what a naturalized rationalist like Chomsky should embrace.

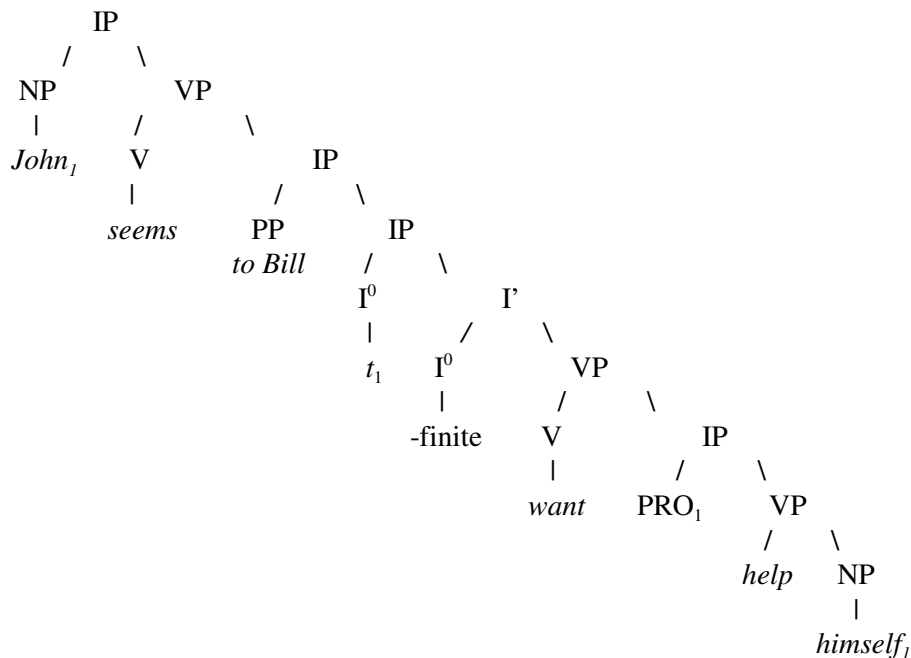
1. The Non-existence of SLEs

1.1 Syntax and Morphology

It would seem to be a commonplace that tokens of SLEs are regularly produced and heard by language speakers (see, e.g., Devitt and Sterelney 1987:59, Bromberger and Halle 20??). Over the years, however, this has been repeatedly challenged by linguists (e.g., Saussure (1916/66), Sapir (1933/63), Chomsky and Halle (1968), Jackendoff (1987)). Some of the most interesting cases concern phonology, but there is not time to consider their complexities here. The linguists' general point can be appreciated by considering merely the typical structure of a sentence, such as one indicated by:

(U) John seems to Bill to want to help himself.

In standard generative grammars this sentence bears *at least* something like the following analysis:



Thus, not only is there an elaborate tree structure in (U), there are also “empty” categories: trace (t_1) and PRO that indicate a node in the tree that for one reason or another has no SLE attached to it, for example,

because an SLE has been “moved” to a different node in some “transformation” of the sentence, and/or needs to be co-indexed with elements at other nodes: thus t_1 is co-indexed with John, PRO₁, and with *himself*, capturing the fact that the subject of *want* and *help* are both *John*, and not *Bill*.

Now *what actual thing in the world possesses this structure?* Well, what is the thing I take myself to have just uttered? A token of an English sentence. This, according to the commonplace, is something in space/time. But does anything I actually produced in *space and time* have the above structure? I think not. Compare (U) to another human artifact that admits of a type/token relation, my trusty Honda automobile. Like (U), it was produced by people intending to produce something with a certain complex structure. However, in the case of my Honda, it’s crucial that the engineers *succeeded* in doing so: lifting up the hood, one finds a very elaborate, causally structured engine that is responsible for my Honda’s characteristic zip. However, examining the spatio-temporal region of my utterance of (U) reveals nothing remotely like the linguistic structure that I intended. One can *imagine* building a kind of Rube-Goldberg tree-structure, replete with morphemes suspended from each terminal node, and with wires and pulleys that permitted some movements or connections here and prevented others there. But it is an interesting fact that the noises we produce when we intend to utter sentences are nothing like this. There is nothing we ordinarily produce that has a tree structure with items that have in fact been “moved” or even in fact “cross-indexed” from one node to another, leaving a nevertheless still existing node “empty.” As we’ll see in the next section, there are not even boundaries in the acoustic stream that could serve to causally distinguish the physical parts of *words*, let alone NPs, VPs and the like. In uttering (U), I simply didn’t succeed in producing something with an actual linguistic structure in the way that the technicians at Honda produced something with an automotive one.

Someone might claim that the structure of an SLE is “more abstract,” in the way that, for example, a computer can have an abstract computational structure not easily visible in the array of transistors and the pieces of metal and plastic of which it’s made. However, there is an important difference between the two cases: in the case of the computer and for at least *some* of the algorithms it can be literally be said to be running (e.g. the operating system), *there are all sorts of causal claims and true dispositional counterfactuals about what would happen if you were to apply a certain current to certain points in the network, and it’s these causal claims and counterfactuals that make it true, if it is true, that the computer is executing a certain algorithm:* indeed, the steps in the algorithm can, if the computer is actually executing it, be clearly mapped in an explanatorily illuminating way to specific physical states of the machine. Nothing analogous seems to be true of an SLE --at least not independently of human perceivers (to whom I’ll return in §2).

This latter contrast seems to me crucial. The structures of (biological) trees, rivers, neural nets and automobile engines play a real causal role in the world; but, I submit, the tree structures of sentences do not. All that need be true for the noises a speaker makes to have their intended effect is that they be *perceived* to have the tree structure that the speaker intended; and something can be *perceived* to have a structure without it actually *possessing* that structure, as in the case of the familiar Kanizsa triangle:




fig. 1

There is no triangle in fig. 1, anymore than there is a "moving arrow" in a row of consecutively flashing lights pointing into the Las Vegas casino, or any sort of "mouse" in a Mickey Mouse cartoon (though

notice, interestingly enough, that we can “refer” to such “things” nonetheless).³ Just so: there are no real *sentences* in the acoustic stream of “speech.”

3. Here I think I take issue with Zenon Pylyshyn's (2003:268ff) treatment of his FINGS, or the “virtual objects” of early visual perception, which he insists on regarding as real objects to which our visual system bears causal relations.

1.2 Phonology and Phonetics

What about the more basic SLEs, such as phonemes and phones? Here, to be sure, the issues are more complex, but not, I think, in a way that changes the general point.

Especially after the pioneering work of Jakobson, Fant and Halle (1952/63) and Chomsky and Halle (1968), attention in phonology has been directed more upon phonological and phonetic "*features*" than on phonemes or phones themselves. How exactly to identify such features is a topic of continuing controversy, there being considerable tension between *gesticulatory* and *acoustic/perceptual* criteria, a tension that, of course, isn't surprising, given the dual demands on natural language of production and comprehension. Suffice it to say that an ecumenical view is that a phonological feature is widely regarded as "abstract" in the sense that it can be connected to any of a variety of articulatory gestures or acoustic patterns.

There are a number of peculiarities of the relation between the acoustic stream and at least the intended *phonemic sequence* that render the identification of even any of these SLEs with physical phenomena at least extremely awkward.

(i) **failure of boundaries:** in normal speech, there is no acoustic difference in such pairs, *attack/a tack*, *an ashla Nash*, *a parent/apparent*, *this guy/the sky* (see Jackendoff 1987:57, 1994:56). As the phoneticist John Laver (1997) points out:

The stream of speech within a single utterance is a continuum. ... [S]egmentation is mostly an imposed analysis, and not the outcome of discovering natural time-boundaries in the speech continuum, is a view that deserves the strongest insistence. --(Laver:p101)

(ii) **failures of segment correspondence** (FBG:296-7). A sequence of the formants for individual phones may not be perceptible as the sequence of those phones: one doesn't produce the monosyllable "bag" by producing the trisyllabic sequence of phones of which it is composed (see Liberman 1996:33).

(iii) **unruly complexity of physical typing:** In a way that is now familiar in discussions of many macro-theoretical entities, phonetic *types* do not correspond to physical acoustic types. For example, the /d/ is a significantly different type of acoustic event depending upon the vowel it immediately precedes: the /d/s in, for example, /di/, /de/, /da/, /do/, /du/ are all different formants (FBG:296). In many contexts, a phone gets "assimilated" to adjoining sounds, /n/, example, often being pronounced as /m/ in "seven pens." More dramatically, mere silence (on the order of 30-40ms. appropriately placed) is sufficient to distinguish the pairs *slit/split*, *say shop/say chop*, and the triple *rapid/rabid/ratted* (Liberman 1996:325,346). Perhaps mere silence should be counted as a token of "p", "ch" and "t"? More plausibly, Liberman and Mattingly (1985) write:

There is simply no way to define a phonetic category in purely acoustic terms. A complete list of the cues --surely a cumbersome matter at best-- is not feasible, for it would necessarily include all the acoustic effects of phonetically distinctive articulations. --(Liberman and Mattingly 1985:12)

(iv) **displacement of cues:** sounds are heard as different phonemes at one point in the speech stream as a result not of differences at that point, but at earlier or later points in the stream,

sometimes from one end of a complex syllable to the other (Lieberman 1996:p36). Jackendoff (1987) discusses the example of *pat* and *pap*, where

the end effects of the vowel provide the only acoustic cue for the correct choice of consonant. The acoustic signal may be roughly segmented into temporal intervals that contain cues for *p + a* and *a + t* but none that consist of all the cues for *p* alone or *a* alone or *t* alone. –(Jackendoff 1987:57-8)

Similarly, in normal speech, the difference between "rider" and "writer" is heard as a difference between /d/ and /t/. But in fact the /d/ and /t/ are pronounced identically (as flaps), the difference in how they're heard being due to the longer sound of the preceding vowel (see Fodor, Bever and Garrett 1972:292-3).

Why should the phonological system be so complex? Liberman points out that a system allowing for continuous assimilations, co-articulations and miscellaneous, displaced context-dependent cues nicely solves the problem of a communicating long strings of discrete, categorical units *quickly* and in a fashion that can be manageably decoded by a hearer:

The problem is that, at normal rates, speech produces from eight to ten segments per second, and, for short stretches, at least double that number. But if each of those were a unit sound, then rates that high would strain the temporal resolving power of the ear, and...also exceed its ability to perceive the order in which the segments had been paid down. ... –(Lieberman 1996:pp33)

The phonological order a hearer recovers from a stretch of sound is thus determined by the overall "shape of the acoustic signal, not by the way the pieces of sound are sequenced in it" (p34).

1.3 Philosophical Moves

Philosophers keen on saving ordinary SLE ontology (e.g. Fodor (1998), Devitt (forthcoming)) will be quick to point out that many phenomena in the world are relationally identified: something is poisonous iff it causes death upon ingested; something is a carburetor iff it is the part of a motor that mixes the oxygen and the fuel. This is, of course, the strategy behind Strong Externalist strategies that would identify the content of an SLE representation with the SLEs to which that representation is or would be causally related, under normal, ideal or perhaps evolutionarily selective conditions. For example, on Fodor's (1987, 1990) "asymmetric dependency" theory, a symbol "S" would mean [fricative] iff its being caused by non-fricatives depends upon a law that it's caused by fricatives; and, by his (1998) proposal, something's being a fricative is constituted by the fact that it is the sort of phenomena to which hearers generalize on being presented with prototypical fricatives..

In the longer version of the present discussion, I explore Fodor's and a number of other such efforts that have been or could be made to save SLEs, and a number of difficulties that have been encountered in trying to work them out. I'll raise two problems here: variability (§1.3a) and circularity (§1.3c), both of which are made vivid by considering one of the most plausible candidates for an actual SLE, *the articulatory feature* which I'll discuss in between the two (§1.3b). The problems should, I think, lead one to wonder why anyone should bother with such efforts (§1.3d).

1.3 Variability

The problem of variability is the obvious variability of people's perceptions and executions of their intentions within different speech contexts, a variability that renders actual acoustic or motoric patterns unfit for general linguistic theory. In addition to dialectical and regional differences, there are differences merely in pronunciation between people due to, e.g., speed, age, gender, anatomy, speech impediments, personality, social class, and, even within a single person at certain stage of life, differences due to social style, circumstances, emotional intensity and relative inebriation (think of whispering, yelling, singing, auctioneering). And, of course, people's abilities to "correct" for all this variability itself varies widely, with the consequence that one would be very hard put indeed to specify a "normal" listener, delivery and social context that would suffice for all speakers of "the same language" or even dialect.

To be sure, one shouldn't infer underlying from mere superficial chaos: as Chomskyans are the first to insist, there are *systems* of syntax, semantics and phonology constituting a linguistic competence that underlies the diversity of dialects and pronunciations. But the difficult question is to say precisely what that system is, at what level it is to be described, and, in particular, *whether there's any reason to suppose that it is tied to type identities among external stimuli*. Obviously, one would want to abstract from inebriation, and maybe emotional intensity. But what of the other "interferences"? Should linguistics be grounded by BBC newscasters, *l'Academic Française* or how sufficiently well educated lawyers entone in court? Phoneticians do not await such analyses any more than vision theory waits upon a dispositional analysis of worldly color, adjudicating the blue/green differences over which men and women regularly contend (see Hardin 19??). To press Jespersen's famous quip, a phonology defined by relations to auditory stimuli would seem at best an idiolect with a gun-boat --and an obsessive, *a priori* philosopher at the helm!

1.3b Phonological Features

Both the above difficulty and the circularity problem to be raised next can be nicely illustrated by considering a plausible candidate for one particular sort of real, physical SLE that has emerged in the light of the above phenomena, *an abstract phonological gesture*. This is a proposal that has its roots in the articulatory feature proposals of Halle, Jakobson and Fant (1951) and Chomsky and Halle (1968). In the form it takes in Liberman and Mattingly's (19??/96, hereafter "L&M") recent "Revised Motor Theory":

A phonetic gesture...is a class of movements by one or more articulators that results in a particular, linguistically significant deformation, over time, of the vocal tract configuration. The linguistic function of the gesture is clear enough: phonetic contrasts... -(L&M:p252)

They hasten to add:

What is not so clear is how the gesture relates to the actual physical movements of the articulators and to the resulting vocal-tract configurations... We would argue...that the gestures do have characteristic invariant properties, as the motor theory requires, though these must be seen, not as peripheral movements, but as the more remote structures that control the movements. These structures correspond to the speaker's intentions. -(L&M:253-4)

They go on to suggest that there is a special linguistic module that represents in some detail how these intentions eventuate in motor movements. Along lines inspired by Fodor (1983), they propose an informationally encapsulated system that engages in automatic inference-like computations from the various cues presented at various points in some perceptual unit.⁴ Such a system

incorporates complete information about the anatomical and physiological characteristics of the vocal tract and also about the articulatory and acoustic consequences of linguistically significant gestures...[T]he module has then merely to determine which (if any) of the small number of gestures that might have been initiated at a particular instant could, in combination with gestures already in progress, account for the signal. --(L&M:256).

This "analysis by synthesis" strategy (originally proposed by Stevens and Halle 1967) permits the complexity and variability of the mapping observed between the abstract articulatory gesture and actual motor movements and acoustic signals. The abstract phonetic gesture that the hearer's module infers in this way from the acoustic effects of those movements corresponds to no specific movements in this or that person's vocal tract, but to "some very abstract conception of vocal tracts in general" (Liberman, Borman and Rafael, 1977:348). On this view, a phonetic feature is a little like an ideal Euclidean triangle, an entity that might be instantiated in splendid isolation under some sort of ideal circumstances (as in Eimas 1971), but one seldom if ever encountered in normal continuous speech, which exploits all the context dependent tricks we've noted.

1.3c Circularity

Even if we could decide on what to count as the right speech produced and heard in the right circumstances, we would still face the problem of avoiding circularity in specifying the disposition to such productions and perceptions. In particular, we would need to delineate the SLEs without relying on the very characterization of them a dispositional or aetiological theory is trying to supply. Leading theories in phonology seem driven to characterize SLEs either in terms of a hearer's "categorical" *perception of a noise as an SLE*, or, as we just saw in the case of the revised motor theory, in terms of a speaker's *intention to produce one* (see also Bromberger and Halle 2000:pp24-5). But how are these perceptions or intentions to be understood on a relational analysis? As an intention to produce a sound that will be produced by just such an intention? To avoid vicious circularity, a dispositional theory would seem to need a characterization of the content of SLE representations (perceptions, intentions) that is independent of the objects dispositionally specified. As I argue in my earlier papers, this independence would not appear to be satisfiable by currently fashionable Strong Externalist theories of content, which

4. Note that it is precisely such a system that makes intelligible how later cues within a unit of perception could determine assignments for earlier portions of the unit, as we observed occurs in the case of displacement phenomena. Indeed, It is worth comparing the process to a related phenomenon in vision whereby, if two differently illuminated dots are flashed on a screen in succession, the first appears to "move" to the position of the second, changing to the color of the second before it reaches it (see Dennett and Kinsbourne 1977 for discussion).

identify the content of a (primitive) representation with the external phenomena to which it bears some specific causal relation (as in Fodor 1990, Millikan 1998, 2004).⁵

1.3d Why Bother?

Perhaps here is some ingenious combination of dispositional and aetiological approaches that could circumvent these difficulties and yield relational analyses of some SLEs. But, given the difficulties, one might wonder why it's so important to provide one. Chomsky himself dismisses such efforts as "just spinning wheels" (Chomsky 2000:129). I think what he means is that, if one looks at the theoretical work SLEs are supposed to perform in a theory, any sort of definition in terms of acoustic phenomena –indeed, their very existence– is entirely *needless*. Outside of the acoustic phonetics of the sort that we have quoted, linguists almost never discuss formants. Rather, they discuss how SLEs enter into complex sequences: phonetic and phonological features into phonemes, phonemes into morphemes, morphemes into words and syllables, words into phrases, sentences, and sentences into discourse. Along the way, the sequences may actually be structured trees, involving additional "null" elements, posited to round out certain structures, record "movement" or co-indexing, and so forth. There are abundant theories and disputes about the specific structures, rules, principles and parameters for these structures –i.e., the substance of linguistics, whose details do not wait upon relational analyses of the sort philosophers might expect.

2. Representations and Intentional Content

Chomsky's dismissal of efforts to identify SLEs as phenomena in the external environment leads him, however, to make some further extremely puzzling claims about the nature of the "representations" to which his and virtually all linguistic theories constantly appeal. Defending what he thinks is the proper "internalist" theory of language, he writes:

The internalist study of language also speaks of "representations" of various kinds, including phonetic and semantic representations at the "interface" with other systems. But here too we need not ponder what is represented, seeking some objective construction from sounds to things. The representations are postulated entities, to be understood in the manner of the mental image of a rotating cube, whether it be the result of tachistoscopic presentations of a real rotating cube or stimulation of the retina in some other way; or imagined, for that matter. Accessed by performance systems, the internal representations of language enter into interpretation, thought, and action, but there is no reason to seek any other relation to the world, as might be suggested by a well-known philosophical tradition and inappropriate analogies from informal usage.
--(2000:158??)

5. What I call "Strong Externalism" is the insistence that content is constituted by a real phenomenon to which primitive internal representations bear some causal relation. "Weak Externalism" is the much less problematic view that the content of *some* representations depends in part upon *some or other* relations those representations bear to the external world. I think it's important to notice that the standard examples marshalled by Kripke, Putnam and Burge invite Externalism only of this weaker sort. Strong Externalist theories are, I think, rash conjectures about how to explain such examples.

What then is an SLE?

An expression E of [language] L is a pair <PHON, SEM>, where *PHON(E)* is the information relevant to the sound of E and SEM(E) to its meaning. PHON and SEM are constructed by computational operations on lexical items. ... PHON(E) and SEM(E) are elements at the "phonetic level" and "semantic levels" respectively; *they are phonetic and semantic "representations."* The terms have their technical sense; there is nothing "represented" in the sense of representative theories of ideas, for example. --(2000:173)

Indeed, he spurns the "representative theory of ideas" not only for his internalist linguistics, but for science generally:

If "cognitive science" is taken to be concerned with intentional attribution, it may turn out to be an interesting pursuit (as literature is), but is not likely to provide explanatory theory or to be integrated into the natural sciences. --(Chomsky 2000:22-3)

One has, of course, heard much of this before. Skinner and Quine were fond of inveighing about the "emptiness of a science of intention" (Quine 1960:221). But this is strange company for Chomsky to keep.⁶ Does the rejection of the external reality of SLEs really drive us to such claims? Just what is this alternative conception of "representation" that he is proposing?

I think there are quite a number of issues driving Chomsky at this point, from scepticism about intentionality that he inherited from his teacher, Nelson Goodman, to rejection of various semantic programs in contemporary linguistics, to exasperation with much recent philosophy. I want to focus here on what seem to me certain crucial confusions, between use and mention (§2.2) and between different senses of "psychological reality" (§2.3). I then want to show how a traditional understanding of intentionality, as allowing intentional inexistents, will provide an illuminating way to think about SLEs and Chomsky's program in linguistics generally (§2.4).

2.2a The Use/Mention Confusions Themselves

I shall understand the use/mention distinction as the distinction between deploying an expression to express a certain meaning, and sometimes to refer to certain phenomena, and deploying it, usually within quotation marks, to refer to the expression itself. Thus, I use "fish" when I say that fish swim, and I mention the word when I say "The word 'fish' is monosyllabic." To confuse the two is to confuse the meaning or reference of (the use of) an expression with the expression itself, e.g., confusing fish with the word "fish"; or, to take an example without a real world referent closer to our present concerns, confusing *the meaning* of the word "unicorn" with the word "unicorn" itself. *Meanings*, whatever they are, aren't *words*, any more than fish are. Meanings are something like properties that words have, or

6. I might mention that, among the many audiences to which I've presented some of this material over the last five years, about 3/4 of them –some of them linguists who take themselves to be Chomskyan– are astonished when I quote these passages to them, even though their eyes have passed over them in print, so incongruous do they seem to the understanding of Chomsky with which they were raised (this was indeed my own reaction when I first encountered them).

perhaps abstractions to which they are related. After all, different words can have the same meaning; and "ambiguous" words have several meanings.

When Chomsky writes in the second passage quoted above that "PHON(E) is the information relevant to the sound of E and SEM(E) to its meaning," and then also that "PHON(E) and SEM(E) are phonetic and semantic "representations," this would seem on the face of it a use/mention confusion of this sort. Whatever "information" may be it is not to be identified with the representations that convey it. Thus, while it may be true that, per our earlier discussion, "PHON(E)" may not refer to real phenomena in the external world, this doesn't imply that, even in its technical sense, it mightn't "represent" that information in precisely the way that at least certain standard "representative theories of ideas" might claim. To take up a suggestion Chomsky (2003) endorses elsewhere, PHON(E) might represent "instructions to the articulatory system" (I'll return shortly to the confusing diversity of direct objects the verb "represent" can take).

This sort of use/mention confusion crops up with surprising frequency in his work. For the record I cite a number of other, what seem to me equally significant passages (I italicize the confusing talk):

The language involves three kinds of elements:

- (i) the properties of sound and meaning, called "features";
- (ii) the items that are assembled from these properties, called "lexical items";
- and (iii) the complex expressions constructed from these "atomic" units.

It follows that *the computational system that generates expressions has two basic operations: one assembles features into lexical items, the second forms larger syntactic objects out of those already constructed....* --(2000:10-14)

The computational procedure maps an array of lexical choices into a pair of symbolic objects, phonetic form and LF....*The elements of these symbolic objects can be called "phonetic" and "semantic" features, but we should bear in mind that all of this is pure syntax and completely internalist.* --(2000:125)

The theory of I-language attributes such *features* as [+voice] to the mind/brain. The I-language generates *expressions E constructed from these elementary features*; E is accessed by [the Sensori-Motor system] and used to carry out articulatory gestures... --(2003:@16)

Indeed, in his great, early work, *The Logical Structure of Linguistic Theory* ("LSLT"), he makes an explicit policy of the confusion: within a page of setting out the primes of P_n as phonetic *symbols* that will be mapped by Φ^n to "physical descriptions of phones," he stipulates the convention:

We will henceforth apply the term 'phones' to symbols of **P_n**, as well as to utterance tokens represented by them" --(LSLT:159)

This actually just makes explicit a relaxing of use/mention that was in fact in force already (see LSLT:106). I submit that, with a few exceptions to be noted shortly, this relaxation convention has persisted down to the present day. It's probably a good idea on the whole (keeping track of use/mention with quotations marks would drive linguists blind!). It only becomes problematic when one tries to make

theoretical hay of it, too conveniently evading the problem of intentionality by denying the difference between a representation and what it represents.

But, on second thought, given by our earlier discussion that SLEs don't exist in the external world, perhaps such a denial might nevertheless be a good idea. Since they're nowhere else, why shouldn't SLEs be identified with their representations in the brain?

2.2b Why SLEs Can't Be the Vehicles of Internal Computation

There seem to me two important reasons to resist the identification of SLEs with cranial representations:

- (1) Phonological explanations regularly advert to phenomena that are at least represented as being realized in the oral cavity, not in the brain; and
- (2) Insofar as phonetic and other linguistic features are the objects of computations, they must be represented within the brain.

On behalf of (1), one need only open any phonological text to observe from the start rich discussion of the geometry and mechanics of speech production, and how it provides a basis for phonological classification. Explanation of allophonic variation, displacement, assimilation –why, for example, /n/ is pronounced as /m/ in “seven pens”, or why the plural /s/ in English is voiced after voiced consonants– would be unintelligible without reference to the mouth and the inertia that constrains rapid speech production. Or consider Kenstowicz's (1994) explanation of why there aren't nasal fricatives --"so much airflow is diverted to the nasal cavity that not enough remains to generate the turbulence required of a fricative" (p16); or how consonantal constructions tend to occur in languages where the resonances of the front and back oral cavity reinforce one another (p180). How would these explanations go if [+fricative] or [+consonantal] were representations in the brain?

To be sure, some phonologists, e.g., Hale and Reis (2000), argue that serious phonological theory should be pursued purely *formally*, abstracting from properties of oral articulation; insisting otherwise amounts to what they call "substance abuse." Although I gather this is something of a controversial view (see Idardi (20??) for discussion), it is important to distinguish it from the view being discussed here, of whether SLEs should be *identified* with structures in the brain. It is one thing to claim that one gains greater explanatory power by ignoring the interpretation of a certain formalism, quite another to *identify* the interpretation *with* the formalism. A physicist may find it illuminating to view her theory formally; but that doesn't commit her to regarding leptons and quarks as, *themselves*, formal syntactic objects!

It is important to notice, though, that even a purely formal understanding of linguistic theory doesn't correspond to Chomsky's own linguistic practice. Thus, when he and Halle (1968) explicitly address the question "What exactly is a phonetic representation?" (in a whole section devoted to the topic!), they answer it thus:

A phonetic representation has the form of a two-dimensional matrix in which rows stand for particular phonetic features; the columns stand for the consecutive segments of the utterance generated; and the entries in the matrix determine the status of each segment with respect to the features. In a full phonetic representation, an entry might represent the degree of intensity with which a given feature is present in a particular segment.
–(1968:5)

No use/mention confusions here: the features are not *identified* with the matrices; they are what the entries *stand for*, for example, a "degree of intensity" that an entry "might represent."

The second reason for resisting the identification of SLEs with their cranial representations has to do with a fundamental fact about the kind of computational explanation standardly invoked in linguistic theory, as in most every other area of cognitive science. When Alan Turing provided his famous characterization of computation in terms of Turing Machines, he quite sensibly didn't make such things as NPs or VPs the primitive symbols on the tapes, or postulate a machine that operated by detecting features like [+voice], [-continuant] —much less [+past tense], [+agent] or [+goal]. And this is in part because of some of the very difficulties that I reviewed in §1, of treating such phenomena as *physical* phenomena, not to mention *local* ones. *Given that such linguistic phenomena are not local, physical features of the acoustic signal, there would be no reason to think that a Turing machine could be built that was sensitive to whether such features were instantiated on the tape.* What Turing did, of course, was to put on his "tapes" *formal symbols* that have some or other *local and physically detectable* features, in the ways that `0's and `1's do, so that it would be perfectly clear that a mere machine could respond to their presence and proceed accordingly. Without the stipulation of local, physical properties, a computational theory of mind exploiting Turing's proposal would risk the standard "homunculus" objection raised against computational theories, needing some sort of mind to "read" the symbols on the tape.

Notice, by the way, how this computational/representational conception provides a perfectly innocuous answer to what might otherwise be a puzzle about what "moves" or gets "concatenated" when there's "movement," say, of a "wh" from one position to another. If SLEs were themselves actual structures that were being "derived" in the brain, then one might well worry, say, about the hypothalamus getting in the way. On a representational view, all such talk merely means that *a representation whose content assigns a feature to one location is revised or replaced by a representation whose content assigns it to a different one*, rather in the way the moves of a possible chess game could be described as "1.P-K4, P-K4, 2. P-Q4, PxP ...", or an address book might indicate that someone has moved by changing the entry for the person's address.

In short: we need to conceive phonological features as phenomena in the mouth to explain their intricate relations, and we need their representations in the brain to be different from them in order explain how those relations are computed. But for these representations to do this work they must *represent* these features at least as their intentional content: it is this content that makes a representation a representation *of* one feature rather than another, of [+fricative] or of [-voiced], whether or not those features actually exist.

Why do Chomsky and many of his followers resist this perfectly natural move? As I said, I think there are quite a number of reasons, but what I think underlies the use/mention confusions and the insistence on SLEs in the brain are some peculiar difficulties underlying the issue of "psychological reality" and the ways we discuss intentional content.

2.2c "Psychological Reality"

In an interesting passage shortly following the above discussion of matrices, Chomsky and Halle (1968) proceed to emphasize the "perceptual reality" of the phenomena the matrices describe:

We do not doubt that the stress contours and other phonetic facts...constitute some sort of *perceptual reality* for those who know the language in question. In fact we are suggesting a principled explanation for this conclusion. A person who knows the language should "hear" the predicted phonetic shapes... We take for granted, then, that *phonetic representations describe a perceptual reality*. --(SPE:25, italics mine)

Again, no use/mention confusions: the "perceptual reality" that the phonetic representations describe is presumably not the two-dimensional matrices themselves: *they* are not what hearers "hear."

Such a "perceptual reality" is, of course, precisely what Edward Sapir had in mind when, along the lines of our earlier discussion of SLEs, he described the phoneme as having only a "psychological reality."⁷ However, it is crucial to distinguish this sense of "psychological reality" from another that has equally occupied Chomsky throughout his career, and that is the sense in which he regards his grammatical theories as describing and explaining the psychological facts underlying human linguistic competence:

What is "psychological reality," as distinct from "truth," in a certain domain? ... I am not convinced that there is any such distinction, and see no reason not to take our theories tentatively to be true at the level of description at which we are working. --(Chomsky 1980a:107)

Indeed, he takes the question of *psychological* reality to be identical to the question of a *realistic* interpretation of his theory generally:

The "realist interpretation" of linguistic theory is assumed throughout, and it is argued that the competence attained by the normal speaker-hearer is represented by a transformational generative grammar, which determines the representation of each sentence... These representations are then employed in the use and understanding of language... --(1975c:45)

This last passage affords a nice way of appreciating the distinction we need in the two senses of "psychologically real," which I will provisionally distinguish by "explanatorily real" vs. "phenomenally real." When Chomsky claims that a speaker's competence is well represented by a grammar, and a speaker employs representations of sentences in the use and understanding of language, I take him to claiming that the grammar and those representations are explanatorily real. This accords well with the computational/representational understanding of the theory I advocated above: the representations of SLEs would be causally efficacious phenomena in the brain, just as representations of bank balances are causally efficacious phenomena in computers (leave aside for now whether grammar itself is internally represented). But the *representations* of a phenomena may, of course, be causally efficacious without the *represented phenomenon itself* being so. A visual representation of a dagger may explain why Macbeth thinks there's a dagger in front of him, but that doesn't mean *there really is* such a dagger, entering anywhere into the causal stream. An "hallucinated dagger" is not a real dagger, neither in the air nor in the hallucinator's brain. *It* is not explanatorily real; only the representation of it is. But I think many

7. Tom Washow tells me it was actually Sapir's use here that was the origin of the phrase.

might describe “the dagger” nevertheless as “psychologically real” in the sense being “part of Macbeth’s visual experience”; it is, that is, “phenomenally real.” I think it is this latter sense that Sapir had in mind in calling the phoneme “psychologically real,” and that Chomsky and Halle had in mind when they claimed that “*phonetic representations describe a perceptual reality.*” And I think it is indeed an important truth that people “hear” SLEs as having been actually produced by themselves and other speakers. But, like hallucinated daggers, this doesn’t entail that they exist either in the acoustic stream or anywhere in their brains. Being “psychologically real” in the phenomenal sense does not entail being psychologically real in the explanatory one. (The same distinction arises for the case of the “mental image of the rotating cube” to which Chomsky alluded in the first passage: although *representations of a cube* may well play a role in visual processing, this is no reason for thinking that the *phenomenal mental image* does. As Pylyshyn (2003) has argued at length, the representation in the brain may be as discursive and non-imagistic as any other standard computational vehicle.)

2.3 Empty Representations

The “phenomenally real” sense of “psychologically real” does invite a worry about just what we *are* talking about when we talk about phenomenally real things. If what I have said so far in this paper is correct, then, when we talk about and take ourselves to hear SLEs, we are not talking about or hearing *anything real*. But does this imply that linguistics is about *nothing*, that it’s no better off than some Macbethian theory of the daggers he sees, or a theologian’s theory of angels?

Well, yes and no. The trouble lies with odd features of intentional talk. There are two ways that representations can be “about nothing”: they can be *meaningless*, as in the case of nonsense expressions, like ‘brillig’; or they could be perfectly meaningful, but, as in the case of ‘Zeus’, there might no *real things in the world* that they represent. These two ways of being “about nothing” give rise to a crucial ambiguity in ways of talking about “what representations represent”:

(REP) (i) If we are talking about a representation, *x*, of some *real thing y*,
then *x* represents that real thing *y* —thus ‘Nixon’ represents the actual man Nixon.

(ii) When there isn’t, as in the case of ‘Zeus’,
then we rely on talk about the *content* of the expression ‘y’

I call the first usage the “existential,” the second the “(purely) intentional” usage of ‘represent’. For purposes here, I shall be confining myself to the latter, purely intentional use. That is, I shall not presume “representation of *x*” entails the actual existence of *x*. And, as I said at the beginning, I shall take a provisional characterization of intentional content to be: *however we take x to be in a purely intentional usage*, i.e. when we talk about “a representation of *x* when there is no *x*.”

I was deliberately cagey in the way I phrased the second-clause of (REP): I said we “rely on talk about the content of the expression ‘y’.” It would be tempting to say “so a purely intentional use of ‘representation of *y*’, for lack of any *y*, is really about an intentional content.” This would however get the subject matter wrong: someone thinking about Zeus and his philandering ways is not thinking about the philandering ways of an intentional content. Speaking more carefully, we should rather say something like: when *x* represents a *y* that doesn’t exist, a person is standing in the thinking relation to the intentional content [*y*]; but this doesn’t entail she is thinking *about that intentional content* [*y*].

But if Zeus is neither a God, a representation of one, or an intentional content, what in the world is He? Unlike many who’ve addressed this problem (e.g. Meinong (19??), Parsons (1980)), I am loathe

to rely here on any special metaphysics. I agree with Cartwright's (1960/87) nice quip: "unreality is just that: it is not another reality" (p30). For me, Zeus is *nothing*, nowhere, nohow. So far as I can see, (REP) simply captures an odd way we have of talking and thinking about intentional content, such as {Zeus}: in talking in any detail about the Greek's ideas about Zeus we will easily fall into thought and talk about –*Zeus*: how he lived on Mt. Olympus and cavorted in the woods. Constantly prefixing such talk with "the Greeks ideas" or "representations of" would be as cumbersome as always keeping track of use and mention. Thus, it seems to me, do linguists talk about SLEs, when all that's really at stake are the contents of speaker's representations "of SLEs."

Indeed, at least provisionally, I suspect that we can understand talk of intentional inexistents along the following, metaphysically deflated lines:

(DEF) x is an intentional inexistent for a representational system S iff_{df}:
there is a representation in S that has the content $[x]$ and x doesn't exist.

where a system can be a community-shared work of fiction, a theological tradition, or --as in the cases that concern me here-- a system in a human or animal psychology, e.g. the visual system, or a language module. When the system is a perceptual system, and the illusion of existence is vivid and stable across people, one might talk about *perceptual inexistents*; in a linguistic system, *linguistic inexistents*. As this definition makes clear, the ontology remains purely an ontology of representations and their intentional contents, which are patently needed in any case for an adequate psychology.⁸

3. *Folieism*

So, on my view, SLEs don't exist. But what on earth then do I think is going on when people speak, hear, write and read? On the *folieist* view I'm proposing, a speaker intends to utter a certain SLE, and produces various representations of it in various sub-systems of her mind; this causes her to contract her articulatory system in certain ways that in turn cause certain wave forms to traverse the air, impinge upon the linguistic/auditory modules in hearers, which are caused to produce representations of the very SLEs the speaker intended to utter.⁹ This process is so stable and reliable, involving specific processes in production and perceptual modules, that everyone concerned seems actually to "hear" tokens of those SLEs. However, it's an illusion: examination of the acoustic events reveals no entities with the structures that linguists have reasonably argued that SLEs possess. As I like to put it, the whole process is a kind of *folie à deux* (or *à n*, for the n speakers of a "common language"): the speaker has the illusion of uttering an SLE that the hearer has the illusion of hearing, with, however, the happy result that the hearer is usually able to determine precisely what the speaker intended to utter. And since, unlike the case of

8. I don't mean to underestimate the difficulties in providing a semantic theory of talk about intentional inexistents –or, for that matter, for natural language generally, which, n.b., seems to cheerfully include such talk (there is not even a suspicion of contradiction in such ordinary assertions as "Some things don't exist, for example, Zeus).” A step in the right direction seems to me to have been taken recently by Colin McGinn (2000:ch 2), who urges toleration of such talk by removing existential import from the so-called "existential" (what he prefers to call a "partial") quantifier, *some*.

9. For brevity, I shall throughout this discussion speak of intentions (and intending) to utter SLEs in a very weak sense, entailing only the instigation of action, and not entailing anything to do with conscious deliberation or the like. Indeed, many of the intentions I have in mind may well be simply imperatives of a speech production module, insensitive to the "central," reflective thoughts of the very sort I am defending in this paper, viz., that, try as one might, one is very unlikely to produce an SLE.

many other human artifacts, such as tables or cars, there are no causal/explanatory roles independent of human psychology for the supposed actual tokens of SLEs to play, there is no reason to posit them as real objects.

This last remark is intended to reassure the metaphysical realists among you that, no, I haven't gone "post-modern," to the extent of becoming an idealist about *everything*, along lines pursued by Jackendoff (19??, 20??) and sometimes even Chomsky (2000).¹⁰ I am concerned only to deny the existence of things like SLEs, where the only reality is so-called "psychological reality," a crucial phrase I'll discuss below. Other examples that I discuss elsewhere and I'll mention in passing are Kanizsa figures, mental images, and so-called "qualia." One thing I'm anxious to do is distinguish these cases from cases like tables, chairs, automobiles, electrons and stars in which, *pace* Chomsky and Jackendoff, it seems to me we have reasons *independent* of psychology to believe in their existence.

The analogy with vision is quite close. Much as certain "pacman" figures trigger the "perception" of standard Kanizsa figures, so do certain forments give rise to perceptions of SLEs. Indeed, entire alphabets have been constructed Kanizsa-style, where the letters are, as it were, intimated by various "pacman" corners, without there being any actual tokens of them at all. Just as people might routinely communicate by exploiting a Kanizsa alphabet, so, on my view, do they routinely communicate by providing each other similarly fragmentary auditory cues (another example, is, of course, handwriting, particularly as practiced by doctors).

Despite his resistance to intentionality, this story ought really be music to Chomsky's rationalist ears. Precisely as the rationalists claimed, many of our most fundamental ideas, such as those of Euclidean forms and articulatory gestures are not derived from sensory stimulation, but are understood independently of it;¹¹ they are ideas of certain sorts of computationally tractable ideals, applied in only approximate ways to the complex, ragged stimulation of sense, involving "things" that often do not and, (in the case of Euclidean forms) sometimes couldn't possibly exist. I differ from Chomsky only in insisting that such ideas are, indeed, *ideas* –intentional contents– not mere representations of them.

References:

Bromberger, S. and Halle, M. (2000), "The Ontology of Phonology (Revised)," in Burton-Roberts, N., Carr, P., and Docherty, G., *Phonological Knowledge: Conceptual and Empirical Approaches*, Oxford University Press

Cartwright, R. (1960/87), "Negative Existentials," in his *Philosophical Essays*, Cambridge: MIT Press

10. See his endorsement of "naming as a kind of 'world-making,' in something like Nelson Goodman's sense" (2000:181), which is followed by expressions of sympathy with various forms of 17th-18th Century idealism, according to which "The world as known is the world of ideas," which Chomsky (2000:182) quotes approvingly from Yolton's (1994) account of that period. See also McGilvray (2000:5-6) for extended interpretation of Chomsky along these lines.

11. Cf. "[the] link between articulatory and acoustic aspects of speech is...innately specified, requiring only epigenetic development to bring it into play" (Halle and Stevens 1991, citing Liberman and Mattingly (1985, 1989); Halle 2002:185).

- Chomsky, N. (1955/75), *Logical Structure of Linguistic Theory*, New York: Plenum
- , (1980), *Rules and Representations*, New York: Columbia University Press
- , (2000), *New Horizons in the Study of Language*, Cambridge University Press
- , (2003), “Reply to Rey,” in L. Antony and N. Hornstein, *Chomsky and His Critics*, Oxford: Blackwell
- Chomsky, N. and Halle, M. (1968), *The Sound Pattern of English*, New York: Harper and Row
- Dennett, D. and Kinsbourne, M., (19??), "Time and the Observer," BBS ??
- Devitt, M. (1996), *Coming to Out Senses*, Cambridge University Press
- Devitt, M. (forthcoming), *Ignorance of Language*, Oxford UP
- Evans, G. (1982), *Varieties of Reference*, Oxford University Press
- Fodor, J., (1975), *The Language of Thought*, New York: Crowell
- , (1990), *A Theory of Content*, Cambridge (MA): MIT Press
- , (1998), *Concepts: Where Cognitive Science Went Wrong*, Cambridge (MA): MIT Press
- Fodor, J. , Bever, T., and Garrett, M. (1974), *The Psychology of Language*, New York: McGraw Hill
- Goodman, N. (1956), *Fact, Fiction and Forecast*, Indianapolis: Bobbs-Merrill
- Jackendoff, R. (1987), *Consciousness and the Computational Mind*, MIT
- Jutronic, D. (2005), “Chomsky Amongst the Philosophers,” *Croatian Journal of Philosophy*
- Kenstowicz, M. (1994), *Phonology in Generative Grammar*, Cambridge (MA): Blackwell's
- Laver, J. (1993), *Principles of Phonetics*, Cambridge UP
- Lieberman, A. (1996), *Speech: A Special Code*, Cambridge: MIT Press
- Lieberman, A., Borman, M. and Rafael, L. (19??/96), "Some Experiments on the Sound of Silence in Phonetic Perception," in Liberman (1996), pp325-49
- Lieberman, A. and Mattingly, I. (1985/96), “The Motor Theory of Speech Revised,” *Cognition* 21:1-36
reprinted in Liberman, A. (1996), pp237-265 (references are to this edition)
- McGinn, C. (2000), *Logical Properties: Identity, Existence, Predication, Necessity, Truth*, Oxford University Press.
- Meinong, A. (????),
- Millikan, R. (2000), *On Clear and Confused Ideas*, Cambridge University Press

Millikan, R. (2005), in Schantz, R., *The Externalist Challenge: New Studies on Cognition and Intentionality*,

New York: de Gruyter, pp??

Neander, K. (forthcoming), teleosemantics book, MIT Press

Parsons, T. (1980), *Nonexistent Objects*, New Haven, CT: Yale University Press, 1980

Peacocke, C. (1994), *Concepts: A Philosophical Study*, Cambridge (MA): MIT Press

Pyllyshyn, Z. (2003), *Seeing and Visualizing: It's Not What You Think*, Cambridge (MA): MIT Press

Rey, G. (1997), *Contemporary Philosophy of Mind: a Contentiously Classical Approach*, Oxford: Blackwell

Rey, G. (2003a), "Chomsky, Intentionality and a CRTT" in *Chomsky and His Critics*, ed. by L. Antony and N. Hornstein, Oxford: Blackwell

— (2003b), "Representational Content and a Chomskyan Linguistics," for *Epistemology of Language*, ed. by Alex Barber, Oxford University Press

— (2005), "Mind, Intentionality and Inexistence: an Overview of My Work," with "Replies to Commentators," *Croatian Journal of Philosophy* ??

— (2006), "The Intentional Inexistence of Language –But Not Cars," *Debates in Cognitive Science*, ed. by R. Stainton, Oxford: Blackwell

Sapir, E. (1933/63), "The Psychological Reality of the Phoneme," in D. Mandelbaum (ed.), *The Selected Writings of Edward Sapir*, Berkeley: University of California Press

Stevens, K. and Halle, M. (1967), "Analysis by Synthesis," ??

Taylor, K. (2003), *Reference and the Rational Mind*, Stanford: CSLI Publications