

*W(h)ither Semantics!(?)*¹

JEFFREY C. KING

Department of Philosophy, Rutgers University

1. Introduction

Call a semantics for a given language *externalist* just in case it assigns to any expression of the language in question an “entity in the world” as its semantic value (perhaps relative to a context or other parameters). Thus for example, the (extensional) semantics in the now standard semantics textbook Heim, Kratzer [1998] is externalist since it takes names to be type e and one place predicates to be type $\langle e, t \rangle$. This means that the semantics assigns individuals to names and sets of individuals to one place predicates. So externalist semantic theories posit a semantic relation (perhaps relative to a context or other parameters) between expressions of the language and entities in the world. Such relations have gone by many names: ‘___having___ as its semantic value (relative to context c)’; ‘___refers to ___ (relative to c)’; ‘ $\| _ \| ^c = _$ ’; etc. I think it is safe to say that much recent semantic theorizing is externalist in this sense.²

Chomsky would call such theories (*externalist*) *reference-based semantics* and would say that on e.g. the Heim Kratzer semantics described above names and one-place predicates *refer* to things in the world. For our purposes in discussing externalist semantics we can stick with names and one-place predicates. Notoriously, Chomsky has been a vocal opponent of externalist semantics. For example, Chomsky [2000b] writes:

As for semantics, insofar as we understand language use, the argument for a reference-based semantics (apart from an internalist syntactic version) seems to me weak. It is possible that natural language has only syntax and pragmatics; it has a “semantics” only in the sense of “the study of how this instrument, whose formal structure and potentialities of expression are the subject of syntactic investigation, is actually put to use in a speech community,” to quote the earliest formulation in generative grammar 40 years ago, influenced by Wittgenstein, Austin and others (Chomsky 1955/1975; 1957: 102–3). In this view, natural language consists of internalist computations and performance systems that access them along with much other information and belief, carrying out their instructions in particular ways to enable us to talk and communicate, among other things. There will be no provision for what Scott Soames calls “the central semantic fact about language, . . . that it is used to represent the world,” because it is not assumed that language is used to represent the world, in the intended sense.³

As this quotation itself indirectly suggests, Chomsky seems to think that what he calls his *internalist* approach to the study of language precludes giving an externalist semantics for a language. The goal of the present paper is to assess whether Chomsky is right about this. The paper proceeds as follows. In section 1 I spell out the details of Chomsky's internalism. Though I spell out Chomsky's view in a fairly complete way, I concentrate on features of his view that he appeals to in arguing against externalist semantics. In section 2 I state some of Chomsky's arguments against externalist semantic theories and respond to them. In section 3 I discuss some Chomsky inspired arguments against externalist semantic theories from Pietroski [2003]. Section 4 is the conclusion.

2. Chomsky's internalism

One sometimes hears it said that for Chomsky linguistic theory is a branch of psychology. Roughly, something like the following is meant by this. For Chomsky, a theory of language, including semantics, is a theory of what a speaker knows when she knows a language. The idea is that a speaker of a language is so in virtue of having an internal representation of a theory of her language. Linguistic theory, including semantics, is that branch of psychology whose aim it is to give an explicit characterization of this internally represented theory. In this section I want to say a bit more about how to understand these Chomskian claims and give a more exact characterization of the sense in which linguistic theory is a branch of psychology on the view in question. As promised, I'll then return to the question of whether the claim that linguistic theory is a branch of psychology in the relevant sense conflicts with externalist semantic theories many practicing semanticists currently formulate.

Details and terminology aside, I think that from 1965 to the present day Chomsky has been consistently pretty clear about the sense in which theory of language is psychology. Consider first the following passage from Chomsky [1965]:

To learn a language, then, the child must have a method for devising an appropriate grammar, given primary linguistic data. As a precondition for language learning, he must possess, first, a linguistic theory that specifies the form of the grammar of a possible human language, and, second, a strategy for selecting a grammar of the appropriate form that is compatible with the primary linguistic data. As a long-range task for general linguistics, we might set the problem of developing an account of this innate linguistic theory that provides the basis for language learning. (Note that we are again using the term "theory" - in this case "theory of language" rather than "theory of a particular language" - with a systematic ambiguity, to refer both to the child's innate predisposition to learn a language of a certain type and to the linguist's account of this.)⁴

Though Chomsky does not here say that linguistic theory is a branch of psychology, he does set as its goal the characterization of a speaker's internally represented theory that specifies what it is to be a possible human language *and* the characterization of the speaker's internally represented theory that selects a particular possible human language given primary linguistic data. The former Chomsky calls *universal*

grammar (UG) and it constitutes the initial state of our biologically endowed *faculty of language* (FL); the latter is a mechanism for selecting what Chomsky calls a particular *I-language* on the basis of primary linguistic data. I'll say more about I-languages shortly. One goal of linguistic theory, then, is to provide a characterization of UG and of the mechanism by means of which speakers select I-languages based on primary linguistic data. Chomsky doesn't say so here, but another goal is to give characterizations of particular I-languages.

So Chomsky's claim here is that speakers of a language possess, as a matter of their biological endowment, an internally represented theory of UG (the initial state of FL) and an internally represented mechanism for selecting a particular I-language given primary linguistic data. Of course once such an I-language has been selected, it too is internally represented. Linguistic theory has as its aim characterizations of UG, the mechanism for selecting an I-language and particular I-languages.

Let's look at a more recent statement of what I take to be very similar points. First, a preliminary: Chomsky [1986] characterizes an *I-language* as follows:

A rather different approach was taken, for example, by Otto Jespersen, who held that there is some "notion of structure" in the mind of the speaker "which is definite enough to guide him in framing sentences of his own," in particular, "free expressions" that may be new to the speaker and to others. Let us refer to this "notion of structure" as an "internalized language" (I-language). The I-language, then, is some element of the mind of the person who knows the language, acquired by the learner, and used by the speaker-hearer.⁵

So an I-language is the speaker's internally represented theory of her particular language that allows her to speak and understand it. A bit later Chomsky writes:

Summarizing, then, we have the following general picture. The language faculty is a distinct system of the mind/brain with an initial state S_0 common to the species . . . and apparently unique to it in essential respects. Given appropriate experience, this faculty passes from the state S_0 to some relatively stable state S_s , which then undergoes only peripheral modification . . . The attained state incorporates an I-language (it is the state of having or knowing a particular I-language). UG is the theory of S_0 ; particular grammars are theories of various I-languages.⁶

He goes on to say:

UG and theories of I-languages, universal and particular grammars, are on par with scientific theories in other domains . . . Linguistics, conceived as the study of I-languages and S_0 , becomes part of psychology, ultimately biology.⁷

Before saying more about what is meant here, it is important to point out, as Chomsky [1965] does himself in the quotation above, that Chomsky uses the terms 'grammar', 'theory of language', 'UG' and 'I-language' ambiguously. Sometimes they are used to mean the internally represented theories of speakers; and sometimes they are used to mean the theories linguists offer of the internally represented theories:

The term “grammar” was then used with systematic ambiguity, to refer to what we have here called “I-language” and also to the linguist’s theory of I-language; the same is true of the term UG, introduced later with the same systematic ambiguity, referring to S_0 and the theory of S_0 .⁸

That caveat having been addressed, let’s return to the passages quoted above from Chomsky [1986]. Chomsky claims there is an initial state S_0 of a speaker’s biologically endowed language faculty. One job of linguistic theory is to give a characterization of this state. Linguists’ attempts to do this are the theories they offer of UG. Being presented with an appropriate amount of primary linguistic data, speakers pass from the state S_0 to a relatively steady state S_s , which is the state of having or knowing or “cognizing” a particular I-language. A second job of linguistic theory is to give characterizations of the various S_s ’s speakers internally represent. Linguists’ attempts to do this are the syntactic, phonological and semantic theories they offer of particular languages. These states S_0 and S_s are states of the human mind/brain. Hence, in trying to characterize these states linguistic theory is doing psychology. To the extent that linguistic theory concerns itself with how these states are realized physically in the brain, linguistics is essentially neuropsychology or even, as Chomsky says, biology. Chomsky frequently compares the work linguists are doing in trying to characterize S_0 , S_s and the language faculty generally with empirical work psychologists are doing in attempting to formulate theories of the human visual system.⁹ Chomsky certainly thinks of the attempt to formulate theories of the human visual system as work in psychology, neuropsychology and biology. Hence, the same goes for attempts to characterize S_0 and S_s . This is the sense in which linguistic theory is psychology for Chomsky and to the extent to which semantics is part of this project, it too is psychology for Chomsky in this sense.

We’ll get back to semantics in a moment, but before doing that let me emphasize some general themes in Chomsky’s view here. First, as we saw, a given I-language is the relatively stable state S_s of a particular speaker’s mind/brain attained as a result of exposing the initial state S_0 of FL to primary linguistic data, (or the correct theory of the state S_s of a particular speaker’s FL—recall the systematic ambiguity of ‘UG’, ‘I-language’ etc.). Hence, theories of particular languages are theories of certain mental states of particular speakers. This, of course, is in part why Chomsky uses the term ‘I-language’¹⁰: to mark that the objects of study for a theory of language are not “public languages” like English, Chinese and so on. This brings me to a second point. There is a sense in which Chomsky denies the existence of any public languages like English and Chinese.¹¹ Perhaps it would be better to say that Chomsky thinks that there is no non-arbitrary way to demarcate a purported public language like English and that in any case, such a thing is not the proper object of study for a scientific investigation of human languages. When people like Annie and Mary are (truly!) said to both speak English, this is made true by the fact that the relatively stable states of their respective FLs are sufficiently similar.¹² A final point: because for Chomsky the theory of language is concerned with characterizing internal states of speakers’ minds/brains, he often calls his approach to the study of language *internalist*.¹³

3. Chomsky's arguments against externalist semantics

Let's just grant that Chomsky is right about all the views I have attributed to him to this point. Does any of this conflict with externalist semantics? That is, is there any argument from Chomsky's internalism as just described, together with some additional Chomskian claims to be outlined below, to the conclusion that externalist semantics is misguided?

You might think that it is just trivially true that Chomsky's internalist approach to studying language rules out externalist semantics. On Chomsky's view, a theory of language is trying to characterize states of the minds/brains of particular speakers. If that's right, how could things in the real world outside of people's brains figure in semantics and so in the study of language? But for the most part, Chomsky doesn't seem to claim that his internalist approach just rules out externalist semantics from the start. After all, he gives *arguments* against externalist semantics. So I think the way to take our current question is this. Let's grant Chomsky's internalist approach to language as I have characterized it to this point. Note I've said almost nothing about semantics. Our question is: assuming Chomsky is right to this point, when we turn to semantics are there reasons for rejecting externalist semantics or is as much externalism as externalist semantics requires a legitimate addition to his internalist approach? As I've said, Chomsky argues against externalist semantic theories for natural languages.¹⁴ In the remainder of the paper, I want to examine whether Chomsky has given good reason for rejecting externalist semantics. It's not possible to address *every* argument Chomsky gives against externalist semantics and I'll mostly stick to the arguments he gives in Chomsky [2000b], which is the main repository for such Chomskian arguments. But I do think I will have covered the main *types* of arguments Chomsky and his supporters employ.

I begin with a bit more background. Lexical items, and linguistic expressions generally, have a sound and a meaning.¹⁵ Because of this, Chomsky holds FL interacts with at least two other systems in the human mind/brain: the articulatory-perceptual system, which deals with producing and perceiving the sounds of linguistic expressions; and the conceptual-intentional system, which deals with linguistic meaning. Focusing now on lexical items, they must somehow encode instructions about their sound to feed into the articulatory-perceptual system and about their meaning to feed into the conceptual-intentional system. Chomsky holds that lexical items do this by simply being bundles of semantic and phonological (and other—formal—) *features*.¹⁶ The phonological features of a word provide information about a word's sound to the articulatory-perceptual system; and the semantic features provide information about its meaning to the conceptual-intentional system.¹⁷ Here we want to focus on Chomsky's views about semantic features of words.

Chomsky [2000b] does not give a lot of examples of semantic features, but he does say several things about them. Following the form of many of Chomsky's examples, I'll use *that* clauses to specify semantic features; further, especially for our purposes, I'll take the paradigm case of a semantic feature to be: *that [the word] is/can be used to refer to . . .*¹⁸ A word like 'book', Chomsky says, has as semantic features that it is used to refer to an artifact and that it is nominal. The

word ‘bank’ in the sense of financial institution has as semantic features that it can be used to refer to concrete buildings of a certain sort (‘The bank burned down.’) and that it can be used to refer to a certain sort of financial institution (‘The bank raised interest rates.’).¹⁹ The word ‘bottle’ has as semantic features that it can be used to refer to containers of a certain sort and that it can be used to refer to the contents of containers of that sort (‘The baby finished the bottle.’).²⁰

A few caveats before continuing. Note that in expositing Chomsky’s view here, I have been talking about *words being used to refer* and not *words referring* or *having referents*. To hold that *words* have (semantic) referents is to endorse an externalist semantics and we are in the process of explaining Chomsky’s arguments *against* such approaches. Further, though I have glossed semantic features as *that ‘book’ can be used to refer an artifact*, where I am using an expression ‘refers’ that picks out what is for Chomsky at least a three-place relation between a speaker, a word and something in the world, semantic features must be understood completely internally, as being part of the I-language and so part of the state of the brain of the language user. Recall that semantic features must be “legible” to the conceptual-intentional system and so provide information to it. To foreshadow things I am about to turn to, that ‘book’ has a semantic feature like *that ‘book’ can be used to refer to an artifact* has to mean something like ‘book’ has an I-language internal semantic feature that when “read” by the conceptual-intentional system defeasibly allows the agent to know that it can be used to refer to an artifact.

Chomsky seems to say that semantic features play a role in explaining at least two things: 1. That a word is used by X to refer to a certain thing(s)²¹; and 2. certain sorts of entailments.²² Taking the former first, I think the idea is that if X uses ‘bank’ to refer to certain sorts of buildings, then ‘bank’ in X’s I-language has the semantic feature that it can be used to refer to such buildings. Chomsky is less explicit on the latter point, but the entailments he has in mind are things like *being a house* entails *being a building*. I suspect that the idea is that if I refer to a certain kind of structure using the word ‘house’, I do so by exploiting one of its semantic features like *that ‘house’ can be used to refer to such-and-such sort of building*. But then this will guarantee that the thing I referred to is a building.

Chomsky also says that ‘book’ has as semantic features that it is used to refer to a material thing and that it is used to refer to an abstract thing.²³ Similarly, Chomsky says that it is a semantic feature of ‘house’ that it can be used to refer to something material and that it can be used to refer to something abstract.²⁴ These comments about the semantic features of ‘book’ and ‘house’ suggest that a word’s features can *conflict* in the sense of placing incompatible demands on what it can be used to refer to (something material vs. something abstract). Chomsky confirms that this is his view. In various places, he talks about semantic features providing *perspectives* by means of which we can think and talk about the world; and in several places he notes that a single term can provide perspectives that conflict in the sense just described.²⁵ In such cases, speakers can use a single word and by exploiting different conflicting features in different contexts *refer* to different kinds of things. To illustrate with ‘book’, I can exploit its semantic feature that it can be used to refer to a material thing, and truly say that John and Peter checked out *different*

books when each checks out a copy of *War and Peace*. I thereby use ‘books’ to refer to material things that are distinct. But I could also exploit the semantic feature of ‘book’ that it can be used to refer to an abstract thing and say that John and Peter checked out the *same* book. I thereby use ‘book’ to refer to an abstract thing that is not the thing referred to in my previous use of the word. Chomsky’s idea here is that language users can attend now to one and now to another of the conflicting semantic features of a word, or even to both at the same time. Chomsky [2000b] illustrates this latter possibility with ‘book’ in the following passage and there are similar passages involving ‘door’, ‘bank’ and even ‘London’:²⁶

Investigating language use, we find that words are interpreted in terms of such factors as material constitution, design, intended and characteristic use, institutional role, and so on. Things are identified and assigned to categories in terms of such properties – which I am taking to be semantic features – on a par with phonetic features that determine its sound. The use of language can attend in various ways to these semantic features. Suppose the library has two copies of Tolstoy’s *War and Peace*, Peter takes out one, and John the other. Did Peter and John take out the same book, or different books? If we attend to the material factor of the lexical item, they took out different books; if we focus on its abstract component, they took out the same book. We can attend to both material and abstract factors simultaneously, as when we say that “the book that he is planning will weigh at least five pounds if he ever writes it,” or “his book is in every store in the country.”²⁷

Concerning the examples involving Peter and John and the unknown author last mentioned, call him ‘Bill’, consider the following sentences:

1. Peter and John took out the same book.
2. Peter and John took out different books.
3. Bill’s book is in every store in the country.
4. The book Bill is planning to write will weigh five pounds if he ever writes it.

As I’ve indicated, Chomsky thinks that the apparent truth of 1 and 2 when uttered in appropriate contexts with no change in what Peter and John did (each checking out a copy of *War and Peace*) shows that speakers can exploit the perspectives offered by *different* semantic features of ‘book’ to think and talk about different aspects of the world/refer to different things. Chomsky thinks the fact that we can take a sentence like 3 to be true shows that we can simultaneously exploit the conflicting semantic features that ‘book’ is used to refer to something material and that ‘book’ is used to refer to something abstract: we must be talking about something abstract if one thing is in every store, but we are also talking about something concrete since we are claiming that some material thing is in every store. I am not sure whether Chomsky wants to say that we *refer* to two different things with one use of ‘Bill’s book’ in cases like 3. But in similar cases with anaphoric pronouns such as

5. The bank burned down and it moved across the street.

Chomsky thinks that ‘The bank’ refers to a building and ‘it’, which is anaphoric on it, refers to an institution.²⁸

This at last puts us in a position to state Chomsky’s first argument against externalist semantic theories according to which names and 1-place predicates (the *expressions*) refer to things in the world. One more caveat before doing that. Chomsky sometimes rails against the idea of an expression having a referent in a *public language*. The latter, of course, contrasts with Chomsky’s favored notion of an I-language. But then part of the reason Chomsky objects to the notion of *having a referent in a public language* is because he objects to the notion of a *public language*. To focus more on the question of whether there is a legitimate externalist notion of *reference*, let’s grant Chomsky the claim that I-languages are the proper objects of scientific study and that public languages are dubious entities. We can then ask whether he offers any arguments that show that we can’t or shouldn’t have externalist semantics *for I-languages*. Obviously, he can’t object to the notion of *having a referent in an I-language* because he objects to the notion of an *I-language*. Henceforth, then, assume we are considering a particular I-language that we would call English, in the sense that it is Jone’s I-language and we would ordinarily say that Jones speaks English. Here is a quote in which Chomsky gives an argument that Jone’s I-language cannot be given an externalist semantics:

Contemporary philosophy of language follows a different course. It asks to what a word refers, giving various answers. But the question has no clear meaning. The example of “book” is typical. It makes little sense to ask to what thing the expression “Tolstoy’s War and Peace” refers, when Peter and John take identical copies out of the library. The answer depends on how the semantic features are used when we think and talk, one way or another. In general, a word, even of the simplest kind, does not pick out an entity of the world, or of our “belief space.” Conventional assumptions about these matters seem to me very dubious.²⁹

This passage comes shortly after the passage quoted previously in which Chomsky claims that by exploiting different semantic features of ‘book’, speakers can refer to different things (a material thing vs. an abstract thing). Again, Chomsky takes sentences like 1 and 2 to show this. In the present passage his point seems to be that since what a speaker refers to in using a word depends on the semantic feature he exploits in using it, as shown by 1 and 2, it makes no sense to ask what *the word itself* refers to. Presumably that question only has an answer if we specify a word and a semantic feature being exploited in its use. Hence, the question what a word refers to *simpliciter*, and so any answer to it, has no clear sense. But then externalist semantic theories that claim that words have referents are making claims with no clear sense.

This argument is only as strong as the claim that data like 1–4 above cannot be accounted for by an externalist semantics and must be accounted for by a theory that posits semantic features in such a way that the claim that words have references ends up having no clear meaning. This claim strikes me as rather weak. 1 and 2 appear to be an instance of polysemy. The natural thing for the externalist to say in this case is that count nouns like ‘book’ quite generally have one meaning on which

they refer to book tokens and another meaning on which they refer to book types. Of course this does mean that strictly speaking there is no straightforward answer to the question “what does ‘book’ refer to *simpliciter*?” But that fact doesn’t force us to reject an externalist semantics any more than the phenomenon of ambiguity does. Further, we can even go along with the idea that ‘book’ has two different I-language internal semantic features corresponding to the two polysemous external meanings. With respect to things like 3, there are several things one could say on an externalist account that posits polysemous externalist meanings. One is that these are cases where both polysemous meanings are invoked simultaneously so that one is making claims both about book tokens and book types (compare ‘Glenn’s name is on the front page of every newspaper in the country’). But another is that only the book type meaning is invoked and that a book type can have a property like *being in every store in the country*. It’s just that the way it has that property is by having a token in each store. Finally, turning to 4 and 5, let’s focus on 5. The natural thing to say here is that when a pronoun is anaphoric on a polysemous noun where the noun is interpreted with polysemous meaning M_1 , the pronoun can be interpreted with polysemous meaning M_2 of the antecedent noun, at least in cases where M_1 and M_2 are “sufficiently close” meanings. Hence ‘the bank’ in the first clause refers to a building (has meaning M_1) and ‘it’ in the second clause refers to an institution (has meaning M_2).

The explanation just given of the data 1–5 posits polysemous meanings of ‘book’ corresponding to *book type* and *book token*. Similar explanations could be given of the relevantly similar data involving Chomsky’s other examples ‘bank’, ‘bottle’, ‘door’ and so on, except that for these examples the relevant polysemous meanings won’t correspond to the type/token distinction.³⁰ Hence it appears that an externalist semantics that posits multiple polysemous meanings for words like ‘book’ etc. can account for 1–5 above. Though I don’t claim that polysemy is completely understood on an externalist semantics, there seem to be sensible, promising things to say here. Hence, it seems to me that to jettison externalist semantics upon encountering 1–5 is a serious overreaction.

A second argument Chomsky offers against externalist semantics is that positing real world references for words will involve us in inextricable puzzles. As before, Chomsky points out that a word can have different semantic features that allow a speaker to refer to different things by means of it. The example I’ll focus on here is ‘door’. Chomsky notes that you can paint the door white, and so here the word is used to refer to a concrete material object. But you can also walk through the door, and so here the word is used to refer to something like an opening of a certain sort. He makes similar observations about ‘house’, ‘water’ and ‘London’.³¹ He then writes

Note that the properties of such words as “house,” “door,” “London,” “water,” and so on do not indicate that people have contradictory or otherwise perplexing beliefs. There is no temptation to draw any such conclusion, if we drop the empirical assumption that words pick out things, apart from particular usages, which they constrain in highly intricate ways.³²

The thought is that our talk about painting and walking through the door does not indicate that we have “perplexing” beliefs about doors to the effect that there is this thing that we can both paint and pass through. However, Chomsky suggests, if we assume that ‘door’ has a real world referent D, people’s “door talk” *will* tempt us to ascribe to them the belief that D is something we can both paint and pass through. The thought is, then, that positing real worlds referents and ignoring I-language internal semantic features will lead us to ascribe to people contradictory or otherwise strange beliefs that they simply don’t have. Such puzzles are avoided by dropping the assumption that words have real world referents and attending to the fact that words have I-language internal semantic features that allow speakers to think and talk about different aspects of the world.

It is easy to see that the externalist semantic theory that posits multiple polysemous meanings that we invoked to respond to Chomsky’s first argument provides a response to this argument as well. ‘Door’ is polysemous and has both a *material barrier* meaning M_1 and an *opening* or *passage* meaning M_2 .³³ Ordinary people are tacitly aware of this and can easily be made explicitly aware of it. So they believe that they can paint the door given ‘door’ is understood as meaning M_1 and can walk through the door given ‘door’ is understood as meaning M_2 . But they don’t think that there is some one thing that they can both paint and pass through and we need not attribute any such belief to them. So positing external real world references does not involve us in puzzling ascriptions of belief in such cases.

I illustrated Chomsky’s second argument against externalist semantics just discussed, which claims that an externalist semantics will result in perplexing attributions of beliefs and other attitudes, with examples involving the word ‘door’. The response there was that an externalist semantics that posits multiple polysemous meanings for expressions will have no such problem, (a similar response would handle Chomsky’s similar examples involving ‘bank’, ‘house’, ‘water’ and so on). Chomsky sometimes suggests that a similar argument can be given involving expressions such as ‘Joe Sixpack’, ‘the average man’ and the like:

Suppose Peter says that Joe Sixpack voted for a living wage because he’s worried about his child’s health. Are we entitled to conclude that Peter believes the world to be constituted of such entities as Joe Sixpack, living wages, and health, and relations like voting-for and worrying-about that hold among them? Would the parallel inference be legitimate when Pete says that Tom visited Boston? If Peter says that the bank moved across the street after it was destroyed by fire, does he believe that among the things in the world there are some that can be destroyed but still be around, so that they can move?³⁴

We have already seen what to say about the case involving ‘bank’ at the end of the quotation here. As we saw, we can say that the antecedent ‘the bank’ has as its meaning what we can call *the institution meaning* M_1 and the pronoun ‘it’ anaphoric on it has what we can call *the building meaning* M_2 of ‘bank’, the latter being polysemous (the meanings here are reversed for 5 above). When polysemous expressions are antecedents for pronouns, the antecedent and pronoun can have *different* polysemous meanings of the antecedent. Thereby, the clause containing

the antecedent talks about an institution and the clause containing the anaphoric pronoun talks about a building. Speakers are tacitly aware of this and so do not believe in such a case that there is a thing that both was destroyed and subsequently moved. But what should we say about a case in which Peter said that Joe Sixpack voted for a living wage? Does an externalist semantics require us to say that Peter believes that the world contains an entity *Joe Sixpack*, so that again an externalist semantics saddles us with perplexing attitude attributions? Obviously not. The fact that we assign real world references to expressions does not mean we have to assign to all expressions that superficially look the same references of the same type. In particular, the fact that we take names to be type e, and so have individuals as references, does not require us to take ‘Joe Sixpack’ to be functioning semantically as a name here and so have a referent of type e. Similarly, ‘the Black Mamba’ designates the former basketball player Kobe Bryant and it looks like a definite description. Supposing that definite descriptions are generalized quantifiers, we nonetheless don’t have to assign ‘the Black Mamba’ an $\langle\langle e, t \rangle, t \rangle$ type denotation in its use to designate Kobe. It’s true that externalists about semantics need to say something about how ‘the average man’ and ‘Joe Sixpack’ function semantically. But that problem doesn’t seem intractable.³⁵ Finally, it seems quite clear that ordinary speakers are well aware that ‘Joe Sixpack’ is not just an ordinary name and so there is no motivation to attribute attitudes about some entity *Joe Sixpack* to people who claim that Joe Sixpack voted for a living wage.

Chomsky gives another argument against externalist semantics that involves ‘the average man’. Actually, it is more of an argument to the effect that though there is an empirically motivated justification for positing a relation R (“reference”) between expressions and elements of a specified domain D, R cannot be a relation between expressions and things in the world. Hence we are left without any empirical motivation for the externalist’s *reference* relation between expressions and real world referents.

Chomsky [2000b] considers the following sentences:³⁶

6. a He thinks the young man is a genius.
- b The young man thinks he is a genius.
- c His mother thinks the young man is a genius.

noting that in 6b and c, ‘he’/‘His’ may be interpretatively dependent on ‘the young man’, but in 6a ‘He’ is obligatorily not interpretatively dependent on ‘the young man’. This of course is explained by the principles of Binding Theory. Chomsky is imagining stating the principles of Binding Theory in terms of coreference and being disjoint in coreference, understood as co-R/non-co-R above, and writes:³⁷

Within internalist semantics, there are explanatory theories of considerable interest that are developed in terms of a relation R (read “refer”) that is postulated to hold between linguistic expressions and something else, entities drawn from some stipulated domain D (perhaps semantic values). The relation R, for example, holds between the expressions *London* (*house*, etc.) and entities of D that are assumed to have some relation to what people refer to when they use the

words *London* (*house*, etc.), though that presumed relation remains obscure . . . Explanation of the phenomena of example [6] is commonly expressed in terms of the relation R. The same theories of binding and anaphora carry over without essential change if we replace *young* in example [6] by *average*, *typical*, or replace *the young man* by *John Doe*, stipulated to be the average man for the purposes of a particular discourse. The same theories also carry over to anaphoric properties of the pronouns in examples (3) and (4):

- (3) a It brings good health's rewards.
 b Good health brings its rewards.
 c Its rewards are what make good health worth striving for.
- (4) a [There is a flaw in the argument], but it was quickly found.
 b [The argument is flawed], but it was quickly found.

In terms of the relation R, stipulated to hold between *the average man*, *John Doe*, *good health*, *flaw*, and entities drawn from D, we can account for the differential behavior of the pronoun exactly as we would with *the young man*, *Peter*, *fly* ("there is a fly in the coffee"). The relations of anaphora differ in (4a and 4b), though there is no relevant difference in meaning between the bracketed clauses. And it might well turn out that these expressions, along with such others as "the argument has a flaw" (with the anaphoric options of (4a)), share still deeper structural properties, possibly even the same structural representation at the level relevant to the internal semantics of the phrases, a possibility that has been explored for some years (see Tremblay 1991). The same is true in more exotic cases. It would seem perverse to seek a relation between entities in D and things in the world – real, imagined, or whatever – at least, one of any generality.³⁸

Focusing on 'the average man' and 'John Doe' here (though for Chomsky similar remarks apply to 'good health' and 'flaw'), Chomsky's point is that the relation R ("reference") he mentions above will be used to state the principles of Binding Theory. But since those principles work equally well in a case in which 'the average man' or 'John Doe' replaces 'the young man' in 6, and since we know that 'John Doe'/'the average man' don't designate an individual in the way 'the young man' does, the relevant notion of "reference" R cannot be the relation between words and real world referents posited by the externalist. Rather, as Chomsky says, it is some relation between words and elements of some specified domain D (discourse referents?), whose relation to real world objects is "obscure". So though we have some justification for the internalist "reference" relation R in terms of its use to state the principles of Binding Theory, we so far have no justification/theoretical motivation for the externalist's reference relation that words allegedly stand in to real world objects.

I actually pretty much agree with what Chomsky says here. I don't think the need to state Binding Theory motivates or requires the externalist's reference relation. One way to see this is that Binding Theory was originally stated in terms of coindexing and not coreference/noncoreference.³⁹ And the relation between coindexing and real world reference is quite indirect. So though in 6a 'He' is obligatorily not

coindexed with ‘the young man’, they may in fact designate the same real world thing. And this does not violate Binding Theory. Further, we saw above that on the externalist semantic theory that invokes polysemy, we may well wish to say that an antecedent and an anaphoric, and so coindexed, pronoun are interpreted with *different* polysemous meanings. Hence, though coindexed they are not coreferential in any real world sense of ‘coreferential’. So to repeat, I agree that stating Binding Theory may require only internalistically acceptable notions and so doesn’t provide a motivation for the externalist’s real world reference relation. But that is far from a general argument *against* such a relation. At most it is an invitation to say something about what the motivation *is* for externalist semantics. Let me all too briefly discuss that.

In order to motivate *any* way of approaching semantics, we should begin by asking what the point of a semantics for a natural language is. What are we trying to do or explain in formulating such a semantic theory? There is a range of data that our semantic theories should play a significant part in explaining:

1. Productivity: speakers can produce and understand novel sentences, where such understanding consists at least in part in grasping the truth conditions of the sentence, which often involve objects, properties and relations.⁴⁰
2. Transfer of information: speakers are able to efficiently and systematically convey information about the external world to each other using sentences of the language they share (related to 1).⁴¹
3. Logical relations: speakers make robust judgments about entailment relations between sentences.
4. Truth and falsity: speakers make robust judgments about the truth or falsity of sentences taken relative to contexts and circumstances of evaluation, where such judgments are based on whether certain objects possess certain properties or stand in certain relations in those circumstances.

Let me say a few words about some of these claims. First, I take the objects, properties and relations mentioned in 1 to be mind independent entities in the external world. Suppose now I say to you ‘George W. Bush narrowly defeated Al Gore in the 2000 U.S. Presidential election by unscrupulous means.’ This sentence may well be novel to you and yet if you are a competent English speaker you may well understand it. In so doing, you must know that it is true iff one object (G.W. Bush) bears a certain relation to another (Al Gore). Regarding 2, by uttering as I did I might have conveyed to you a new piece of information about events in the external world. 1 and 2 entail that in understanding novel sentences speakers grasp information about the real world—real world truth conditions—and that they can systematically and efficiently exchange such real world truth conditions. Finally, concerning 4, having heard me utter as I did and supposing now that you knew the facts about the 2000 election, you have a firm view about whether the real world truth conditions obtain or not.

It is now easy to see how to motivate an externalist semantics on which names are assigned objects and predicates are assigned properties (or sets of individuals)

and relations and sentences are assigned truth-values relative to worlds considered as circumstances of evaluation.⁴² The motivation for such a semantic theory is that it assigns semantic values to expressions of our language in such a way as to give a simple, direct explanation of the above facts 1–4. With respect to 1, our externalist semantic theory compositionally assigns real world truth conditions to sentences of the language.⁴³ Speakers have tacit knowledge of that theory. Thereby, when a speaker hears a novel sentence, she is able to grasp its real world truth conditions. Regarding 2, speakers can also thereby systematically and efficiently exchange information about the real world—real world truth conditions—by uttering sentences. As to 3, a sentence S entails a sentence S' iff the real world truth conditions of S' obtain at any world at which the real world truth conditions of S obtain. If our theory assigns real world truth conditions correctly, it will explain speakers' judgments about entailment by assigning truth conditions to sentences in such a way that speakers judge that S entails S' iff the real world truth conditions of S' obtain at every world where those of S obtain. Similarly, in the case of 4 if our semantic theory assigns real world truth conditions to sentences correctly, speakers judge that S is true (false) relative to a world just in case S's real world truth conditions obtain (don't obtain) at the relevant world. To repeat, our externalist semantic theory simply and directly explains 1–4.

By contrast, Chomsky's internalist theory, including the account of semantic features discussed above, can't explain 1–4. The problem, of course, is that in not allowing semantic relations between expressions and things in the world like objects, properties (or sets of individuals) and relations, the internalist theory will not assign sentences real world truth conditions at all. Chomsky is explicit about this as is Pietroski [2003] and elsewhere in defending a Chomskian account.⁴⁴ But then something will need to be added to the Chomskian account to explain things like 1, 2 and 4 above. That the externalist semantic theory can simply and directly explain these things, while the Chomskian internalist account will require additional and unknown resources, provides motivation for an externalist account. Of course there may be insurmountable problems with externalist accounts and we are now canvassing arguments against them. But the point is we can respond to Chomsky's (correct) claim that an externalist reference relation isn't required to state the principles of Binding Theory and the (false) impression he gives that there is no motivation for such a reference relation by noting that being able to provide a simple, direct explanation of 1–4 above, unlike Chomsky's internalist account, *does* provide substantial motivation for pursuing an externalist semantics.

There is a further point that I can't go into in detail, but I'd like to briefly mention. Chomsky correctly notes that any notion of *reference* or *semantic value* or etc. employed by externalist semanticists to label a relation between words and real world entities is a technical notion. It is a purported theoretical term in a theory. Chomsky concludes from this that we can have no intuitions about what things stand in the relation the term "refers" to since the theory just tells us what things do.⁴⁵ He thereby seems to be chastising externalist semanticists who in defending semantic theories appeal to intuitions about what words refer to in various cases (e.g. Kripke [1980] in Gödel/Schmidt cases; Putnam [1975] on

‘water’).⁴⁶ I agree that the notions of *reference* or *semantic value* employed by externalist semanticists are technical terms in a semantic theory. But it doesn’t follow at all that the appeals externalist semanticists make to speakers’ intuitions, I would prefer *speakers’ judgments*, are illegitimate. After all, in constructing my semantic theory, I am among other things trying to capture certain judgments speakers make as indicated above in discussing 3 and 4. That in turn means that one way of assessing the theory is to ask whether it really *does* capture speakers’ judgments. To take a toy case, imagine I construct a semantic theory for English that claims in part that ‘bachelor’ refers to human males. Given other parts of the semantic theory (e.g. about what names refer to), this claim makes predictions about the truth conditions of sentences of the form ‘X is a bachelor’. Consulting the judgments of speakers whose I-languages are sufficiently like mine to count as English speakers, I discover that their judgments about the truth conditions of sentences of this form are not captured by my semantic theory. Surely, this is an important bit of evidence to consider in assessing my semantic theory. Could Chomsky really think otherwise?

4. Pietroski [2003] against truth-conditional semantics

Let me close with a type of argument against a semantics that assigns truth conditions to sentences (even relative to contexts) that is due to Paul Pietroski [2003] channeling Chomsky [1977, 2000b]. In defending the view that semantics does not assign truth conditions to sentences, Pietroski is arguing against externalist semantics and in favor of a view on which semantics is concerned only with Chomskian internalist semantic features of linguistic expressions.⁴⁷ The idea is that subsentential expressions don’t have the kinds of semantic values (even relative to context) that will compositionally determine truth conditions for sentences. That an utterance or use of a sentence has a truth condition (in a context) is taken by Pietroski, again following Chomsky, to be an *interaction effect* between internalist semantic features of sentences (and their constituents) and “a host of facts less amenable to theorizing, like facts about how “reasonable” speakers would *use* the sentence.”⁴⁸ Only the former are proper objects of scientific study in semantics. Henceforth, I’ll just discuss Pietroski, but it should be borne in mind that he takes himself to be closely following Chomsky, especially Chomsky [1977] and [2000b].

How would one defend the view that a semantics does not assign truth conditions to sentences relative to contexts and that subsentential expressions don’t have the kinds of semantic values in contexts that compositionally combine to yield truth conditions for sentences? Pietroski’s strategy is to give a bunch of cases in which he claims the truth conditions of sentences shift in certain ways where such shifts seem to compel us to assign different truth condition contributing “semantic values” to subsentential expressions. Pietroski seems to think that this sort of “contextual sensitivity” of semantically assigned truth conditions is implausible. Better, he thinks, to explain the shifts in truth conditions by saying that the semantic values of sentences and subsentential expressions are constant across the cases he considers and to say that differences in truth conditions are generated by the interaction

of these constant, internalist semantic values and various features of the context, including what a reasonable speaker is likely to mean in the relevant situation and so on. I can't here of course consider all the examples Pietroski and Chomsky give, but I want to consider a few representative types of examples and say what I think an externalist semanticist should say about them.

The first example concerns the following four sentences⁴⁹

7. Unicycles have wheels.
8. Jim's unicycle has wheels.
9. Cars have wheels.
10. Jim's car has wheels.
11. I have wheels.

Pietroski claims that 7 is true and 8 is false, so that the truth of 7 does not ensure the truth of 8. But he thinks that both 9 and 10 are true and that the truth of 9 implies that "each normal car has wheels" and hence that if Jim's car is normal it has wheels. Pietroski also claims that 11 is false as uttered by a unicycle. Pietroski's challenge is for the externalist to complete the following characterization of the function from individuals to truth-values she claims is the semantic value of 'has/have wheels' in such a way as to capture the data given by 7–11:

(HW) $\lambda x.true$ iff . . .

Prima facie there are two options for the externalist:

(HWs) $\lambda x.true$ iff x has at least one wheel.

(HWp) $\lambda x.true$ iff x has at least two wheels.

In Pietroski's view, neither works. HWs incorrectly makes 8 true. And HWp doesn't seem to explain how 7 manages to be true since we can't say that 7 is true just in case every normal unicycle has at least two wheels or just in case $\{x: x \text{ is a unicycle}\}$ bears the right generic relation to $\{x: x \text{ has at least two wheels}\}$. Pietroski considers the idea 'has wheels' denotes HWp, which makes 8 false and 10 true as desired, while 'have wheels' denotes HWs making both 7 and 9 true as desired. But he notes that this would incorrectly make 11 true as uttered by a unicycle. Finally, he considers the view that 'has wheels' denotes HWp, and that 'have wheels' is ambiguous between HWs and HWp, and generally denotes HWp but in generic sentences like 7 and 9 denotes HWs. To this Pietroski responds "One begins to wonder, though, if this is a theory of meaning or just a stipulation of the facts." I agree that as stated, the view in question does sound pretty *ad hoc*, especially in positing an ambiguity in 'have wheels' where the one meaning only manifests itself in generics. So the challenge is for the externalist to say something about the semantics of 'has wheels' and 'have wheels' that accounts for 7–11 without seeming *ad hoc*.

It seems to me, however, that there is a non-*ad hoc* view the externalist can defend here after all. As before, let 'has wheels' denote HWp. This makes 8 false and 10 true as required. Let 'have wheels' denote HWp with singular subjects

and HWs with plural subjects. Putting these two claims about ‘has wheels’ and ‘have wheels’ together, we have that ‘has/have wheels’ denotes HWp with singular subjects and HWs with plural subjects.⁵⁰ This makes 7 and 9 true and 11 false as required.

It also makes correct predictions about additional data. It correctly predicts that in all of the following, ‘have wheels’ denotes HWs and so the sentences are true (14 as uttered by the representative member of a group of unicycles; 13 and 15 uttered pointing at some unicycles):

12. All/Some/Most/Many unicycles have wheels.
13. Those unicycles have wheels.
14. We have wheels.
15. Those have wheels.

Further, it correctly predicts the contrasts between the following sentences (17a uttered addressing a single unicycle; 17b uttered addressing a group of unicycles):

- 16a. Each unicycle has wheels.
- 16b. All unicycles have wheels.
- 17a. You [2, singular] have wheels.
- 17b. You [2, plural] have wheels.

16a and 17a are predicted to be false and 16b and 17b true. Just so, it seems to me.

Finally, an anonymous referee pointed out that on my semantic proposal here, the following two sentences fail to be equivalent:

- 18a. One or two unicycles have wheels.
- 18b. One unicycle has wheels or two unicycles have wheels.

He/she thought this to be undesirable. But far from being undesirable, I think this is exactly the right result and so *supports* my semantic proposal. For I think it is clear that the sentences *are not* equivalent. In particular, when one unicycle has a wheel but two don’t 18a is true and 18b isn’t (since ‘One unicycle has wheels’ is clearly false, as is ‘Two unicycles have wheels.’—so both disjuncts of 18b are false). When I asked (sophisticated but not philosophical) informants whether the sentences were equivalent and told them to think hard about it, they said that they weren’t. So this is additional evidence in favor of my proposed semantics.

On this view, 9 does not entail 10, but recall that Pietroski only claimed 9 entails that normal cars have more than one wheel. Hence, he claimed that if 9 is true and Jim’s car is normal, 10 is true. I reject this because I think it reflects a naïve account of generics like 9. Suppressing exactly how to understand ‘Cars’ in 9, I think 9 means essentially that cars have at least one wheel. Further, keeping other elements uniform in 7 and 9 (how the generic is interpreted, etc.) 7 puts the same demand on unicycles that 9 puts on cars (that each unicycle/car in question have at least one wheel). Any feeling to the contrary results from our world knowledge that

cars generally have 4 wheels. World knowledge can creep in in other cases where it would be very strange to think there really are truth conditional differences, as in

- 19a. All humans have hands.
- 19b. All humans have heads.

You could get the feeling that the truth of 19a requires a given human to have multiple hands and 19b requires only at least one head. But surely a better explanation is that both require that a given human have at least one thing of the relevant sort and that we know that humans generally have two hands but a single head.

It seems to me that the resulting externalist semantic theory is not *ad hoc* at all and makes correct predictions about data like 7–18. Further, once we expand the data we look at beyond Pietroski’s 7–11 to include 12–18, the data looks quite systematic. If 7–11 were the result of ephemeral and protean interaction effects, we would hardly expect the data to exhibit the robust patterns we see in 7–18 taken together. I conclude that the externalist can meet Pietroski’s challenge here.⁵¹

Chomsky [1977] and Pietroski [2003] often cite complex data involving generics in favor of their views. Consider the following:

- 20a. Beavers are mammals.
- 20b. Beavers build dams.
- 20c. Dams are built by beavers.

20a seems to require all beavers to be mammals, whereas 20b does not require all beavers to be builders. The passive of 20b (20c) seems false whereas 20b seems true, and so on. I won’t address such arguments here except to say that it has been known for quite a while that generics are quite complicated.⁵² However, much progress has been made and one shouldn’t despair of getting a good theory this early in the semantics game.⁵³ To cite a well-known unsolved problem for a theory is not a very strong argument against it, especially when those giving the argument make no attempt to explain the data themselves (“see how unsystematic and crazy interaction effects can be?”). So to this sort of argument I want to simply say: yes, we’ve known for a long time this is going to be a hard one. But we have made progress and it’s a bit early to throw in the towel.

The final argument against a semantics that associates truth conditions with sentences I’ll consider is again due to Pietroski [2003]. As in the case of ‘has/have wheels’ above, the argument purports to show that there is no good externalist semantics for a subsentential expression. The argument concerns the following sentence

- 21. The bathhouse will be cleaned at 10AM.

Pietroski reports that he and his wife saw this sentence displayed in a Swiss campground and that when 10 AM came around they found a maintenance team cleaning the *outside* of the bathhouse. In fact, the inside was cleaned as well, but Pietroski

asks us to imagine that *only* the outside had been cleaned. Would 21 have been true? If we say that 21 would have been false in this situation, as Pietroski seems inclined to say, it seems that the externalist must say that ‘cleaned’ must be assigned a function that maps things like bathhouses to true iff their insides are cleaned. But this amounts to saying that ‘cleaned’ means *cleaned the inside of*, which makes 21 equivalent to the following:

22. The inside of the bathhouse will be cleaned at 10 AM.

But, Pietroski says, this seems wrong since it would mean that 21 can *never* be true if only the outside of the bathhouse is cleaned; and Pietroski apparently thinks that 21 *could* be true in some contexts when only the outside is cleaned. Further, *cleaning* can’t amount to *cleaning inside* generally on pain of the following being false when Norbert dusts the surface of his globe:

23. Norbert cleaned his globe today.

So here again, it appears that which function from objects to truth-values ‘cleaned’ denotes depends on the kind of thing you are talking about, what you are interested in and so on. Once again, the data look too unsystematic to be captured by an externalist semantics that assigns expressions like ‘cleaned’ functions from objects to truth-values. Better to assign ‘cleaned’ some more austere, internalist semantic value and have the various functions from objects to truth-values determined in context by interaction effects.

In response, the externalist should agree with Pietroski that *something* has gone wrong in the case in which the maintenance crew washes only the outside of the bathhouse in the campground where 21 is displayed. But she should resist the idea that she should capture this by having 21 come out *false* in that situation. Rather, she should say that there are many ways to clean a bathhouse: you can do so by scrubbing the exterior, scrubbing the interior or both. So 21 is just true in the imagined situation. By the externalist’s lights, Pietroski is multiplying truth conditions associated with 21 instead of saying that there is a single truth condition (polysemy aside) that can be satisfied in various ways. Still, I said that the externalist should agree that *something* went wrong in the scenario Pietroski considers. One way of seeing this is the following. Having seen 21 displayed and noting the less than immaculate condition of the inside of the bathhouse, suppose we decided to wait for the cleaning crew and enter after they had finished cleaning. We would be surprised and amused, if not disturbed, to see them scrubbing only the outside and leaving. These reactions indicate that something has gone wrong in this scenario.

Why do we react so? Well, a natural thought is that the fact that the campground displayed 21 indicates they took themselves to be informing campers about something likely to be of significance to them. That the inside of the bathhouse is to be cleaned daily at a certain time would have significance for campers, since they could not use the bathhouse at that time. Not so for an exterior scrubbing. Further, the usual way of cleaning a public bathhouse daily is to clean its inside. Thus, having seen 21, it is natural for the campers to expect that a cleaning of the *inside* of the

bathroom will occur at every 10 AM. But none of that makes 21 *false* if only the exterior gets cleaned at 10 AM. My judgment, at least, is that if I were to complain to campground management in such a case, I should complain not that 21 is false, but that it is misleading. In defense of this view, I'll only half in jest note that a Miami Herald article from August 29, 2014 read "HOUSE CLEANING MEANS THE EXTERIOR TOO" and went on to extol the virtues of and methods for cleaning the exterior of your home.

Similar points apply to the case of 23 and Norbert's dusting his globe's exterior. If Paul were to utter 23 and I walked into Norbert's office at the relevant time to find him scrubbing the *inside* of his globe, I perhaps will feel misled but should judge 23 true. But why should I feel misled in *this* case? Here I think an element is at work that I hinted at in the bathroom case and that needs further elaboration. In the case of globes, there is a *normal* or *stereotypical* way of cleaning them. One removes foreign material from the surface in some manner or other (try googling *cleaning globes*). Now it is a quite general feature of communication that when one describes a situation in a simple, straightforward way without qualification, if there is a stereotypical way of satisfying the description, one *conveys* that it is satisfied *in that way*. Many researchers have articulated something like this idea. For example, though I would disagree with much of what he says about it, Levinson [2000] has what he calls the *I-principle* that he thinks of as a heuristic for pragmatically conveying information and that he glosses as: *what is expressed simply is stereotypically exemplified*. 23 expresses a cleaning of the globe very simply, and hence conveys that it was done in stereotypical manner. Thus my feeling surprised and even misled in finding Norbert non-stereotypically cleaning the inside of his globe. But cleaning his globe, he was.

I'll close by saying that I think these points apply to lots of allegedly puzzling cases in the literature in which people cut grass with sowing scissors and so on. In such cases it is true that people are cutting the grass, but in saying only that you *convey* something false: that they are doing it in the stereotypical way. As evidence that people think the relevant sentences are literally true in non-stereotypical situations, I'll note the existence of *Literal Genie* or *Make A Wish* threads on the internet. The first person on the thread makes a wish and the next person grants it in a way that the initial person won't expect or like and makes her own wish. Here's a typical example:

A: I wish I owned an airplane.

B: Granted, but it has no wings or engine.

If we assume that "granters" are trying to grant the literal contents of the wishes (it's called *Literal Genie* after all), this example seems to illustrate the phenomenon in question. To say 'I own an airplane' is to convey that you do so in the normal stereotypical way: you own a functioning airplane that you or a pilot flies with some regularity etc. In this case, the granter B made A's wish literally true ('A owns an airplane' is true given B's grant), but failed to do so in the stereotypical way A suggested by her simple description of the plane-owning situation. Thus, B literally granted A's wish but not in the way A (merely) *conveyed* it should be granted.

Many, many examples on these threads exemplify this pattern, suggesting that speakers employ something like the pragmatic principle described above. Further, the following would be a completely appropriate, if less than hilarious, entry on such a thread suggesting that ordinary folk take 21 to be true in the scenario imagined:

A: I wish they would clean the bathhouse at the campground.

B: Granted, but they only cleaned the outside.

In addition to being careful what we wish for, when literal genies are around we should be careful *how we say* what we wish for.

5. Conclusion

Chomsky and his followers like Pietroski argue that Chomsky's internalist approach to language, including his account of semantic features, is incompatible with externalist semantics. In section 1 I outlined some central features of Chomsky's internalist approach to the theory of language. In section 2 I explained his account of semantic features. In sections 2 and 3 I outlined arguments Chomsky and Pietroski give against externalist semantic theories and provided responses to those arguments. This shows that one can accept the central features of Chomsky's internalist approach to the theory of language outlined in sections 2 and 3 and still endorse an externalist semantics for I-languages.

Notes

¹ A version of this paper was given at the University of Michigan Spring Colloquium March 12-13 2015, at the University of Notre Dame on April 17, 2015 and at the Institut Jean Nicod/École Normale Supérieure on June 4, 2015. Thanks to Daniel Drucker, my commentator at Michigan, and members of the audiences for helpful comments.

² Examples include Elbourne [2005], Glanzberg [2011], Heim and Kratzer [1998], Jeshion [2015], King [2001], Ninan [2010], Schwarzschild [2002] and Yalcin [2007] among myriad others.



⁴ p. 21–22.

⁵ Chomsky [1986] p. 25.

⁶ Chomsky [1986] p. 27.

⁷ Chomsky [1986] p. 29.

⁸ See for example Chomsky [1965] p. 8; Chomsky [1980] p.5, 11, 14; Chomsky [1986] p. 40; Chomsky [2000a] p. 90, p. 97.

⁹ “internal, individual and intensional”, Chomsky [2000b] p. 26.

¹⁰ E.g. see Chomsky [2000b] p. 148 where he claims that the notion of *public language* has not been formulated in coherent terms; and p. 155 where calls the notion “mysterious” and denies that ‘Chinese’ picks out anything in the real world. See also Chomsky [1986] pp. 24–31 for similar sentiments.

¹¹ Chomsky often compares talk of people *speaking the same language* with talk of people *looking alike or living near each other*. See Chomsky [2000b] p. 30–31. See also pp. 72–73 and 155.

¹² E.g. Chomsky [2000b] p. 15.

¹³ The restriction to natural languages is important here, since Chomsky allows that terms in statements of scientific theories may refer and so be given an externalist semantics. The linguistic media in which scientific theories are stated are importantly different from natural languages/I-languages according to Chomsky. See Chomsky [2000b] p. 22–24; 42–43; and 131–32.

¹⁵ An anonymous referee pointed out that this leaves out expressions of sign languages. I suspect Chomsky meant that expressions have a phonology and a meaning.

¹⁶ Chomsky [2000b] p. 10, 120.

¹⁷ Chomsky seems to hold that phonological and semantic features are innate and part of our biologically endowed FL. See Chomsky [2000b] pp. 33–34, 61–64, 120, 185.

¹⁸ Chomsky [2000b] describes semantic features using both formulations, as will I. See p. 10, 21, 34.

¹⁹ Chomsky [2000b] p. 16, 180.

²⁰ Chomsky [2000b] p. 128.

²¹ Chomsky [2000b] p. 187.

²² Chomsky [2000b] p. 34, 174.

²³ Chomsky [2000b] pp. 15–16.

²⁴ Chomsky [2000b] p. 36.

²⁵ Chomsky [2000b] p. 16, 36, 125, 126, 180. Chomsky also often talks about referring to something *as a certain kind of thing* (p. 20; 42) or *from a certain perspective/standpoint/point of view* (p. 20; 21; 36). I suspect that what Chomsky means by this is that speakers exploit particular semantic features of a word in referring to things with it and thereby refer to the thing from the perspective provided by the semantic features exploited.

²⁶ p. 36–37, 126–27, 180–81. Note that the inclusion of ‘London’ shows that Chomsky holds that names have semantic features just as do nouns like ‘book’ and ‘door’. See p. 180–81. He makes clear that he thinks names of persons such as ‘Peter’ have semantic features as well. See p. 10.

²⁷ pp. 15–16.

²⁸ See Chomsky [2000b] p. 126, 180.

²⁹ Chomsky [2000b] pp. 16–17. He appears to give the same argument on p. 181.

³⁰ As we saw, in the case of ‘bank’, it will have one polysemous meaning on which it refers to buildings and another on which it refers to institutions. For ‘bottle’, one meaning will have contents as referents and another will have containers of the relevant sort as referents. It is perhaps worth mentioning that the *Oxford English Dictionary* explicitly recognizes distinct, related meanings in the latter case.

³¹ Chomsky [2000b] p. 125–128. One finds similar remarks on p. 37 and 127.

³² Chomsky [2000b] p. 129.

³³ Again the *Oxford English Dictionary* notes both meanings.

³⁴ Chomsky [2000b] p. 172. See related remarks on p. 150, 178.

³⁵ See Kennedy and Stanley [2009] for some ideas about these issues.

³⁶ p. 35. Examples renumbered to fit in with the numbering in the present paper.

³⁷ Chomsky [1995] p. 100.

³⁸ Chomsky [2000b] p. 39.

³⁹ Chomsky [1988] pp. 184–89.

⁴⁰ I suppress contextual sensitivity here since it isn’t relevant to present concerns.

⁴¹ Of course, what this really means given our focus on I-languages is that two speakers whose I-languages are sufficiently similar can efficiently and systematically convey information to each other.

⁴² This will be similar to the semantics in Heim and Kratzer [1998] pp. 299–312.

⁴³ It is an open question what flavor of compositionality our semantic theory enjoys. Indeed, a proper explanation of 1 and 2 might even proceed (in part) with a semantics that is not compositional in any of the ways usually considered. See for example Fine [2007] where he argues that the usual notions of compositionality must be given up. It may even be that natural languages aren’t compositional and 1 and 2 need to be explained by some weaker notion.

⁴⁴ In the first quote I gave from Chomsky, he wrote: “There will be no provision for what Scott Soames calls “the central semantic fact about language, . . . that it is used to represent the world,” because it is not assumed that language is used to represent the world, in the intended sense.” (Chomsky [2000b] p. 132). He is here precisely denying that on his approach sentences represent the world as being a certain way and so impose conditions that a world must meet for it to be true there (i.e. he is denying that on his approach sentences will be assigned truth conditions). In Chomsky [1996] he is even more explicit: “For similar reasons, we cannot assume that statements (let alone sentences) have truth conditions. At most they can have something more complex: ‘truth indications’, in some sense . . . There is good evidence that words have intrinsic properties of sound, form, and meaning; but

also open texture, which allows their meanings to be extended and sharpened in certain ways; and also holistic properties that allow some mutual adjustment.” (pp. 83–84). Pietroski [2003] in defending a Chomsky style internalist approach is equally explicit: “We should be sceptical of the idea that a theory of meaning for a natural language will have theorems that specify the truth-conditions of all the declarative sentences of that language. The successes in semantics suggest that the theoretical action lies elsewhere; semantics is concerned with “internalist” features of linguistic expressions, rather than truth per se. The fact that (an utterance of) a sentence has a certain truth-condition is typically an interaction effect whose determinants include (i) intrinsic properties of the sentence that we can isolate and theorize about, and (ii) a host of facts less amenable to theorizing, like facts about how “reasonable” speakers would use the sentence. If this is correct—if our best semantic theories turn out to be theories of linguistic features that do not determine truth conditions—and meaning is what our best semantic theories are theories of, then the meaning of a sentence doesn’t determine its truth-condition.”

⁴⁵ Chomsky [2000b] pp. 38–42; 130.

⁴⁶ See Chomsky [2000b] pp. 40–42.

⁴⁷ See , 2.

⁴⁸ p.

⁴⁹ Examples renumbered to fit in with the numbering in the present paper.

⁵⁰ Because Pietroski’s challenge was to give a non-*ad hoc* account of the denotations of ‘has wheels’ and ‘have wheels’ that gets the data 7–11 right (as well as other data I am about to discuss), I have done only this and have not given an account of the meanings of ‘have’, ‘has’ and ‘wheels’ that yields the stated semantics for ‘has/have wheels’. However, it would be natural to think that the way to get the VPs to have the denotations described would be to assign two different meanings to the bare plural ‘wheels’: an *at least one* meaning and an *at least two* meaning. However, Zwieg [2009] argues against such an approach, effectively arguing that ‘wheels’ is univocal (having a “number neutral” meaning), as are ‘has wheels’ and ‘have wheels’, and that the “readings” of sentences like 8, 10 and 11 where they require things to have more than one wheel arise from the combination of the semantics of the sentences and a scalar implicature. I don’t want to rule out such an account and I’ll just note that such an account comprises an even stronger response to Pietroski than mine, since such a semantic account is obviously completely non-*ad hoc* and simpler than the approach I’ve outlined. Thanks to Benjamin Spector for discussion

⁵¹ In responding to Pietroski here, I have gone along with his claim that 8 and 11 above are false, so that ‘has/have wheels’ cannot univocally have as its semantic value HWs. Jeff Speaks and Rich Thomason independently noted that there is at least some reason to doubt Pietroski’s claim. If I am addressing an audience and ask ‘Who here has children?’, it seems parents with a single child should raise their hands. Similarly, if I ask you ‘Do you have children?’ you should answer ‘Yes’ if you have a single child. So, Speaks and Thomason thought, it appears that ‘has/have children’ in the questions denote the analogue of HWs. I agree with their judgments here (though I think that when I ask you ‘Do you have children?’ and you have a single child you should perhaps say ‘Well, I have a son/daughter.’), but I am still not convinced that ‘has/have children’ univocally denotes the analogue of HWs. For one thing, should the one-child parents say ‘I have children.’ that now sounds false to me. Similarly, if someone says ‘John has children.’ for one-child parent John. So the questions Speaks and Thomason consider seem to somehow produce noise. But the important point is that if Speaks and Thomason are right, there is an easy response to Pietroski: ‘has/have wheels’ univocally denotes HWs. Hence, in going along with Pietroski’s claim that 8 and 11 are false, we have given him the best argument he can have against the externalist here. And I have argued she can still meet his challenge. A final point in favor of Pietroski’s judgments here. In the recent, highly acclaimed BBC television production of ‘Sherlock’, a contemporary version of Sir Arthur Conan Doyle’s detective stories in which Benedict Cumberbatch plays Holmes and Martin Freeman plays Watson, there is an episode (season 2 episode 2 ‘The Hounds of Baskerville’) in which Holmes tells Watson he doesn’t have friends. Later in the episode, Holmes says the following, conveying to Watson that Watson is indeed his friend: ‘Listen, what I said before John, I meant it. I don’t have friends. I’ve just got one.’

⁵² Kamp [1981]: “Generics, however, are among the most recalcitrant constructions known to me. They will not be treated in this paper.” P. 220.

⁵³ See Nickel [2008] for some illuminating recent work on the topic.

References

- Chomsky, Noam, 2000a, 'Minimalist Inquiries: The Framework' in *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik*, Martin, R., D. Michaels, J. Uriagereka eds., 89–155, MIT Press, Cambridge, MA.
- . 2000b, *New Horizons in the Study of Language and Mind*, Cambridge University Press, Cambridge, UK.
- . 1996, *Powers and Prospects*, Pluto Press, London.
- . 1995, *The Minimalist Program*, MIT Press, Cambridge, MA.
- . 1988, *Lectures on Government and Binding* (fifth edition), Foris Publications, Dordrecht.
- . 1986, *Knowledge of Language: Its Nature, Origin and Use*, Greenwood Publishing Group.
- . 1977, *Essays on Form and Interpretation*, North Holland, New York.
- . 1965, *Aspects of the Theory of Syntax*, MIT Press, Cambridge, MA.
- Elbourne, Paul, 2005, *Situations and Individuals*, MIT Press, Cambridge, MA.
- Fine, Kit, 2007, *Semantic Relationism*, Blackwell Publishing Ltd., Oxford.
- Glanzberg, Michael, 2011, 'More on Operators and Tense', *Analysis* 71.1 112–123.
- Jeshion, Robin, 2015, 'Names not Predicates' in *On Reference*, Bianchi (ed), Oxford University Press, Oxford.
- Kamp, Hans, 1981, 'A Theory of Truth and Semantic Representation' originally published in *Formal Methods in the Study of Language*, Groenendijk, Stokhof and Janssen eds., Mathematische Centrum, University of Amsterdam. Reprinted in *Formal Semantics: The Essential Readings*, 2002, Partee, Portner eds., Blackwell Publishers, Oxford, 189–222. I use the pagination of the latter.
- King, Jeffrey C., 2001, *Complex Demonstratives: A Quantificational Account*, MIT Press, Cambridge, MA.
- Kennedy, Chris and Jason Stanley, 2009, "On 'Average'", *Mind*, 118: 583–646.
- Kripke, Saul, 1980, *Naming and Necessity*, Harvard University Press, Cambridge, MA.
- Levinson, Stephen C., 2000, *Presumptive Meanings: the theory of generalized conversational implicatures*, MIT Press, Cambridge, MA.
- Nickel, Bernhard, 2008, 'Generics and the Ways of Normality', *Linguistics and Philosophy*, v. 31, n. 6, pp. 629–648.
- Ninan, Dilip, 2010, 'Semantics and the Objects of Assertion', *Linguistics and Philosophy*, 33, 355–380.
- Peitroski, Paul, 2003, 'The Character of Natural Language Semantics' in *Epistemology of Language* Alex Barber (ed.), Oxford University Press, New York.
- Putnam, Hilary, 1975, "The meaning of 'meaning'" in *Language, Mind and Knowledge* K. Gunderson ed., University of Minnesota Press, Minneapolis, MN.
- Schwarzschild, Roger, 2002, 'Singleton Indefinites', *Journal of Semantics*, 19, 289–314.
- Yalcin, Seth, 2007, 'Epistemic Modals', *Mind*, Vol. 116, 464, 983–1026.
- Zweig, Eytan, 2009, 'Number-neutral bare plurals and the multiplicity implicature' *Linguistics and Philosophy* 32: 353–407.