

Open Texture and Schematicity as Arguments for Non-referential Semantics

Christopher Gauker

University of Salzburg

Forthcoming in *Meaning, Context and Methodology*, eds., Sarah-Jane Conrad and Klaus Petrus, de Gruyter Mouton.

Abstract: Many of the terms of our language, such as “jar”, are *open-textured* in the sense that their applicability to novel objects is not entirely determined by their past usage. Many others, such as the verbs “use” and “have”, are *schematic* in the sense that they have only a very general meaning although on any particular occasion of use they denote some more particular relation. The phenomena of open texture and schematicity constitute a sharp challenge to referential semantics, which assumes that every non-logical term has a definite extension. A different, non-referential approach to formal semantics defines truth as relative to a context and defines contexts as built up from exclusively linguistic entities. For any given utterance of a sentence, there will be one of these contexts that *pertains* to it. In this framework, open texture and schematicity can be understood as consequences of the complex nature of the pertaining relation between contexts and utterances.

Keywords: Referential semantics, extension, context-relativity, open texture, schematicity

1. Introduction

According to a commonplace approach to semantic theory, every basic lexical item of a language has an *extension*. The extension of a singular term is an object. The extension of a one-place predicate is a function from objects to truth values. The extension of a two-place predicate is a function that takes an object as input and yields as output a function from objects to truth values. In other cases, the extension may be a function that takes some kind of function as its input and yields some kind of function as its output.¹ The extensions of compound subsentential expressions, such as verb phrases, and the truth values of sentences, are somehow defined on the basis of the extensions of the basic lexical items. For instance, “Socrates is snub-nosed” is a true sentence if and only if, when the extension of “Socrates” (namely, the man Socrates) is the input to the function that is the extension of “is snub-nosed”, the output is truth. Inasmuch as “reference” can refer to the relation between a predicate and its extension, as well as to the relation between a name and its referent, it is commonplace to call this approach to semantics *referential semantics*.

Although I said that the extension of a one-place predicate is a function from objects to truth values, I will, for simplicity, often conflate this function with the set of objects which, when the function is applied to them, yield the value *true*. Although I said that the extension of a two-place predicate is a function from objects to functions from objects to truth values, I will conflate this function with the set of pairs such that the first member yields a function that, when applied to the second member, yields the value *truth*. On this way of talking, “Socrates is snub-nosed” is true if and only if Socrates is a member of the set of snub-nosed things.

¹ For an introduction to this way of thinking about semantic theory, see Heim and Kratzer’s textbook (1998).

An embarrassment for referential semantics is that a lot of terms do not seem to have any definite extension. In particular, there are two kinds of indeterminacy that should concern us. The first kind is what I will call *open texture*. To say that a term is open textured in my sense is to say that whether it applies to a thing is not always determined by the nature of that thing and past usage of the term. For example, the term “jar” is open textured because whether a given container belongs to the extension may be something that has not yet been decided. There may be a kind of glass vessel that could equally well be called a “jar” or a “bottle”. But if a soft drink manufacturer starts selling large quantities of a soft drink in such a vessel and in advertisements describes the beverage as packaged in a “jar”, then no one could object to the claim that “Choco-cooler is sold in jars”.²

The second kind of embarrassing indeterminacy is what I will call *schematicity*. There are many verbs that seem to tell us very little but which we use a lot. Consider, for instance, the verb “to use”. The eater uses a knife, the light bulb uses electricity, the bird uses the air (to fly on), the viewer uses the television (for entertainment), the moon uses the sun (to reflect light). A true utterance of a sentence of the form “X uses Y” does not, in virtue of a meaning that that use of “use” shares with every other use, tell us much of anything about the relation between X and Y. And yet, we will usually take an utterance of a sentence of that form to be telling us something more.

My first task in this paper will be to say something about the expectations that drive us toward referential semantics. Next, I will try harder to show that the kinds of indeterminacies that I have identified above really constitute a decisive reason to

² Ludlow (2014) also appeals to phenomena of open texture to motivate some doubts about referential semantics.

entirely abandon referential semantics. Finally, I will introduce my alternative. The alternative will be a kind of non-referential semantics. The composition of semantic values will be handled by means of a recursive definition of truth in a context. For this purpose, contexts will be a kind of formal structure built up from linguistic entities. As I will explain, the job of grounding truth in the structure of the world will take the form of explaining the conditions under which a context *pertains* to a particular utterance.

A variant on referential semantics takes the extensions of predicates to be not sets but *properties* and *relations*, conceived of as something other than sets. In this paper I will be concerned with the version of referential semantics that takes the extensions of predicates to be sets (or, more precisely, functions). The problems of open texture and schematicity are equally problems for the version that takes the extensions of predicates to be properties and relations, but I will not take the space to prove it.

There are other aspects of natural language lexical semantics that we might struggle to incorporate into the framework of a referential semantics. These include:

- *Indexicality*: The reference of words like “this”, “you” and “now” may vary from situation to situation, and two or more occurrences of any one of them may refer to different things even in a single situation.
- *Intensionality*: In “Dottie is waiting for Santa Claus”, the extension of “is waiting for” cannot very well be a relation between existing objects.
- *Polysemy*: The extension of a word like “book” might shift from a set of concrete particulars to a set of abstract creations (Pustejovsky 1998: 28; Asher 2011: 89).

- *Vagueness*: There is no very sharp dividing line between a movie that counts as “long” and one that counts as “not very long”, or between one that counts as “exciting” and one that counts as “unexciting”.
- *Variable types*: In “She began to read a book”, the object of “begin” is an event, the reading, but in “She began a book”, the object is not an event (cf. Pustejovsky 1998: 115; Asher 2011: 14-18).
- *Permissible syntactic settings*: We can say both “Mary likes reading books” and “Mary enjoys reading books”, but not both “Mary likes to read books” (which is fine) and “Mary enjoys to read books” (which is not fine) (cf. Pustejovsky 1998: 135; Asher 2011: 90).

Each of these phenomena has been addressed in various ways in the literature while adhering to the basic paradigms of referential semantics. I do not assume that they have all been addressed successfully, but in this paper I will not attempt to discredit referential semantics by appeal to them.

2. *Desiderata for a semantic theory*

Among the many tasks for a semantic theory, there are two that a referential semantics seems especially well suited to perform. First, referential semantics seems to provide a means to define the conditions under which a sentence is true. Second, referential semantics seems to provide a means to define logical properties, such as the logical validity of arguments.

As I said at the start, referential semantics aims to define the truth conditions of sentences of the language in terms of the extensions of the terms from which they are composed. Actually, this statement has to be modified to accommodate the context-variability of sentence truth. For instance, the truth of the sentence

“Everyone is present” will vary with a contextually-determined domain of discourse. The truth of “Dumbo is small” will vary with a contextually-determined standard of size. The truth of “Tipper is ready” will vary with a contextually-determined activity for which Tipper may or may not be ready. What we can expect of a sentence, then, is not that it be true simpliciter but only that it be true relative to a context.

A context may be defined as a structure that provides values to parameters that we think of as variable across different situations. For example, domain of discourse and standard of size are two such parameters. The sentences of a natural language will generally have truth values only relative to a context in this sense. However, particular utterances or inscriptions of sentences may have truth values absolutely. There will usually be a single context that we can say *pertains* to the utterance. In terms of the pertaining relation, we can explicate the truth value of utterances in terms of the truth values of sentences relative to contexts, for the truth value of an utterance can be identified with the truth value of the sentence uttered relative to the context that pertains to it.

Underlying the expectation that we can define truth relative to a context in the manner of referential semantics is a more general expectation that we can put in its place if need be. When we speak, we expect our words to satisfy certain norms – unless there is a reason to violate them. What we say should be relevant, informative, well-organized, concise, polite, etc. Among these expectations is that our words be adequately responsive to the way the world is. One kind of responsiveness to the way the world really is is the way that we might characterize as that of *being true* in the manner that referential semantics purports to define. But if it turns out that there are intractable problems with referential semantics, then there is a fallback position,

namely, that our semantic theory should define *somehow* the way in which what we say should be adequately responsive to the way the world is.

The other desideratum that a referential semantics seems to satisfy is the need to define logical properties, such as the logical validity of arguments. From the starting point of referential semantics, we may arrive at definition of logical validity, by generalizing our definition of *truth in a context* to a definition of *truth in a model and a context for that model*. A *model* may be defined as a structure that includes a set of objects, the *universe* for the model, and an *assignment* of an extension to each nonlogical constant. An assignment may be conceived as an alternative reference relation. To accommodate context-relativity, we might build into each model a set of contexts as well. Then we can say that an argument (a sequence of sentences, the last of which is the conclusion) is *logically valid* if and only if, for each model and each context in the set of contexts for the model, if the premises are true in that context in that model, then so is the conclusion.

3. Open texture

In his 1945 contribution to a session of the Aristotelian Society, Friedrich Waismann introduced the term “open texture” into the philosophical lexicon (Waismann 1945). Waismann meant to criticize the idea that all empirically meaningful concepts can be defined in terms of sense data. He had a couple of objections, but one of them was that we can imagine situations in which we would not know whether to apply a term or not and that would challenge the definitions we had entertained up to that point. For example, we might think we could define “gold” in terms of the emission spectrum of gold, but we cannot rule out the possibility of a discovery that would force us to modify our definition (Waismann 1945: 122-123). Consequently, for

many terms, even for a term of science such as “gold”, we cannot through lexical fiat decide in advance how we will use the term no matter what happens. This indeterminacy in the definition of a term is what Waismann called *open texture*.

Referential semantics does *not* commit the error that Waismann was concerned to criticize. We can say that the extension of “chair” is the set of all chairs without committing ourselves to being able to precisely define “chair” using more basic terms. The referential semanticist is committed only to there *being* a definite set of things that is the extension of “chair”. There can be such a set though we have no means of specifying its membership and in no sense *know* which particular objects do in fact belong to it.

Still, an analogous problem, which I will also call the problem of *open texture*, confronts referential semantics. I will say that a term *F* truly applies to an object *o* in a situation *s* if and only if a sentence of the form “*t* is *F*” is true relative to the context that pertains to *s*, where *t* is used to refer to *o*. I will say that a term has *open texture* if and only if the true applicability of a term to an object is not always determined by the nature of the object together with the patterns that can be abstracted from our prior true applications of the term or from the past thoughts of those who have used the term. The *problem of open texture* is the fact that some terms do have open texture in this sense.

That terms have open texture in this sense is also one of the lessons one might take away from Wittgenstein’s discussion of rule-following in the *Investigations* (1953). A different, more radical conclusion that one might also draw from that discussion of rule-following is that the applicability of a term is *never* determined by the patterns of use evident in past uses. (That appears to be how Kripke understands Wittgenstein in Kripke’s book on Wittgenstein (1982).) But here I am not

committing myself to anything so strong. We may grant that past usage determines that the “+” sign denotes specifically the addition function. We may grant that the truth of an application of “cat” to a typical domestic cat and the falsehood of an application of it to a dandelion is strictly determined by past usage (and the nature of the cat and the dandelion). Still, in certain other cases, apart from some lexical fiat, there may be nothing in past usage that decides whether or not a term is truly applicable to a given object.

A study by Malt, Sloman, Gennari, Shi and Yuan (1999) investigated the similarity of container classifications between speakers of American English, speakers of Argentinean Spanish and speakers of Chinese (in Shanghai). They found little correlation between the groupings imposed by the three languages. That is, the fact that two containers received the same label in one language did not make it likely that those same two containers would receive the same label in another language. Moreover, in each language the extensions of the container words did not well reflect the groupings that subjects produced when asked to sort containers on the basis of “physical”, “functional” or “overall” similarity. Most importantly for our purposes, when Malt et al. constructed a multidimensional scaling representing the subjective similarities between the various containers, it turned out that the regions occupied by the extensions of the several container words were not disjoint and the region occupied by the extension of a given container word was often not convex.

Malt, et al. speculate that the extension of a container word may be determined by a variety of processes that have little to do with the overall unity of the category. For example, in a process that Malt et al. call “chaining”, object *a* might be called a “juice box” because it is a cardboard carton containing fruit juice; a cone-shaped cardboard carton *b* might be called a “juice box” because, like *a*, it is designed to

contain a single serving of juice sipped through a straw; a cone-shaped plastic container might be called a “juice box” because it is similar in function to *c*; and so on, to the point where a plastic bottle in the shape of Mickey Mouse is also called a “juice box”. For another example (my own, not theirs), our labeling practices may be driven by marketing considerations. Manufacturers of ice cream might prefer to sell their product in “cartons”, rather than “boxes”, just because we tend to think of what we call “boxes” as rather dirty.

So the decision to label a certain container as a “jar” rather than as a “bottle” may be driven not by a difference in the meanings of the two words but rather by a similarity to particular exemplars, whether or not these are paradigmatic of the kind, and the selection of exemplars may depend on such things as marketing considerations. The container that holds the newly invented choco-cooler is similar enough to some things that have already been called “jar” that it too could be called a jar. It is also dissimilar enough from typical jars – the mouth is a bit too narrow compared to the body – that one could also call it a “bottle” instead of a “jar”. But if the manufacturer, in its advertising, describes the container as a “jar”, not as a “bottle”, and manages to get that usage accepted by consumers, then that’s what that container is – a jar.

This same phenomenon of open texture affects terms for natural kinds such as “cat” as well as artifact kind terms such as “jar”. Consider, for instance, the (fictional) case of the Borneo weasel. Deep in the forests of Borneo, English-speaking explorers come upon a hitherto unknown creature. It looks very much like a weasel and in many ways acts like one. But genetic testing reveals that, actually, its genetic code is precisely that which had been taken to characterize the domestic cat (*Felis catus*). Its physical and behavioral features are in many ways weasel-like but in

some ways cat-like. Its claws retract, but not as far as those of a typical cat. A difference in some neglected portion of a chromosome presumably accounts for the weasel-like appearance. Thus, the explorers face this question: Shall we call it a “cat”?

The Borneo weasel poses a challenge to lexicography, because we have various criteria for something’s being a cat that have never come into conflict in this way before. On the one hand, a thing is a cat if and only if it looks and acts like a typical domestic cat. On the other hand, a thing is a cat if and only if it has the genetic constitution of a cat. By the first criterion, the Borneo weasel is not a cat. By the second, it is a cat. It is not reasonable to insist that the correct classification is determined by prior usage or by the thoughts of people who have used the word “cat”. In using the term “cat”, they may never have had any thoughts that would entail a decision about the Borneo weasel. If we assume that “cat” is a so-called *natural kind* term, then we will say that the genetic characterization has to take precedence over the phenomenological characterization. But we could modify the genetic characterization to exclude the Borneo weasel. And whether “cat” is a natural kind term in this sense is one of the questions that cannot be decisively answered on the basis of past usage or past thinking.

Whether the Borneo weasel ends up being labeled a “cat” or not may depend on accidental features of the situation, by which I mean features of the situation that do not seem to provide an objective basis for judging the real extension of the word. If the discovery had first been reported in a popular magazine, the headline might read, “Bizarrely in Borneo: Cat looks like weasel!”, and in that case the label “cat” might stick. But if the discovery had first been reported in a scientific journal, a scientific classification might be introduced, adding perhaps a new subspecies to the

genus *Felis*, and the term “cat” might simply not come up, in which case other terminology might take hold for describing this animal, leaving “cat” for application to members of the genus *Felis* that look more like a typical cat.

Let us call these situations, where something happens that forces us to decide whether to apply a term to an object, though nothing in past usage or past thoughts decides the matter, a *lexicographical choice point* for the term. In the case of “jar” the question is which other things that have been called “jar” should serve as precedents and which similarities should count. In the case of “cat”, as I have imagined it, the question is which of two criteria that have always been satisfied jointly will guide our usage in a case where only one of the two criteria is satisfied. Terms with open texture are those that may face such lexicographical choice points.

Open texture should not be confused with vagueness. As lexical choice points for a word are encountered and resolved one way or another, the questions they pose may be permanently decided and the meaning of a term is developed and enriched. By contrast, the vagueness of a term is a lasting feature of its meaning and does not indicate a similar place for enrichment. Likewise, open texture should not be confused with context-relativity as this is usually understood in referential semantics. The idea is not that the extension of a term with open texture may vary from one situation of utterance to another; rather, the decisions that are made at lexicographical choice points may govern all usage from that point forward.

Can we plausibly maintain that, all along, even before the decision to call the choco-cooler container a “jar”, the term “jar” had an extension that either included the choco-cooler container, or that all along it had an extension that excluded the choco-cooler container (so that the decision to call it a “jar” was a decision to change the extension)? In the first of these cases, there would have had to be something about

past usage and past thoughts that made it the case that the extension of “jar” included the choco-cooler container. In the second case, something about past usage and past thoughts made it the case that the extension of “jar” excluded the choco-cooler extension. But on the contrary, in deciding whether to call the choco-cooler container a jar, we face a lexicographical choice point, not determined by past usage or thoughts.

If we want to say that even before we encountered the Borneo weasel, the term “cat” had an extension that either excluded or included the Borneo weasel, then there are four options. The extension of “cat” consists of exactly those things that satisfy both criteria (appearance and genotype), exactly those things that satisfy either one or the other criterion, exactly those things that satisfy the genotype criterion, or exactly those things that satisfy the appearance criterion. For any of these options however, we should expect that there is something about past usage or the past thoughts of English speakers that makes it the case that exactly that class of things is the extension of “cat”. On the contrary, since nothing has ever been encountered that satisfies one of the criteria and not the other, and no one has ever thought about whether we should call something a “cat” if it satisfies one but not the other, there is no basis for deciding between these four options. Rather, we face a lexicographical choice point, at which we have to decide (deliberately or not) which of the two criteria our future usage will conform to. So we cannot maintain that prior to the discovery of the Borneo weasel, the extension of “cat” included it or that the extension of “cat” excluded it.

Could we possibly deny my assumption that the extension of “cat” must somehow be determined by past usage or past thoughts of English speakers? Perhaps we could deny it by taking seriously the reality of future events. There has always

been a fact of the matter about what would happen when English speakers first encountered the Borneo weasel even if, in some sense, something else could easily have happened. Suppose that in fact what would happen was that they would withhold the term “cat” from the Borneo weasel; then all along those animals have not belonged to the extension of “cat”. The reason the explorers took that decision was not that they had reason to think the extension of “cat” had all along excluded the Borneo weasel; they did not. They might just as well have decided to apply the term “cat” to the Borneo weasel. So yes, that decision was arbitrary in a certain sense. Their decision resolved what was a lexicographical choice point in the sense that past usage and past thoughts did not determine what they would or should say. Still, we can say that the extension of “cat” had always been that which it was taken to be at the point at which the arbitrary decision was made.

This position cannot be sustained. On the one hand, it is supposed, the Borneo weasel does not, and never has, belonged to the extension of “cat”, because what we will in fact decide, when we first encounter it (deliberately or not) is to not call it a “cat”. Since it does not and never has belonged to the extension of “cat”, an utterance of “It’s a cat”, in reference to the Borneo weasel, would not be true. There may be cases where asserting that something is so makes it the case that it is so (e.g., that two people are married), or even cases where asserting that something is so makes it the case that it *always has been* so (although I cannot think of any examples). But surely it is not the case that asserting that the Borneo weasel is not a cat is what makes it the case that *all along* it has not been a cat.

Could we perhaps say that until a lexicographical choice point for a term is surmounted, the extension of the term is, to some extent, *indeterminate*, and that resolving a lexicographical choice point makes the extension of the term to that extent

more determinate? Before the explorers encountered the Borneo weasel, it was just *indeterminate* whether the Borneo weasel belonged to the extension of “cat”, and afterward it either determinately belonged or determinately did not belong. The problem with this approach is that exacts a revenge. If it is indeterminate whether the Borneo weasel belongs to the extension of “cat”, then that means that the Borneo weasel does not belong to the extension of “cat”. But, as before, that is something we have no basis for saying.

The story of the Borneo weasel is a made-up example, but I take for granted that it is representative of what has sometimes actually happened. More realistic examples might include *fish, mass, water, mercury, slime mold, marriage and death*. Liston (2010) poses a challenge to referential semantics similar to mine by means of the history of “rabbit”, which Old World settlers applied to New World hares. Even a word like “chair” could face a lexicographical choice point, if an ingenious furniture designer designed a sufficiently challenging case. To say that nothing in past usage or past thoughts decides which way to go with such terms is not to say that nothing important hangs on our decisions. Terms such as “person”, “marriage” and “death” can present lexicographical choice points that stand at the center of contentious social policy debates. But the moral questions decided by means of such debates do not necessitate any particular resolution of the lexicographical choice points. For instance, we could decide that a body kept alive in vegetative state does not really have to be *dead* before we can harvest its organs for transplants.

The conclusion I draw is that referential semantics cannot accommodate the open-texture of many of the terms of our language. Referential semantics says that each non-logical term has a definite extension. But in many cases, there can be no possible basis for deciding whether a given object belongs to the extension or not.

Even given the facts concerning the past usage of the term and the thoughts of the users of the term, no supposition about the object's properties warrants any conclusion about whether it belongs. So referential semantics faces an intractable problem and is the wrong way to formulate semantic theories.

4. Schematicity

The English language, as well as presumably every other natural language, contains many verbs of which we might like to say that their meaning is very *schematic*. An example is the verb "to use". We may truly say that a person uses a knife, a commuter uses the subway, a light bulb uses electricity, a bird uses the air (to fly), the moon uses the sun (to reflect light), and so on. Other such schematic verbs of English include: "have", "do", "make" and "get". (Compare "faire" in French, "machen" in German, "Κάνω" in Modern Greek, and "하다" in Korean.) Words in other categories, such as adjectives (e.g., "open"), may be schematic as well, but here I will focus just on verbs.

The fact that the meanings of verbs may be schematic in this way poses a challenge to referential semantics. Apparently, all a referentialist can say about the extension of "to use" is just that it encompasses a great deal. It includes every pair of things such that the first in any way uses the second. (Of course, using is something that happens at a time. I will not worry just now about how to acknowledge that fact in our semantics.) But we cannot leave the matter there, because this answer leaves unexplained the fact that particular utterances of sentences containing "use" will usually communicate something more specific. If someone asks how John gets to work, and the answer is, "John uses the subway", then we learn that John rides to a location near his workplace in the subway train. Or suppose we ask how John paid

his restaurant bill and the answer is, “He used cash”. If “to use” just has the very broad extension, we would have to count this utterance as true even if what actually happened was that John handed the waiter his credit card wrapped in a dollar bill. Unless we can explain in what way a particular occurrence of “uses” refers to some more particular relation, referential semantics will not perform its appointed role in defining our obligation to make sure our words are adequately responsive to the way the world is, as I put it earlier.

In answer to this challenge to explain what “uses” means on any particular occasion of use, one might reply that, while the meaning of “uses” is just very general, on any particular occasion of use the speaker will have some more specific relation in mind and the more particular relation that the speaker communicates by saying “uses” is only that which he or she happens to have in mind. If we say that someone *used* a knife, then typically we will have in mind that the person cut something with it. If we say that the commuter *uses* the subway, then we will have in mind that the commuter rides in the subway train. François Recanati has introduced the term *modulation* to describe, among other things, this process by which an interpreter identifies a more particular relation to take as the meaning of a schematic verb on a particular occasion of use (2004: 134). On his view, the correct interpretation is what it would be reasonable for the interpreter to take the speaker to have intended (Recanati 2013: 82). This answer in effect rescues referential semantics by telling us how, on any particular occasion of use, a particular extension attaches to a word.

One objection to this answer is that it ignores the fact that the primary basis for our knowledge of what other people have in mind is what they *tell* us. If from the fact that the speaker says, “John uses the subway”, we learn that the speaker believes that John rides in the subway train, then that’s because there is something about his

use of those words that allows us to make the inference. We can infer that the speaker believes that John rides the subway to work, because we interpret the speaker's words as meaning that John rides the subway to work. So we need to be able to say what makes it the case that the speaker's words have that meaning on the occasion of his or her uttering them other than that that is the meaning that we can reasonably assume the speaker to have had in mind. Another objection to this answer is that it assumes that thoughts cannot be schematic in just the way that spoken words can be. If mental representations can be schematic as well, though on the occasion of any particular tokening they have a more specific meaning, then obviously we cannot equate that more specific meaning of a thought with the meaning of some other thought that underlies it (Gauker 2012).

A second sort of answer would be to say that the meaning of "uses" is context-dependent. Somehow, the meaning of "uses" and a context together determine the more specific relation, such as *cuts* or *rides*, that an utterance of "uses" will denote in any situation to which the context pertains. However, this appeal to context-relativity in the defense of referential semantics is self-defeating. Context-relativity can legitimately play a role in referential semantics in perhaps two ways. First, it plays a role inasmuch as the truth value of an utterance is supposed to depend on something more than the reference of the referring terms and the meaning of the logical terms and particles (such as case markers) that compose the sentence uttered; it depends as well on the content of the pertinent context. For instance, the context may determine a domain of discourse or a standard of size. Second, context-relativity may play a role inasmuch as the reference of an indexical can be constrained in a rule-governed way by the context. For instance, there might be a defeasible rule according to which

“now” has to be a period of time including the moment of speaking.³ If we now extend the notion of context-relativity to such an extent that even the extension of basic nouns and verbs is relative to context, then it is not so clear that this is a way of defending referential semantics rather than a way of depriving it of all content. If for each token of a so-called referring expression, we now have to consider what it is about the circumstances of tokening that determines that the pertinent context assigns this extension rather than that one, then it is not clear that the appeal to extension serves an explanatory purpose. Rather, the truth value of an utterance will be explained somehow by appeal to these circumstances and their relation to the utterance.

A third possibility (in addition to the idea that “uses” is just very general and the idea that “use” is context-relative) is to say that a verb like “uses” is a kind of pro-verb. Much as a pronoun may *pick up* the reference of a prior proper noun or definite description, so too a pro-verb, such as “uses” may pick up the reference of a prior verb.⁴ One problem with this is just it does not seem to be the case that every understandable use of “uses” is preceded by the use of some other verb whose reference is picked up by that use of “uses”. When not, it might be said, “uses” has a *deictic* meaning, just as a pronoun without a prior anchor in overt speech might refer to an object *deictically*. The problem with this is that it is just not clear what it could mean to speak of deictic reference to a relation, as opposed to deictic reference to an object.

³ In saying this, I am simply bowing to convention. In fact I am skeptical about this second purported role for context (cf. Gauker 2015).

⁴ This answer was suggested to me by Reto Gubelmann.

A fourth possibility is to say that “use” is just *ambiguous*. The problem with this answer is that if “uses” is ambiguous, it is in at most finitely many ways ambiguous. We cannot write a recursive semantic theory for a language containing a term that is in infinitely many ways ambiguous. So for each of the different meanings of “uses”, we should be able to find, or invent, a synonym that has exactly that meaning and no other. For example, the many meanings of “uses” might include the meanings of “cuts”, “rides in”, “drives”, “is powered by”, and so on. But that just goes to show that positing ambiguities does not address the challenge, because each of these synonyms is schematic as well, though to a lesser degree (Searle 1983).

5. *A non-referential semantics*

So far, I have been playing along with the assumption that contexts are structures that, so to speak, supplement models. A context supplies a domain of discourse, which is a set of objects, relative to which we can evaluate quantified sentences. A context supplies a standard of size relative to which we can evaluate the adjective “small” as denoting the set of things that are smaller than that standard. And so on. But there is another way to think about contexts. A context, as I propose to define it, will contain a variety of structures built up entirely from linguistic entities.

To illustrate the basic idea, let me start with a very simple example. Consider a language L having the usual syntax of the languages of first-order logic. Suppose it contains a finite number of one-place predicates, F, G, \dots , and a finite number of two-place predicates, R, S, \dots , a finite number of individual variables, x, y, \dots , and a finite number of singular terms a, b, \dots (We may find it helpful to think of these singular terms as subscripted demonstratives, such as “this₁” and “this₂”.) The logical operators of L are confined to the universal quantifier \forall , the negation symbol \neg , and

the disjunction symbol \vee . For such a language we can define a context as a pair consisting of a *base* and a *domain*. For a context Γ , the base B_Γ is a set of atomic sentences and negations of atomic sentences (literals). The base must be *consistent* in the sense that for no atomic sentence is it the case that both it and its negation is a member. But the base need not be maximal; there may be atomic sentences such that neither they nor their negations are members. (It can even be empty.) For a context Γ , the domain N_Γ is a set of singular terms (not objects such as we might think of singular terms as denoting); it must be nonempty, and it must contain at least every singular term that occurs in any member of B_Γ .

Given these definitions of the language and set of contexts for the language, we can recursively define truth in a context as follows:

- (T0) If $p \in B_\Gamma$, then p is true in Γ .
- (T \neg) If p is false in Γ , then $\neg p$ is true in Γ .
- (T \vee) If p is true in Γ or q is true in Γ , then $(p \vee q)$ is true in Γ .
- (T \forall) If, for every $n \in N_\Gamma$, $\Phi[n/x]$ is true in Γ , then $\forall x\Phi$ is true in Γ .
- (TCI) No other sentence is true in Γ .

- (F0) If $\neg p \in B_\Gamma$, then p is false in Γ .
- (F \neg) If p is true in Γ , then $\neg p$ is false in Γ .
- (F \vee) If p is false in Γ and q is false in Γ , then $(p \vee q)$ is false in Γ .
- (F \forall) If for some n , $\Phi[n/x]$ is false in Γ , then $\forall x\Phi$ is false in Γ .
- (FCI) No other sentence is false in Γ .

Here “ $\Phi[n/x]$ ” stands for the result of substituting the singular term n for the variable x wherever x occurs free in Φ . For example, suppose $\Gamma = \langle B_\Gamma, N_\Gamma \rangle$, where $B_\Gamma = \{Fa, Fc, Gb, \neg Ka, Rab, \neg Hbc\}$, and $N_\Gamma = \{a, b, c\}$. Then the following sentences will all be true in Γ : $Fa, \neg Ka, \forall x(Fx \vee Gx), \neg \forall xHbx$. Evidently, this semantics is three-

valued. There will be sentences and contexts such that the sentence is neither true nor false in the context.

Next, we may define logical validity very simply, as follows: An argument is logically valid if and only if for each context, if the premises are all true in the context, then so is the conclusion. The resulting logic is not classical. For example, sentences of the form $(p \vee \neg p)$ are not valid; they are not true in every context. The rule of universal instantiation (from $\forall x\Phi$ infer $\Phi[n/x]$) is not valid. (That is, arguments having that form are not valid.) I have argued elsewhere (see esp. my 2003) that these results are defensible as a representation of the logic of natural language. It is noteworthy that we can in this way get a defensible definition of logical validity merely by considering the values of the premises and the conclusion in every possible context. We do not need to employ any analogue of the set of models with all their various assignments of extensions to non-logical terms. In particular, we do not need to countenance one particular model that assigns to each nonlogical term the extension it really has. In this respect, this kind of semantics can be characterized as *non-referential*.⁵

The basic idea behind this approach to semantics is that every time we add something of, so to speak, logical interest to the language, we accommodate it by

⁵ This definition of logical validity in effect generalizes on the truth-value semantical definition explored by Hugues Leblanc (1976). An objection to this sort of semantics, with its substitutional interpretation of the quantifiers, is that it seems to validate the *omega rule*, which is an undesirable result. Leblanc shows that this objection can be addressed with modifications in the definition of logical validity, but there is more to say about this issue than I can take up here.

adding structure to our definition of a context and adding clauses to our definition of truth/falsehood in a context. For example, if we wish to add a conditional operator having a logic that approximates to the logic of the natural language indicative conditional, we need to define contexts in such a way that contexts can, in a sense, have other contexts as members. Then we can reformulate the truth conditions for other sorts of sentences. For instance, an atomic sentence will be true in such a context if and only if it is true in every member of the context. A conditional, we can say, is true in such a context Γ if and only if for each context Δ that is either a member of Γ or identical to Γ , if the antecedent is true in Δ , then so is the conclusion. (For details, see my 2005 monograph.)

The definition of truth in a context, so understood, is, in one sense, quite uninformative. It simply defines a relation of truth between two abstract entities – sentences and contexts. However, it can be made to yield an account of the truth values of actual utterances when we add to it an account of the *pertaining* relation between contexts and utterances. For as I said in section 2, an *utterance* of a sentence is true (simpliciter) if and only if the *sentence uttered* is true relative to the context that *pertains* to the utterance. We can hope that our account of the pertaining relation will avoid invoking the reference relation that we are attempting to banish, but we cannot expect our account of it to be a tidy definition.

My hope is that we can identify the context pertinent to a conversation along the following lines: Each interlocutor in a conversation will have his or her own *take* on the context. This we can think of as literally consisting of linguistic items written in the brain. A person's take on the content pertinent to his or her conversation drives his or her behavior in cooperation with other interlocutors. *The* context objectively

pertinent to a conversation can be identified in terms of the *takes* on a context that would be most effective in driving cooperation interaction. In light of this general conception, we can identify a variety of desiderata concerning the interlocutors and their environment that the context pertinent to a conversation ought to satisfy; the context that really does pertain to a conversation will be that which provides the best overall fit to these desiderata. (For further discussion, see my 2005.)

As I have already explained, in these terms we can define logical validity, and so this kind of semantic theory satisfies our requirement that our semantics be able to do that. It also satisfies our requirement that it defines a way in which what we say should be adequately responsive to the way the world is. The definition of truth in a context does not do this, for as I said, that is just a relation between two abstract entities. But when we add the relation of pertaining between a context and an utterance, we get what we require: The normative expectation that we can define in terms of this semantics is that, other things being equal, the sentences we utter should be true relative to the context that *pertains* to the situation in which our discourse takes place.

6. *Open texture and schematicity revisited.*

A word has what I call *open texture* when its true applicability to an object in a given case is not always determined by the nature of the object, past uses of the term or the thoughts of speakers. When that is so, our use of the term may face a *lexicographical choice point*, at which point the question whether to apply the term to an object or withhold application can only be answered by a more or less arbitrary decision. That words have open texture in this sense is an embarrassment for referential semantics,

because referential semantics demands that terms such as “jar” and “cat” have a definite extension.

The non-referential approach to semantics described in the previous section faces no such embarrassment. Because it posits no extensions, no question can arise about whether the choco-cooler container belongs to the extension of “jar”. Likewise, no question can arise about whether the Borneo weasel belongs to the extension of “cat”. In any situation in which the question arises whether to utter an atomic sentence of the form “*x* is a cat”, the question, “Does that sentence belong to the base of the context that pertains to our situation?” will still have to be answered (though not necessarily in just those terms). That, as I have already acknowledged, may be a messy, difficult question. In the “cat” case, part of what is involved will be thinking about similarities between the present objects and things we have applied “cat” to in the past. And part of what is involved will be thinking about what we expect to achieve by calling something a “cat”. But no part of the difficulty is deciding whether the referent of the singular terms belongs to the extension of “cat”. We can allow that a factor that may figure into the answer may be just the sort of arbitrary decision that I have called a lexicographical choice point.

What I here call *schematicity* is a feature of verbs that seem to have only a very general meaning in the abstract, but which seem to denote more specific relations on any particular occasion of use. Schematicity is an embarrassment for referential semantics, because referential semantics does not seem to be able to offer any reasonable account of the relation between the general meaning of a schematic verb and the particular relation it denotes on any particular occasion of use.

The non-referential approach to semantics faces no such embarrassment. Let us say that a context in which a sentence of the form “*X* uses *Y*” is true is a *uses-*

context. Even on the non-referential approach to semantics one can acknowledge that “uses” is schematic *in the sense* that the features of situations that make it the case that a *uses*-context *pertains* to a situation may vary widely. In one situation it may be the fact that something *cuts* something in that situation, in another that something *rides in* something in that situation, etc. That the pertinent feature of a situation that determines that a *uses*-context *pertains* to the situation may vary in this way from situation to situation is in no way contrary to expectations, given that, in general, the relation between a context and the situation that it *pertains* to may be a complex one, involving a competition of factors and amounting in the end to an optimality of fit.⁶

References

- Asher, Nicholas. 2011. *Lexical meaning in context: a web of words*. Cambridge: Cambridge University Press.
- Gauker, Christopher. 2003. *Words without meaning*. Cambridge, MA: MIT Press.
- Gauker, Christopher. 2005. *Conditionals in context*. Cambridge, MA: MIT Press.
- Gauker, Christopher. 2012. Inexplicit Thoughts. In Laurence Goldstein (ed.), *Brevity*, 74–90. Oxford: Oxford University Press.
- Gauker, Christopher. 2015. How many bare demonstratives are there in English? *Linguistics and Philosophy* 37. 291–314.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Oxford: Blackwell.
- Kripke, Saul. 1982. *Wittgenstein on rules and private language: An elementary*

⁶ I thank audiences in Tartu, Zürich and Konstanz for stimulating questions that led to many improvements. Questions from Tobias Henschen, in Konstanz, and Matej Drobnák, during his visit to Salzburg, led to significant improvements.

- exposition*. Cambridge, MA: Harvard University Press.
- Leblanc, Hugues. 1976. *Truth-value semantics*. New York: North Holland Publishing Company.
- Liston, Michael. 2010. Rabbits astray and significance awandering: Review essay on Mark Wilson's *Wandering significance*. *Philosophy and Phenomenological Research* 89. 809–817.
- Ludlow, Peter. 2014. *Living words: Meaning underdetermination and the dynamic lexicon*. Oxford: Oxford University Press.
- Malt, Barbara C., Steven A. Sloman, Silvia Gennari, Meiyi Shi & Yuan Wang. 1999. Knowing versus naming, Similarity and the linguistic categorization of artifacts. *Journal of Memory and Language* 40. 230–262.
- Pustejovsky, James. 1995. *The Generative Lexicon*. Cambridge, MA: MIT Press.
- Recanati, François. 2004. *Literal Meaning*. Cambridge, MA: Cambridge University Press.
- Recanati, François. 2013. Reply to Gauker. *Teorema* 32. 81–84.
- Searle, John R. 1983. *Intentionality: An essay in the philosophy of mind*. Cambridge, MA: Cambridge University Press.
- Waismann, Friedrich. 1945. Symposium: Verifiability. *Proceedings of the Aristotelian Society: Supplementary Volumes* 19. 119–150.
- Wittgenstein, Ludwig. 1953. *Philosophical investigations*. Macmillan.