The Semantic Lexicon and Analyticity

Ashley Feinsinger

In this paper, I investigate which pieces of information will, according to the lexical semanticist's methodology, count as part of the meaning of a word. This issue is often framed in terms of what information is in the *semantic lexicon*. I first outline Fodor and Lepore's "atomistic view" that no information is semantically associated with a word beyond a specification of the word's denotation. I then evaluate specific cases based in empirical linguistic data against Fodor and Lepore's strategies. I conclude that for some words, it is plausible that their semantic lexical entries are complex. Finally, I offer the beginnings of an account of analyticity based on information in the semantic lexicon and suggest that this might be the only kind of account that could derive analyticity directly from meaning.

Keywords   lexical semantics, the semantic lexicon, lexical atomism, word-meaning, analyticity

Introduction

Philosophers have had a deep rooted interest in the distinction between analytic and synthetic sentences.[1] On many traditional views, this is the distinction between sentences that are true in virtue of meaning alone and sentences that are true in part because of how the world is. Consider sentences (1) and (2) and the intuitive difference between them.

(1) All bachelors are less than fifteen feet tall.

(2) All bachelors are unmarried.

While both sentences are true, we might say that each owes its truth, at least in part, to something different. Both owe their truth partially to the meanings of the words involved: if 'bachelor' had meant *spouse*, then the truth of both (1) and (2) would have been affected. But (1), as opposed to (2), seems to owe its truth in part to the way the world is. Had the world been different, the truth of (1) might have

---

[1]I adopt the view that sentence types are the objects of analyticity. I do not want to say that truths are the objects of analyticity since on my account it is possible that some analyticities are false. Swapping talk of analytic sentences for talk of analytic inferences or analytic beliefs will be fine for the purposes of this paper.

been different. Alternatively, it seems that there is no way the world could be such that the truth of (2) would be affected.

Out of the Kantian, Fregean, and empiricist traditions, analytic sentences were supposed to be necessary, tautological, *a priori*, and true by linguistic convention. But after Quine (1951) any account based in these traditions seemed either untenable or uninteresting.[2] One strand of Quinean attack was to question the notion of a sentence's being "true in virtue of meaning alone". Harman, following Quine, points out that plausibly all sentences are true or false in part because of how the world is:

> "There is an obvious problem in understanding how the truth of a statement can be independent of the way the world is and depend entirely on the meaning of the statement. Why is it not a fact about the world that copper is a metal such that, if this were not a fact, the statement 'copper is a metal' would not express a truth? And why doesn't the truth expressed by 'copper is copper' depend in part on the general fact that everything is self-identical?" (Harman 1999:p.119)[3]

It seems that all sentences owe their truth value to two factors: (1) the meanings of their words and (2) the way the world is. Even the sentence 'All bachelors are unmarried' is made true in part because the world is such that all bachelors are unmarried.

Another strand of Quinean attack was to object to the notions upon which traditional accounts of analyticity were based (e.g., the notions of definition and synonymy)(Quine 1951). Quine claimed that sentences traditionally deemed "analytic" fall into two classes: those like the sentence in (A):

(A) No unmarried man is married.

which remains true under any reinterpretations of 'man' and 'married', and those like the sentence in (B):

(B) No bachelor is married.

which is not a logical truth, but can be turned into one by substituting synonyms for synonyms. But, Quine argues, this way of understanding the analyticity of (B) type sentences rests on our understanding of synonymy, and we do not have a respectable explanation of synonymy. We might try to say that (B) is reducible to (A) by definition, but, as Quine notes, definitions just report previously established synonymies and so cannot serve to explain synonymy (nor analyticity). We might try to say that two expressions are synonymous if they are interchangeable *salva veritate*, but this suggestion only seems plausible in languages rich enough to contain adverbs like "necessarily". But, Quine thinks, to make sense of these adverbs is to suppose that we have already made sense of the notion of analyticity. As Quine puts it:

---

[2]I provide a brief overview of the Quinean attacks on analyticity here mainly so the reader can see that focusing on linguistic meaning might lead to an account of analyticity which survives such attacks. But, since analyticity is not the main point of this paper, I leave a proper treatment of these issues for another discussion.

[3]See also Quine (1976).

"Analyticity at first seemed most naturally definable by appeal to a realm of meanings. On refinement, the appeal to meanings gave way to an appeal to synonymy or definition. But definition turned out to be a will-o'-the-wisp, and synonymy turned out to be best understood only by dint of a prior appeal to analyticity itself. So we are back at the problem of analyticity." (Quine 1951:p.31)

Many have thought that these attacks on the analytic/synthetic distinction show that the traditional distinction is doomed. But as I see it, we are currently in a dilemma: Quine's arguments are compelling but so is the idea that there is an interesting meaning-related distinction between sentences like 'All bachelors are married' and 'All bachelors are less than fifteen feet tall'.[4]

Perhaps the problem has been that previous accounts of analyticity have failed to consider the important advances made in the study of word-meaning by those who work in lexical semantics. Lexical semanticists find ways to study what information is part of word-meaning without relying on the notions of *definition*, *synonymy*, and *understanding a language*. According to the lexical semanticist, if some piece of information about the entities denoted by a word is needed to explain some aspect the word's empirical (syntactic, morphological, or transitional) behavior, then this is taken as evidence that such information is part of the meaning of the word (Asher and Lascarides 2003: section 6.2). The suggestion of this paper is that by using these new ways of figuring out what information counts as part of the meaning of a word, we might have new ways of grounding analyticity understood as relating directly to meaning.

In what follows, I first investigate what information will, according to the lexical semanticist's methodology, count as part of the meaning of a word. This issue is often framed in terms of what information is in the semantic lexicon. In Section 1, I outline Fodor and Lepore's "atomistic view" of the semantic lexicon and the attack they launch on competing views. FL claim that there is no information about the entities denoted by a word that is semantically associated with that word (Fodor and Lepore 1997). In other words, they think that semantic lexical entries are typically "empty". In Section 2, I evaluate specific cases against FL's strategies and conclude that for some cases, it is plausible that some semantic information is associated with the meaning of a word beyond just a specification of the denotation. In other words, sometimes it is plausible that semantic lexical entries are complex. In Section 3, I offer the beginnings of an account of analyticity based on the information in the semantic lexicon and explore the benefits and consequences of such an account.

The view of analyticity that I advocate in this paper is based on the idea that given certain pieces of empirical linguistic data, we have reason to attribute to speakers an understanding of some semantic information. Our analytic beliefs will be derived from these pieces of semantic information; they will be those tacit beliefs about the

---

[4]For some post-Quinean accounts of analyticity, see Boghossian (1996) and Russell (2008). Neither account is particularly satisfying. Boghossian has trouble saying just what inferences are analytic and his account rests on accepting his conceptual role semantics, while Russell's account derives analyticity from what she calls "reference determiners" thus divorcing analyticity from semantics altogether. I leave a full critique of current accounts of analyticity for another discussion.

world that explain certain aspects of our empirical linguistic behavior. It will become clear that in pursuing this new account of analyticity, we will have to give up some of the properties traditionally attributed to analytic sentences. But since the proposed account is based on a direct empirical investigation into word-meanings, it might be the only kind of account that can offer an explanation of analyticity as properly derived from meaning. I take this to be a serious advantage of my account.

1   Lexical Semantics

In their paper, "The Emptiness of the Lexicon", FL propose a view on which semantic lexical entries are typically "atomic" or "denotational". They gloss such a view as follows: The only thing the semantic lexical entry for a word specifies is the denotation of that word. According to this view, the lexical entry for 'dog' specifies only that it refers to dogs (and not, for example, that dogs are animate). Similarly, the lexical entry for the verb 'boil' says only that it refers to boilings (and not, for example, that if x boiled y, then y boiled)(Fodor and Lepore 1997:p.1). Actually, the way FL put their view is a bit unclear. I think it would be better to see the view as claiming that there is something in a speaker's lexical entry for 'dog' that links the word up to the concept, or *mentalese* word, DOG. How the speaker's DOG concept then refers to dogs will involve some story about standing in the right causal relations to the world. It might be more appropriate to think of FL's view of the semantic lexicon as a "null" view.

FL do not give a more rigorous statement of their position, and at times they say things that seem puzzling. For example, towards the end of their paper, they claim that while lexical *meaning* is atomic, lexical *entries* are allowed to be complex (Fodor and Lepore 1997:p.10). For the purposes of this paper, I'll take the lexical atomist to be anyone who denies either that lexical meaning is complex, or that lexical meaning has structure, or that lexical meaning is decomposable into lexical features, or that the lexical meaning of some words is composed in part by the meanings of other words, or that lexical meaning contains any information aside from a specification of what the word denotes (where this specification is not descriptive, but something like the FL picture).

Stated this way, lexical atomism is more of a negative view: it is the denial of these complex views. One could conceptualize the linguistic atomism/non-atomism debate is as follows. Think of a lexical entry as an item in a speaker's mental dictionary, containing conventionalized information about a word. The debate between atomists and non-atomists then is about the nature of this information and whether any of this information has semantic content. Is the lexicon made up of entirely phonological and syntactic information? Or, do speakers conventionalize semantic information as well? I will advocate a non-atomist position in the paper, though I'd like to remain as neutral as possible with respect to the nature and structure of complex meanings. As Cruse points out, representing complex meanings in terms of simpler ones, while perhaps necessary, is extremely problem ridden (Cruse 1986:p.22). In this paper, I make no claims that the components of complex meanings are innate, or primitive, or that they map onto the meanings of other words. While it might turn out that

any non-atomist view must accept some of these claims, I leave this discussion for another paper. Instead, I will focus on whether there is empirical evidence for any such non-atomist picture.[5]

According to FL, there have been no convincing arguments for the claim that some other information, aside from a specification of the denotation, must be semantically associated with a word. For any data purporting to support a complex semantic lexical entry for a word, FL respond in one of the following two ways:

1. The explanation for that data could be grounded in domain information. The explanation for the data could be based in our general conceptual knowledge about the world and not in our knowledge about the language. So, we could explain the data without appealing to any such semantic information.

2. The explanation for the data could be purely syntactic (i.e. not meaning involving). The information explaining the data could be based purely in some syntactic facts about the words involved and not in any semantic facts. So, we could explain the data without appealing to any such semantic information.

It is helpful to consider that there are two principles behind FL's strategies. The first is that the lexicon is supposed to be a repository of purely "linguistic" information. That is, the lexicon is supposed to contain information about words and not information about the world (Fodor and Lepore 1997:p.3). Notice that this principle assumes that there is some interesting distinction between knowledge one has about the world and knowledge one has about one's language. I find this principle at least somewhat controversial, though I will not develop my concerns here. The second is that the lexicon is supposed to make explicit whatever one has to know "on pain of not understanding the word in question" (Fodor and Lepore 1997:p.3). Notice that "understanding a word" might mean being able to use the word in accordance with its conventional behavior, or it might mean knowing what the word means in some public or standardized language. I think FL intend the latter interpretation, though I prefer the former, since it is controversial whether public or standardized languages, such as English, are suitable objects for scientific linguistic investigation (Chomsky 1986) and (Chomsky 2000).

2   Cases

In order to demonstrate where FL's strategies succeed and where they fail, I would like to look at five different cases based in empirical linguistic data. In what follows, I begin with cases for which FL's position is plausible (or possible, even if

---

[5]I am inclined to think about the components of complex meanings as features, and sometimes I'll talk as if the semantic feature + *animate*, for example, has a content much like that expressed by the word "animate". But this fact remains to be established. The main point of this paper is about the existence of such complex meanings, and leaves working out just what the components of these meanings are for a later time. As one might suspect, carrying out this further project is crucial for fully carrying out the analyticity project. For various theoretical positions concerning the components of complex meanings, see Katz and Fodor (1963), Bierwisch (1970), Wierzbicka (1980), Pustejovsky (1995), Jackendoff (1990), and Jackendoff (2002).

undesirable) in order to show the reader how their different strategies are supposed to work. I then move to cases for which FL's position is insufficient.

2.1 Denominal Verbs

Consider (1) and the intuition that it looks semantically odd:

(1)   *Mary buttered the toast with butter.

Similar examples include: 'I salted the pork with salt', 'I painted the house with paint', etc. *Prime facie*, (1) is not ungrammatical but rather semantically redundant. One way to capture (1)'s oddness is to propose a complex meaning for the verb 'butter'. If 'butter' meant *cover with butter*, then this could explain and predict the redundancy. We could capture a generalization by postulating that all verbs with similar empirical behavior encode in their lexical entries certain analogous pieces of information about their objects.

But FL reply with their domain defense, and claim that we could explain the oddness of (1) without assuming that *cover with butter* is part of the meaning of 'butter'. If the semantic lexical entry for 'butter' says only that it refers to butterings, then anyone who knows this and knows something about the world (e.g., what buttering is) can tell that (1) is redundant (Fodor and Lepore 1997:pp.3-4). The sentence could be odd not because of what 'butter' means but because of some domain knowledge that butterings are always done with butter. On such evidence, FL conclude that the fact that sentences like (1) are odd in a special way is no reason to think that the meanings of the verbs involved contain more semantic information than just a specification of their denotation.

While FL's first strategy is compelling here, it comes at a cost. FL's will have to give up on capturing a generalization about a large class of verbs. Their position will require that for each denominal verb, one has to learn a new rule. In order to explain the fact that 'She salted the pork with salt' is redundant, we will have to know that the activity of salting is always done with salt. Of course, domain knowledge could be structured in this way. FL will also need to ackowledge the fact that saltings are not always done with salt (and that corrallings are not always done into corrals, or that paintings are not always done with paint, etc.). One can salt a mine with diamonds and corral horses onto a truck. Thus FL cannot simply say that the reason 'She salted the pork with salt' is redundant is because of some domain knowledge that salting is always done with salt. In summary, FL's domain account might succeed here at some considerable cost, though not in the simple form that they suggest.

2.2 Near Synonyms

The following sets of sentences are supposed to illustrate that verbs 'eat', 'devour', and 'dine' do not have the same empirical distribution although they purportedly pick out very similar events in the world. That is, some have intransitive forms while others do not.

(2)   a.   Mary ate the apple.
       b.   Mary ate.
       c.   Mary ate something.

(3)   a.   Mary devoured the apple.
       b.   *Mary devoured.
       c.   Mary devoured something.

(4)   a.   *Mary dined the apple.
       b.   Mary dined.
       c.   *Mary dined something.

It doesn't seem likely that differences in domain knowledge alone will explain the differences in empirical behavior of these words. For example, it isn't that when we observe eating events, sometimes there is an affected object and somethines there is not, while when we observe dinings, there isn't ever an affected object. Our conceptual knowledge about these activities seems similar with respect to whether or not some affected object is involved. It is thus likely that the information accounting for these differences does not lie in our general conceptual knowledge about these activities.

But FL might say that the differences between the words' empirical behaviors is due to a difference in their syntax, and so, the data in this section does not give us evidence for thinking that the meanings of the words are complex. We could see the semantics of 'eat', 'devour', and 'dine' as simply specifying the denotation and let the syntax generate the empirical differences. In fact, while the verb 'dine' does not select a direct object, like eat does, it can select an affected theme by using the preposition 'on'.

(5)   Mary dined on an apple.

This fact seems to support a purely syntactic analysis of the argument structure variation among the verbs here.[6]

This case is a good illustration of FL's second strategy and of how for some cases, a purely syntactic account is more plausible. While the main point of this example was to demonstrate where FL's syntax strategy succeeds, I'd like to mention briefly that others have gone a different route. Asher and Lascarides make a distinction between conventionalized domain information and unconventionalized domain information, and suggest that the former counts as semantic information. Utilizing this distinction could help explain in an alternate way why it is that 'eat' behaves differently than its near synonyms:

> "To illustrate the distinction between domain knowledge and word meaning, consider the domain knowledge that eating events involve something that is eaten, and contrast this with the knowledge that the agent of the eating event has a digestive tract. There is linguistic evidence that languages like English have conventionalized the domain knowledge that an eating event involves an object that's eaten: one can use this

---

[6]Thank you to Heather Burnett and to an anonymous reviewer for bringing this point to my attention.

information to predict that eat is a transitive verb, since all transitive verbs share the semantic property that they describe events that involve objects that are 'acted on' in some respect; e.g., love, kick, hit, etc. This link between syntax and semantic is therefore encoded in the lexicon." (Asher and Lascarides 2003:pp.253-254)

According to such a picture, the key to explaining why 'eat' behaves differently than its near synonyms is to see that while the domain information for all the verbs is similar, the domain information that has been stored in the lexicon for each verb is different. The lexical entry for the verb 'eat' encodes the fact that 'eat' picks out an activity that involves an object that is "affected" in some way while this fact is not lexicalized for 'dine'. This linguistic fact can explain and predict why one says 'She ate the meal' but not 'She dined the meal'. Furthermore, adopting this strategy lets us capture generalizations between all transitive and intransitive verbs. However, in order for this explanation to be convincing, one would need a way of testing whether or not this information is in fact conventionalized and stored in the lexicon. In summary, pending further investigation into an AL type account, FL's syntax defense prevails here.

2.3 Blacksmiths

FL sometimes complain that they do not understand what it is for an object to be "affected" or "acted on". They worry that no one has ever adequately explained what such semantic properties are supposed to be (Fodor and Lepore 1997:p.4). Absent any good explanation, FL think it is possible that these properties are syntactic and devoid of any semantic content. However, sometimes there are ways of testing whether the relevant properties explaining the data have semantic content.

Consider that in English[7] some compound nouns are unacceptable even when "pragmatics suggests a relationship between the nouns" (Asher and Lascarides 2003:p. 263):

(6)   *blacksmith hammer, *doctor appointment, *dancer shoe

Many compounds with human denoting modifiers (like 'blacksmith', 'doctor', etc.) and non-human denoting heads require possessives in English:

(7)   blacksmith's hammer, doctor's appointment, dancer's shoe

It is interesting that in German these cases do not require the possessive. Instead of having a possessive construction ('doctor's appointment'), German has the noun 'Arzttermin'. For 'blacksmith's hammer', German has 'Schmiedehammer'. As my German speaking colleague describes it, "We don't have the possessive, we just make a new word".[8] The claim here is that English has a rule sensitive to the human/non-human semantic distinction for some set of compound nouns but in German there is

---

[7]All talk of common or public languages can be swapped for talk of I-languages and groups of I-languages.

[8]Thanks to Denis Bühler for these judgements.

no such rule. What explains the fact that 'blacksmith' and 'hammer' can occur in a possessive construction but not in a compound noun construction is that 'blacksmith' denotes something that is human while 'hammer' does not. So, we might conclude that 'blacksmith' semantically encodes that its referents are human.[9]

FL might try to apply their second strategy and say that the property + *human* is in fact a syntactic property having no semantic content. But we could reply with what I call the Productivity Test.[10] It is possible to introduce a word, specify only what it denotes, and predict whether one can form a compound noun with a non-human denoting noun. Suppose I introduce the word 'blurg' and only tell you that it denotes a person who nails molding up inside houses. If you want talk about the nail gun such a person uses, we can predict that 'blurg's gun' will be acceptable but not the compound noun 'blurg gun'. We could argue that the explanation for this result is not entirely syntactic since we did not specify enough about the syntax of the word to generate the possessive. My description of the denotation revealed that the word is a noun, but this is not enough to explain the occurrence of the possessive since many nouns do not require the possessive when combined with non-human heads (e.g., computer chip, table cloth, cup holder, etc.).[11] Since the behavior is predictable from a specification of the denotation alone, this suggests that the explanation is semantic.[12]

FL might try to respond in the following way: Suppose the underlying information responsible for the possessive construction in 'blacksmith's hammer' is that blacksmiths are human. We claimed that the Productivity Test suggested that this information is semantic, i.e. belonging to the meaning of the word. But, FL might say, the results of the Productivity Test are compatible with the claim that there is some syntactic feature (e.g., +*animate*) which is there because of some domain information that blacksmiths are human. It could be that some domain information merely tells us whether 'blacksmith' (or 'blurg') belongs to a certain syntactic category. So, FL might say, this test does not prove that there is any additional information in the semantics of 'blacksmith', since the information that blacksmiths are human might just be domain information giving rise to some syntactic property.[13]

---

[9]As Gabe Greenberg has pointed out to me, in getting to the claim that 'blacksmith' in English, but not in German, encodes that blacksmiths are human, we need the further claim that to satisfy linguistic restrictions, we do not reach into conceptual knowledge. That is, we need to establish something like the following: in forming the possesive construction, English requires that the modifier be human, and since this is a linguistic rule, this information (that certain words denote humans) has to be in the lexicon. I take it that this is a plausible position, though I've provided no argument here. Here is a shot as to how such an argument might begin: Language processing (or at least certain linguistic tasks on par with forming possessive constructions) happens extrememly fast. This would not be possible if speakers, in carrying out such linguistic tasks, had to constantly reach into their vast corpus of conceptual knowledge.

[10]I am indebted to Sam Cumming for the suggestion of this test.

[11]What about 'father/daughter dance'? Also, there is some controversy over whether the semantic information here is +*human* or +*animate*.

[12]For a discussion of a real case involving the introduction of a new word into a language and predicting its empirical behavior, (e.g., 'download') see (Johnson 2004:especially section 3).

[13]We could interpret the 'eat'/'dine' case as follows: since the words denote such similar slices of reality, the domain information is similar. But some syntactic feature for 'eat' (e.g., that it is a transitive verb) derives from domain information that eatings involve an object that is affected. There is no such syntactic feature for 'dine'.

Here is another slightly different way of putting the point. We might use domain information to group words with similar denotations into similar syntactic catagories. But once we do this, it might be that everything about the word's empirical behavior follows solely from its formal properties (along with domain knowledge). That is, the information that blacksmiths are human could be "off-line".[14] It might be that while we use the semantic information to classify the word into its proper syntactic cataory, it plays no role in explaining the rest of the word's empirical behavior. Perhaps knowing that 'blacksmith' is syntacticly +*human* would be enough to explain and predict the word's behavior in the language, removing any need to access the information that blacksmiths are human.[15]

It is worth taking a moment to reflect on this kind of position. This position claims that all we can show is that sometimes, information about the denotation helps categorize words into their proper syntactic categories. This is to say that "semantics" never plays a role in language processing, but can be helpful in the "pre-processing" stages (in the "setting-up" stages). First, we might reflect on how bad this would be for the lexical semanticist. Second, we might try to find cases in which it is plausible that the semantic information is being accessed. As recently pointed out to me, some piece of information will often show up in multiple areas of the lexical entry for a word.[16] We might see the effects of a word's belonging to a certain noun class, for example, in the syntax, phonology, and/or the semantics of that word. So, it is expected that traces of some information will show up in the syntax as well as in the semantics. We could take the challenge now as trying to prove that, although there is likely some syntactic feature corresponding to a given semantic feature, sometimes it is the *semantic* feature which explains the data.

The next example does just this. For the following case, it is plausible that we must appeal to the semantic information to explain the data. That is, for such a case, the semantic information is "on-line". There are many other pieces of linguistic data which seem to suggest that an understanding of some semantic information underlies some aspect of a word's conventional behavior, and such cases might plausibly survive FL's strategies. Some of these examples include noun class membership, restrictions on the dative shift alternation, markedness and the existence of a negative element in one of the words in antonym pairs, gradable adjectives, the mass/count distinction, and the distribution of negative polarity items (NPI's). I will briefly discuss the the mass/count case and the NPI case below. I leave the other cases for areas of further study.

---

[14]I borrow this terminology from Sam Cumming.

[15]Anyway, it might be that FL think that there is no information about the entities denoted by a word that one needs to understand in order to be a competent user of that word on any level, and that this is why there is no information about the denotation in the semantic lexicon. If this is the issue, then FL should not only resist (1) that the information that blacksmiths are human has to be accessed to explain the empirical behavior of 'blacksmith', but also resist (2) that this information is accessed even to catagorize the word 'blacksmith' into an appropriate syntactic category.

[16]Thanks to Jessica Rett for this point.

2.4   *Dimensions and the Mass/Count Distinction*

When combining with a measure phrase, count nouns occur more freely in what has been called the "attributive construction" than their closely related mass nouns. The attributive construction is a nominal construction composed of a measure phrase and a noun, without including 'of'. This is in contrast with the partitive construction which does include 'of' (e.g., compare 'four foot string' to 'four feet of string'). Taken from Schwarzschild (2006), (8) shows attributive constructions with count nouns and (9) shows the infelicitous attributive constructions with the closely related mass nouns (the intendend dimension is enclosed to the right):

(8)   a.   two hour job (duration)
      b.   two hour trip (duration)
      c.   two milliliter drop (volume)
      d.   two pound bean (weight)
      e.   two page poem (page-count)
      f.   two foot curtain (length)

(9)   a.   *two hour work (duration)
      b.   *two hour traveling (duration)
      c.   *two milliliter blood (volume)
      d.   *two pound coffee (weight)
      e.   *two page poetry (page-count)
      f.   *two foot drapery (length)

Why can count nouns occur in these environments while closely related mass nouns cannot?

Schwarzschild (2006) proposes the following analysis. Sometimes a semantic basis for the distinction between count nouns and mass nouns is given in terms of the part-whole structure of the objects the noun applies to. For example, 'furniture' quantifies over stuff which might overlap (i.e. 'furniture' quantifies over things which might have furniture as parts), whereas 'chair' quantifies over things such that no one part is a part of any other (i.e. 'chair' quantifies over things which do not have chairs as parts). We might represent such a semantic restriction on singular count nouns by saying that the domain of a count noun has no proper parts (where $\leq_{\text{Part}}$ represents the part-whole relation):

Semantic Restriction on Single Count Nouns:

$$\forall x \forall y (x \leq_{\text{Part}} y \rightarrow x = y)$$

If we assume such a semantic account for the distinction between mass nouns and count nouns, then, in combination with a certain analysis of the attributive and partitive constructions, we can explain the curious data presented above.

Sometimes a dimension (like weight) tracks the part-whole structure of the objects it applies to and sometimes it does not. For example, weight tracks the part-whole structure of groups of marbles. Take a group of marbles with a certain weight. Add some marbles, the weight goes up, take away some marbles and the weight goes

down. Color, on the other hand, stays the same whether we add marbles or take some away. When a dimension tracks the part-whole structure of the objects it applies to, the dimension is monotonic on the relevant part-whole relation. Schwarzschild gives the following rules for partitive and attributive constructions:

(i) When a measure phrase is combined with a substance noun in the partitive the interpretation is one in which the dimension is monotonic on the relevant part-whole relation in the domain given by the noun.

(ii) When a measure phrase is combined with a substance noun in the attributive the interpretation is one in which the dimension is non-monotonic on the relevant part-whole relation in the domain given by the noun.

If we consider the semantic restriction on count nouns, we remember that the objects in the domain of a count noun have no proper parts. And so, no dimension will be monotonic on the part-whole relation in the domain of the given count noun because there is no relevant part-whole relation. The fact that the noun is a count noun forces that any dimension will be non-monotocnic on the part-whole relation. This can explain why all the examples in (8) are felicitious. Turning to a mass noun (e.g., 'work'), there are likely to be proper subparts in the noun's extension. For example, for each portion of work, there are likely to be proper subparts of this portion that also count as work. Then, for the attributive 'two hour work' to be felicitious, the duration should be immune to the part-whole relation. But duration increases when you add more work, and decreases when you take away some part. Thus, we cannot say 'two hour work' to express duration, but instead must say 'two hours of work'.[17]

This analysis of the data is particularly interesting for our purposes because it suggests that a semantic fact about count nouns is accessed in the understanding of attributive count noun constructions. In case 2.3, we considered the possibility that a semantic fact about blacksmiths was only used to group 'blacksmith' into a certain syntactic category, leaving a syntactic rule to explain the empirical behavior of the word. But in the current example, such a story, though perhaps not impossible, is highly implausible.

Let's consider how such a story might go. Something we know about the denotation of count nouns (e.g., that they have no proper parts) would be used to group the count nouns into a certain syntactic category. Call this the +*count* category. A similar process would take place for mass nouns. Furthermore, this information would not be accessed in explaining any aspect of the words' empirical behaviors. There would be no rules that make use of the fact that the objects in the count nouns' domains lack proper parts. Instead, we would appeal only to syntactic rules. An example of such a rule might be: if the noun is +*count*, omit the 'of', regardless of the measure phrase used, and if the noun is +*mass* do not omit the 'of'.

Prima facie, such a rule might have some difficulty. The very same noun can sometimes occur with 'of' and sometimes without 'of'. 'Two inch rope' and 'two

---

[17]See (Schwarzschild 2006:p. 9) for a brief discussion of how the fact that durational *in* adverbials can combine with accomplishment verbs but not with activity verbs can also be explained in terms of dimensions and monotonicity.

inches of rope' are both felicitious (though they force different interpretations of the dimension referred to). Furthermore, there are cases in which the very same noun and measure phrase are used while the interpretations differ. 'Five feet of snow' can refer to depth or length depending on the discourse context (Schwarzschild 2006:pp. 8-9). Schwarzschild makes the point that pragmatics can affect the relevant part-whole relation and thus affect which dimensions are referred to. Briefly, the idea is that if one is talking about how much snow has fallen, then the relevant parts are layers of snow, and so the interpretation is one which refers to depth (since depth is monotonic on this part-whole relation). But if one is talking about a line of snow as a boundary between two fields, the relevant parts are now linear segments, and so we get a length interpretation. I assume the general point here is that pragmatic considerations interacts with these constructions, and pragmatics can interface with semantics but not with syntax.

All of this is not to suggest that a purely syntactic account of the differing behaviors of count nouns and mass nouns in attributive and partitive constructions is impossible. It is to suggest that the story on which the semantic fact about the domain of count nouns is accessed is more appealing and has a good chance of surviving FL's strategies.

2.5   NPI's

The NPI 'ever' can occur in environments like those in (10) but not in those like (11):

(10)   a.   No person ever ran.
   b.   Few people have ever made it across the finish line.
   c.   John rarely ever wins.

(11)   a.   *Some person ever ran.
   b.   *Many people have ever made it across the finish line.
   c.   *John always ever wins.

We can make a semantic generalization about the environments: in all the examples in (10), the NPI is within the scope of a downward entailing operator, while in (11) the NPI is not. A downward entailing operator is one that reverses the direction of entailment. In unembedded contexts, entailments proceed from sets to their super sets. For example, since the noun 'brother' denotes a subset of the denotation of the noun 'man', a sentence containing 'brother' will usually entail a sentence with the superset 'man'. A downward entailment is one in which superset values entail subset values. For example, 'no' reverses the direction of entailment: 'No man walks' entails 'No brother walks'. The idea is that NPI's are only licensed if they are in the scope of a downward entailing operator. 'Ever' is acceptable in (10) because it occurs within the scope of 'no', 'few', and 'rarely', which are all downward entailing operators. If we saw the semantic lexical entry for NPI's as specifying this information, we could

explain an aspect of their empirical behavior (Ladusaw 2002).[18]

To explain the data, FL could try to say that the fact that 'ever' can only occur in a downward entailing environment is based in some kind of common domain knowledge that speakers have. The present case would then be similar to the 'butter' case. But the knowledge that 'ever' can only occur in downward entailing environments is fairly tacit. If this piece of information is supposed to be a piece of domain information, then FL would have to extend their notion of domain information to include information that is extremely tacit. FL would need to defend such a move.

FL might try to claim that the information at issue is syntactic. Even if one had a syntactic explanation for the distribution of 'ever' (which some linguistics do)[19] the distribution of 'ever' and many other NPI's is highly correlated with a semantic property that is at work elsewhere in the grammar. FL would have to ignore this fact.

To be fair, FL only talk about what linguists would call "lexical" or "open-class" words. These are usually items which acquire new members fairly easily (e.g., nouns and verbs in English). FL do not say anything about "functional" or "closed-class" words, such as prepositions, determiners, and pronouns. NPI's like 'ever' are usually taken to be closed-class items, and so FL might have something else to say here. If FL do intend for their atomism only to apply to open-class items, then NPI licensors might provide a better case against their theory, as such licensors might have the semantic property of entailment reversal.[20] But, as it stands, FL make no explicit distinction between these classes of items and they say nothing as to how their theory will treat closed-class items. In light of this, one might take this section as an invitation for FL to say whether their atomistic theory applies to both classes of words or only to the lexical class, and to justify whichever position they take. As one might anticipate, trying to apply a FL-like atomistic account to the semantics of words like 'ever' is not straightforward, and seems seriously implausible.

3 Analyticity

A first attempt at giving an account of analyticity based on the previous discussion might look something like the following: A sentence is analytic for a speaker if it follows solely from the information in that speaker's semantic lexicon.[21] Below, I would like to outline what I take to be the results of an account of analyticity based on lexical semantics:

1. Whether a sentence is analytic is an empirical matter. It is an empirical question what information explains the conventional behavior of words in a language.

---

[18]It has recently been pointed out to me by Heather Burnett that the distribution and interpretation of NPIs in English is quite complicated, and does not seem to be adequately captured by the Ladusaw generalization evoked here. However, this does not seem to affect the main point which is that it is plausible that something semantic underlies these syntactic patterns.

[19]See Ludlow (1995) for a syntactic account.

[20]Thanks to Heather Burnett and an anonymous reviewer for bringing these points to my attention.

[21]We could restate the condition without talk of semantic lexicons: a sentence is analytic for a speaker if it follows solely from the information about the entities denoted by a word needed to explain aspects of that word's empirical behavior.

2. It is an open question whether many sentences that were traditionally thought to be analytic are in fact analytic in our sense. Whether 'All bachelors are male' is analytic will depend on whether there is any empirical linguistic evidence that the lexical entry for 'bachelor' encodes that its referents are male.[22]

3. Some analytic sentences might be false. Suppose it is actually the case that all the blacksmiths are machines, yet we believe them to be human. Then it will be false that blacksmiths are human, but the sentence 'All blacksmiths are human' might actually be analytic in our I-languages if this information explains the conventional behavior of the word.

4. One might be *a priori* justified in holding that all blacksmiths are human just by having the word 'blacksmith' in one's I-language and being competent with it.

5. This account of analyticity is not based on the traditionally troublesome notions of definition and synonymy, but rather directly on an empirical investigation into the meanings of words. Consequently, we should be able to avoid the Quinean objections.

6. This account also avoids basing analyticity on the notion of "understanding a language" (e.g., being a competent user of some public language). So our account should be able to avoid objections brought against the legitimacy of such notions as well. In contrast, our account connects analyticity directly to the notion of linguistic competence that linguists find respectable (i.e. competence taken as the capacities of a language faculty considered apart from performance errors). Chomsky characterizes this kind of competence as follows:

   "Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of the language in actual performance...We thus make a fundamental distinction between competence (the speaker-hearer's knowledge of his language) and performance (the actual use of language in concrete situations)." (Chomsky 1965)

   In this way, generating a trillion word long well-formed sentence does not break any of the grammatical rules in one's idiolect, even though no person will likely live long enough to utter such a sentence. More explicitly, the competence relevant to our discussion is competence as related to a speaker's particular I-Language (Chomsky 2000). An I-language consists of the set of rules and generative procedures for forming structural descriptions. In this way, you

---

[22]There might be other kinds of conceptual analyticities related to conceptual containment. To establish the existence of these analyticities, we will need to appeal to considerations outside the scope of this paper. It seems that only the account of analyticity proposed in this paper is (or could be) directly about linguistic meaning.

and I might generate the same set of structural descriptions, but have vastly different rules and generative procedures.

Analyticity is often approached by thinking about some other notion of competence, that is, by thinking about what one would have to know to qualify as speaking English well, or by thinking about what one would know if one knew English. In this sense, one is a competent speaker of English if her grammar matches the grammar of English. But on the I-language picture, there is no notion of what the grammar of a particular I-language should be, no notion of how well one's I-language matches some public language, and no notion of "knowing English" or "being a competent speaker of English".

It is controversial whether there is any such thing as the English language, or whether knowing a language and being competent with respect to some standardized language are respectable notions. These issues are empirical matters. Chomsky famously claims that these notions are not scientific and do not properly belong to the empirical study of language (Chomsky 1986:chapter 2) and (Chomsky 2000:pp.48-49, 150). If he is right, then these claims threaten this normative notion of competence and the notions of analyticity that rely on this notion. Our account of linguistic meaning and of analyticity does not require that there are public languages or that there is some normative notion of adequately grasping or understanding a word in such a language.

7. Our account is based on a notion of word-meaning which allows for semantic externalism. Since what an expression refers to is not properly part of the linguistic meaning of the word, we can allow that what one's words refer to (and what contents one's words express) is often determined by factors outside the mind of the speaker.

8. Not all sentences will be analytic and those that are will express substantive claims (e.g., about blacksmiths).

9. We can preserve both the intuition that analytic sentences are derived from the meanings of the words alone and the intuition that no sentence is true solely in virtue of meaning. Since the lexical semanticist's account of meaning allows domain information into the lexicon, there is sense in which analytic sentences derive from meaning and yet also depend in part for their truth or falsity on the way the world is. The analyticities will flow from the lexicon but this fact does not guarantee or ground their truth.

Conclusion

I have tried to show that there are some cases for which FL's lexical atomism is implausible. For such cases, it is plausible that speakers have an understanding of some semantic information beyond just a specification of the words' denotations. One philosophical upshot of the outcome of the lexical atomism debate is the possibility of groudning analyticity in a new way. I began to outline what the tenets of such an account might be and why such an account should be philosophically interesting.

There are many things about the proposed accounts of lexical meaning and analyticity that remain to be worked out. For one, we haven't addressed how the semantic information is encoded in the lexicon. While I might have suggested a view on which the semantic lexicon is comprised of features, such as +*animate* or +*count*, a proper treatement of this issue still remains. Perhaps the most daunting issue is the relationship between the theoretical linguistic entities (e.g., the semantic features such as +*human*) and the contents expressed by words in a speaker's language. Can we express the fact that the semantic feature +*human* is associated with the word 'blacksmith' for a speaker with the sentence 'Blacksmiths are human'? Will the analyticities derived from a speaker's semantic lexicon be expressible in that speaker's language? We might say that the analyticity has "an information content" similar to that of the sentence, 'All blacksmiths are human'.

The discussion in this paper has only scratched the surface of a vast and rich collection of empirical data. But I hope I have demonstrated that as philosophers, we can and should think directly about the empirical linguistic data and that when we do, various philosophical positions (e.g., that semantic lexicons are empty) seem implausible while others (e.g., that there is a viable notion of analyticity grounded in lexical meaning) seem promising. My aim was to suggest that it might be possible to ground analyticity in the lexical semanticist's notion of meaning and that such an account is deserving of further philosophical and linguistic investigation.

Acknowledgements

References

Asher, Nicholas, and Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press.

Bierwisch, M. 1970. On Classifying Semantic Features. In *Progress in Linguistics*, ed. Bierwisch and Heidolph. The Hague: Mouton.

Boghossian, Paul. 1996. Analyticity Reconsidered. *Nous* 30:360–391.

Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. MIT Press.

Chomsky, Noam. 1986. *Knowledge of Language: Its Naure, Origin, and Use*. Praeger Publishers.

Chomsky, Noam. 2000. *New Horizons in the Study of Language and Mind*. Cambridge University Press.

Cruse, D.A. 1986. *Lexical Semantics*. Cambridge University Press.

Fodor, Jerry, and Ernie Lepore. 1997. The Emptiness of the Lexicon: Critical Reflections on J. Pustejovsky's The Generative Lexicon. `http://ruccs.rutgers.edu/tech_rpt/lexicon27.pdf`. Published in *Linguistic Inquiry* 29(2), 1998: 269–288.

Harman, Gilbert. 1999. The Death of Meaning. In *Reasoning, Meaning, and Mind*, 119–137. Oxford: Clarendon Press.

Jackendoff, Ray. 1990. *Semantic Structures*. Cambridge: MIT Press.

Jackendoff, Ray. 2002. *Foundations of Language*. Oxford University Press.

Johnson, Kent. 2004. Tacit Belief, Semantics and Grammar. *Linguistics and Philosophy* 27:57–91.

Katz, J., and Jerry Fodor. 1963. The Structure of a Semantic Theory. *Language* 39:170–210.

Ladusaw, William. 2002. On the Notion Affective in the Analysis of Negative-polarity Items. In *Formal semantics: The essential readings*, ed. Portner and Partee, 457–470. Blackwell Publishers.

Ludlow, Peter. 1995. The Logical Form of Determiners. *Journal of Philosophical Logic* 24:47–69.

Pustejovsky, James. 1995. *The Generative Lexicon*. Cambridge: MIT Press.

Quine, W.V.O. 1951. Two Dogmas of Empiricism. *The Philosophical Review* 60:20–43.

Quine, W.V.O. 1976. Carnap on Logical Truth. In *Ways of paradox and other essays*. Harvard University Press.

Russell, Gillian. 2008. *Truth in Virtue of Meaning: A Defense of the Analytic/Synthetic Distinction*. Oxford University.

Schwarzschild, Roger. 2006. The Role of Dimensions in the Syntax of Noun Phrases. *Syntax* 9:67–110.

Wierzbicka, A. 1980. *Lingua Mentalis: The Semantics of Natural Language*. Sydney: Academic Press.

Affiliation

Ashley Feinsinger
Department of Philosophy
University of California Los Angeles
feinsinger@humnet.ucla.edu