

Literal Meaning

John-Michael Kuczynski, PhD

To avoid a copyright-violation, I have uploaded this article, instead of the similarly-titled article published in *Communication & Cognition*. Even though there is not so much as a single sentence that occurs in both papers, those papers are virtually content-identical. This paper has never been published.

1.0 The meaning of “meaning”

Meaning #1: The evidential meaning of “meaning”

Meaning #2: The psychological meaning of “meaning”

Meaning #3: The linguistic meaning of “meaning”

1.2 Why Meaning #3 \neq Meaning #2

1.3 Why Meaning #3 \neq Meaning #1

2.0 Sentences as proposition-isomorphs

2.1 Propositions as digital structures

3.0 The three branches of the philosophy of language: syntax, semantics, and pragmatics

3.1. Semantics

3.2 Pragmatics

3.3 Syntax

4.0 The need for the discipline of semantics

4.1 Semantics needed to figure out what is literally meant and what is not

4.2 Semantics needed to figure out what is literally meant and what is not (continued)

5.0 The Nature of Semantic rules

5.1. An important subtlety

5.2 How meaning is assigned to hitherto meaningless and, therefore non-linguistic, entities

5.3 What are semantic rules?

- 5.3.1 Why semantic rules are not functions
- 5.3.2 The Gricean approach
- 5.3.3 More problems with the Gricean approach
- 5.3.4 How some Griceans deal with the problem just described
- 5.3.5 Why Grice's theory is inconsistent with the normative nature of semantic rules
- 5.4 The psychological reality of semantic rules
- 6.0 What is literal meaning?
- 7.0 Tokens vs. types: some preliminary terminological points
- 7.1 Two-dimensional semantics
 - 7.1.1 Demonstratives vs. indexicals
 - 7.1.2 Indexicals (continued)
- 7.2 Definite descriptions
 - 7.2.1 Incomplete definite descriptions
- 7.3 The expressive limitations of indexical-free languages
- 7.4 Tokens, types, and context-sensitivity
- 7.5 Ambiguity and context-sensitivity (revisited) and the type-token distinction (revisited)
- 7.6 Do any expressions have one-dimensional semantics?
- 7.7 Ambiguity and context-sensitivity (re-revisited) and the type-token distinction (re-revisited)
- 8.1 Contextual definition
- 8.2 Contextual definition: its scope and limits
- 8.3 Frege's generalization of the concept of a function

1.0 The meaning of "meaning"

There would be no languages if there were no expressions (words, phrases, sentences, etc.). Nothing meaningless is an expression. For this reason, the concept of an expression must be understood in terms of the concept of meaning, the same therefore being true of the concept of language.

But it isn't much use to be told that words and sentences "have meanings," since the word "meaning" has three different meanings, and only one of these directly relates to the nature of language.

Meaning #1: The evidential meaning of “meaning”

In some cases, to say that x “means” y is to say that x is *evidence* of y—that x and y are causally interrelated in such a way that, given x, it can reasonably be inferred that y. “Smith’s hacking cough means that he has a violent lung infection” means “Smith’s hacking cough is evidence that he has a violent lung infection.” And the latter means that coughs like Smith’s are *causally connected* to violent lung infections in such a way that it may reasonably be inferred that Smith has a violent lung infection.

Smith’s violent lung infection is a *cause* of Smith’s hacking cough. But for x to be evidence of y, it is neither necessary nor sufficient that y cause x.

Why isn’t it *necessary*? First of all, causes can be evidence of their own effects. (Bill’s current drunkenness is evidence of a poor performance on his upcoming economics exam.) Second, if some event or state of affairs z is a cause of both x and y, then x can be evidence of y without being a cause or an effect of y. (Suppose that Bill is slurring his words. This is evidence that he’ll do poorly on the upcoming test. But his slurring his words is neither a cause, nor an effect, of his substandard test-performance. His drunkenness is a cause of (a) his slurred speech and (b) his imminent, substandard test-performance. And it’s because his slurred speech has a causal ancestor in common with his poor test-performance that the former is evidence of the latter.)

Why isn’t it *sufficient*? Given only that the cause of Bill’s failing was that he was drunk, we can’t infer from the fact that he failed that he was drunk. There are many reasons why a person may fail a test. Supposing that y caused x to occur, x is evidence of y only if it can’t reasonably be supposed that anything other than y was the cause. (Only a violent lung infection could be responsible for Smith’s hacking cough, that being why the latter is evidence of the former.)

Meaning #2: The psychological meaning of “meaning”

When we use sentences, we mean things by them. Meaning in this sense is a psychological notion. You tell me that Sally is the most wonderful, decent person you’ve ever known. I respond by saying “things aren’t always as they seem.” What I mean is that Sally is devious. In other words, my intention in making this statement is to say that Sally is devious. Given that intentions are psychological entities, meaning in this sense is obviously a psychological notion.

Meaning #3: The linguistic meaning of “meaning”

The sentence “snow is white” says something about something; it attributes the property of being white to snow. Therefore, it *means* that snow is white. This kind

of meaning is in a class by itself; it isn't identical with either of the two kinds mentioned so far; and there isn't any obvious way to understand it in terms of them. Let us now say why.

Meaning in the psychological sense *involves*, but does not coincide with, meaning in the linguistic sense. Once again suppose that, in response to your telling me that Sally is a wonderful person, I say "things aren't always as they seem." My meaning—what I'm trying to get across—is that Sally is devious. But in my attempt to get this across, I'm taking advantage of the fact that "things aren't always as they seem" has an existing (linguistic) meaning. My meaning that Sally is devious is parasitic on my utterance's meaning that things aren't always as they seem. Meaning in the psychological sense is therefore parasitic on meaning in the linguistic sense.

1.2 Why Meaning #3 ≠ Meaning #2

Some philosophers and linguists have held that for:

(1) *"Snow is white"*

to mean that snow is white is for it to be the case that, in uttering "snow is white," what *people* mean is that snow is white. This view, duly generalized, is that for a sentence *S* to have meaning *M* is for it to be the case that, in uttering *S*, people to mean *M*.¹

This position is false. There are many sentences that have determinate meanings even though they've never been uttered before and, therefore, no one as of yet has ever meant *anything* by them. The sentence:

(2) *"The cube root of three is Sir Lawrence Olivier's favorite irrational number between one and four"*

has a determinate meaning, even though that sentence never probably has been uttered. Thus, meaning in the linguistic sense is not in *all* cases identical with meaning in the psychological sense.

But a stronger point is warranted. Let's say that P_1 and P_2 are the propositions meant by (1) and (2), respectively. In saying that (1) means P_1 and that (2) means P_2 , we are not using the word "means" equivocally. Both occurrences of "means" in the last sentence denote the same relationship. Thus, the relationship that (1) bears to its meaning is the same as the relationship that (2) bears to *its* meaning. Given that, as we saw, (2)'s having P_2 for its meaning isn't identical with P_2 's being what people mean in uttering (2), it follows that (1)'s having P_1 for its meaning isn't identical with P_1 's being what people mean in uttering (1). Of course, what we just said about (1) and (2) can be said of any other sentence. So even if what people mean in uttering a given sentence

happens to coincide with its literal meaning, what it is for sentence S to have proposition P for its literal meaning isn't for people to mean P in uttering S.

Psychological meaning presupposes linguistic meaning. What a person means when uttering a given sentence is a function of, among other things, his beliefs as to what that sentence *already* means. You must believe that "snow is white" means *snow is white* if, intending to speak sincerely and literally, you say "snow is white" with the intention of getting it across that snow is white. If you think that "snow is white" means *grass is green*, you cannot, if your intention is to speak sincerely and literally, believe that "snow is white" means *snow is white*.

Of course, you could know full well what "snow is white" in fact means, but use that sentence to get across something that has nothing to do with the color of snow. Knowing what "snow is white" actually means, you might utter that sentence with the intention of getting it across that the government is controlling our thoughts with alpha waves. And, depending on the circumstances, that could be precisely what an utterance of that sentence would convey.

But whatever the message is that, in uttering a certain sentence, you wish to convey, you must believe that message to have *some* kind of relationship to the one meant by that sentence itself. Furthermore, if you are to *succeed* in saying what it is you wish to say, what you believe to be meant by the sentence you are using must be *right*. If, intending to speak sincerely and literally, you say "snow is white," thinking that it means *bananas are yellow*, you will fail to say what you wanted to say.

Thus, setting aside defective utterances, one cannot, in uttering a given sentence, mean *anything* by it unless one knows what it *already* means. So meaning in the *psychological* sense is parasitic on meaning in the linguistic sense, and the two kinds of meaning are therefore entirely distinct.

1.3 Why Meaning #3 ≠ Meaning #1

The sense in which "snow is white" means that snow is white isn't comparable to the sense in which smoke means fire. The fact that smoke means fire has nothing to do with conventions on the part of human beings.² But the fact that "snow is white" means what it does *is*, at least in part, a matter of convention. It's a matter of convention that "snow" doesn't refer to grass and, therefore, that "snow is white" doesn't mean that grass is white; it's a matter of convention that "white" doesn't mean *green* and, therefore, that "snow is white" doesn't mean that snow is green.

Although the whiteness of snow sometimes causes people to say "snow is white," it doesn't do so in the way that fire causes there to be smoke. Fire happens; smoke happens as a result. The presence of smoke doesn't embody any *judgment* about anything. But, when caused by the whiteness of snow, utterances of "snow is white" *do* embody judgments of various kinds. People see or otherwise come to

believe that snow is white; and, since they know the relevant linguistic rules, they know that, were they to say, “snow is white,” they judge that they’d be making a *correct* statement. Thus, utterances of “snow is white” embody judgments about the color of snow and about how, linguistic conventions being what they are, one can report the color of snow. Also, people don’t say everything that occurs to them. Before deciding to utter a given sentence, people typically make context-based judgments about the appropriateness of uttering that sentence. So various judgments—about snow, about language, and about human psychology—are involved in the causal connection between the whiteness of snow and a given person’s saying “snow is white.” There is thus a *normative* dimension to language use that is absent where purely natural, non-conventional cause-effect relations are concerned.

2.0 Sentences as proposition-isomorphs

The meaning of a true or false sentence is a *proposition*. Propositions are not themselves sentences. That is why different sentences (e.g., “schnee ist weiss” and “snow is white”) can express the same proposition.

Propositions, when true, are *truths*. Thus, propositions have existed as long as there have been truths; which means that they’ve existed as long as there has been anything and, consequently, that propositions are not creations of human creations.

Though distinct from the sentences that express them, propositions are structurally similar to them. Two otherwise dissimilar sentences can share the word “John.” “John loves Mary” and “Sally punched John” are two such sentences. The meanings of those sentences obviously have something in common corresponding to the fact that they share the word “John.” Since they share no other constituents, the thing meant by “John” must be capable of moving *on its own* from sentence-meaning to sentence-meaning. This would not be the case if the thing meant by “John” in the proposition meant by “John loves Mary” were incapable of being disengaged from the things meant by “loves” and “Mary.” It follows that propositions consist of discrete parts; it also follows that those discrete parts correspond to discrete parts of the sentences that express them. Taken together, these two points entail that sentences are structurally like the propositions they express.

2.1 Propositions as digital structures

Given that propositions consist of discrete, isolable entities, it follows that, like sentences and unlike visual perceptions and photographs, propositions are *digital* structures. The sentence “Sally punched Bob” has a unique decomposition into a certain “minimal units of significance,” or “morphemes,” these being “Sally,” “Bob,” etc. Given what we said in Section 2.0, it follows that something similar is true of the

corresponding proposition. Sentences and propositions are *digital* structures, meaning that they have unique breakdowns into minimal significant units.

A visual perception of Sally punching Bob doesn't have a structure comparable to that of "Sally punched Bob" or any other sentence. Unlike sentences, perceptions don't have to decompose into minimal significant units. Visual perceptions, unlike sentences and propositions, therefore have a non-digital or *analogue* structures. Given that at least some thought involves the processing of perceptual information, it follows that thought at least sometimes has a structure very different from language. (See Section 5.4 for further discussion of this.)

3.0 The three branches of the philosophy of language: syntax, semantics, and pragmatics

The study of language is typically divided into three sub-disciplines—semantics, syntax, and pragmatics.³ In addition to denoting a branch of linguistic study, each of these three words denotes dimension of language. So "semantics" refers to a certain discipline and also to a feature of expressions, the same being true of the other two expressions.

3.1. Semantics

The discipline of semantics attempts to make it clear what our utterances literally mean. It has no interest in what is conveyed through suggestion or innuendo.

If a disappointed boss says to a substandard employee, "it might not be a bad idea for you to start thinking about finding a new position," the literal meaning of his utterance is quite innocuous. But the message that is being *sent* is not innocuous—that message is: *you're fired; you're a disgrace; go away*; etc. The utterance's *semantic* coincides with its literal (innocuous) meaning.

3.2 Pragmatics

Pragmatics studies the *use* of language. Sometimes language is used literally. Asked whether I'm over thirty years of age, I say "yes, I'm over thirty years of age." What I mean coincides with what my utterance literally means.

Language is often used non-literally. If, while addressing a pan-handler, I say "you've made a fine life for yourself," what I mean is the antitheses of what my utterance means. But usually the propositions literally meant by our utterances are neither opposed to, nor exhaustive of, the propositions we wish to affirm in producing those utterances. Asked whether there's a place to get food, I say "there's a McDonalds down the road." The proposition I'm affirming is: there is a nearby place to get food, the reason being that there's a McDonalds down the road. Thus, the

proposition literally meant by my utterance is only a part of what it is that I'm saying. Thus, what a sentence literally means is only one of many factors governing what it is used to mean. The discipline of pragmatics tries to identify the remaining factors.

3.3 Syntax

The disciplines of syntax studies the structures of the meanings of *complex expressions*. A complex expression is one that consists of other expressions. (Thus, "the man who ate my cookie" is a complex expression, since it consists of "man," "ate," etc., each of which is meaningful. By contrast, "red" is not a complex expression, since it doesn't have any meaningful proper parts.) The discipline of syntax tries to make it clear how the meanings of complex expressions depend on those of their parts.

Thus, the discipline of syntax doesn't study the meanings of complex expressions *per se*. It studies the *relationships* that such meanings have to those of their constituents. Consider the sentence "Sally hates Bob." The word "hates" *occurs* in that sentence. If that occurrence is replaced with an occurrence of "loves" or "is amused by," the resulting sentence has a very different meaning from the first. This shows that what "Sally hates Bob" means depends on what "hates" means.

Bearing this point in mind, consider the sentence "Larry loves Julie." Obviously this sentence doesn't mean the same thing as "Sally hates Bob." But the *relationship* borne by the meaning of "Sally hates Bob" to the meaning "hates" is identical with the relationship borne by the meaning of "Larry loves Julie" to that of "loves." Exactly similar points hold in connection with each of the remaining two constituents of each of those sentences.

The discipline of syntax studies the *relation* that the meanings of complex expressions bear to the meanings of the simple expressions composing them. Thus, syntax doesn't study the semantics (meanings) of complex expressions. It studies the *structures* of the semantics of complex expressions. Syntax studies semantic structure.

4.0 The need for the discipline of semantics

Even though we all know what is meant by:

(1) "John wants to catch a 20-pound striped bass,"

we don't know what it is that we know in knowing this. Semanticists supply us with the missing metaknowledge. Consider the sentence:

(2) "John wants to punch Bob."

(2) Attributes a certain property to John—that of wanting to punch Bob. Given that (1) and (2) are grammatically isomorphic, it's natural to assume that there exists some 20-pound striped bass x such that the proposition expressed by (1) is:

(3) John wants to catch x .

But this isn't the right analysis. There isn't some *one* fish such that, if the desire ascribed to John by (1) is to be gratified, John must catch *that* very fish. There is thus no fish x such that, if (1) is to be true, John must want to catch x .

The meaning of (1) is:

(1_R) John wants it to be the case that: there exists some fish x such that x is 20-pound striped bass and such that John catches x .

Thus, (1) doesn't describe a relationship between John and some non-existent or quasi-existent fish. It affirms the existence of a relationship between John and a *proposition*. The proposition in question is one that, in English, is expressed by the sentence:

(4) "There exists some fish x such that x is a 20-pound striped bass and such that John catches x ."

If John's fishing-trip is a success, that proposition will be true; otherwise it will be false. But that proposition exists either way.

But we still haven't solved the problem. In (4), the expression "some fish" occurs. Grammatically, that expression is a noun-phrase. But, unlike other noun-phrases, it doesn't refer to anything. ("Some fish" doesn't refer to some fish. There is no fish x such that "some fish" refers to x . That's why, given any particular fish F , if you say "some fish is wet, but x is not," what you are saying isn't self-contradictory.) So the problem we were trying to solve remains.

But to solve the problem, we need only reword (4). The needed rewording is this:

(5) The property of being a 20-pound striped bass that John catches is instantiated.

The property of being such a fish exists. So (1), which seemed to be about a non-existent fish, is about an existent property. (5) says of that property that it's instantiated. Thus, a complete analysis of (1) is given by:

(1_{CA}) John wants it to be the case that the property of being a 20-pound striped bass that John catches is instantiated.

So even though just about every English speaker understands (1), knowing what it is that one knows in understanding it isn't such a trivial thing.

4.1 Semantics needed to figure out what is literally meant and what is not

Despite everything just said, there is clearly a sense in which every English speaker knows what (1) means. What the semanticist is doing in connection with (1) isn't comparable to what you (who, we'll assume, speak Spanish) are doing in connection with it when you tell a monolingual Spanish speaking friend of yours what it means. The semanticist is needed to clarify the structure of the meaning that (1) is already known to have, but he isn't needed to *identify* that meaning. The semanticist isn't a translator. But there are many cases where the semanticist *is* needed to identify literal meaning. In fact, as paradoxical as it may sound, there are cases where he is needed to identify the meanings of sentences *that are perfectly well understood*.

First of all, we must distinguish what is literally meant by an utterance from what it is that the speaker wishes to convey. To give a trivial example: You and I are robbing a bank. I yell: "the cops are coming!" What I wish to convey is that we should hurry up. In this particular case, it's easy to distinguish what is literally meant from what is non-literally suggested, and semantics would therefore have no interest in it. But in other cases, it's exceedingly hard to do this, and it's with these other cases that semantics is concerned.

A story will help us move forward. Somebody who is wearing a ski-mask, and who I therefore don't recognize, deftly snatches my pocket from my wallet. As he's running off, I point at him and yell: "that man is a thief!" Let U_1 be this utterance.

Before moving on, let's take a moment to make it clear what U_1 's literal meaning is. Somebody just stole my wallet. I don't know who that person is. But whoever it is, I am attributing a certain property to him. If that person has that property, I have spoken truly; if not, not. U_1 is correct if, and only if, the person referred to by "that man" has the property expressed by "is a thief." Thus, there is some individual x such that x has just stolen my wallet and such that what I've just said is true exactly if x is a thief. (The underlined part is U_1 's literal meaning.)

The next day, my lovable office-mate Steve eats one of the cupcakes that was on my desk. I jokingly point at him and say: "that man is a thief." Let U_2 be *this* utterance. There is some x such that x just ate my cupcake and such that U_2 is true exactly if: x is a thief.

Unbeknownst to me, Steve is the pick-pocket, and there is some individual x , namely Steve, such that each of U_1 and U_2 is true if and only if x is a thief. Thus, U_1 and U_2 have the very same literal meanings.⁴ But I don't know this, even though I speak

English perfectly and, on each occasion, obviously understand perfectly well what it is that I'm saying.

How this is possible?⁵ Our sense-perceptions describe things. My uttering U_1 was a response to my being given a visual description of Steve. That description was to the effect that:

(i) There is some man x such that x is wearing a ski-mask and such that x is running off into the distance.

But U_1 's literal meaning is not that that effect. There is some man x such that x is a wearing a ski-mask (etc.), such that in uttering U_1 I was saying that:

(ii) x is a thief.

The meaning of U_1 , being identical with (ii), is quite threadbare. But I grasped that threadbare meaning through my descriptively rich visual perception, whose content is given by (i).

My uttering U_2 was a response to my being given a different description of Steve. That description was to the effect that:

(iii) there is some man x such that x is a portly amicable fellow who is sitting over in that chair.

But U_2 's literal meaning is not that that effect. There is some man x such that x is a portly amicable fellow (etc.) such that in uttering U_2 I was saying that:

(ii) x is a thief.

Echoing what we said a moment ago, the meaning of U_2 , being identical with (ii), is quite threadbare. But I *grasped* that threadbare meaning through my descriptively rich visual perception, whose content is given by (iii).

Because I grasped (ii) by way of different bodies of perceptual (descriptive) information, I didn't know, when uttering U_2 , that what I was affirming was the same thing I was affirming in uttering U_1 . Oftentimes, literal meaning is cloaked by the *pre-semantic* information through which it is grasped, and semanticists are needed to uncloak it.

4.2 Semantics needed to figure out what is literally meant and what is not (continued)

Fido is the smartest dog on the planet. I know this well, but some of my friends don't yet know this. I point to Fido and say: "That dog is very smart."

The proposition that it was my intention to affirm and communicate is indeed true. *For a dog*, Fido is indeed smart. Of course, Fido is vastly less intelligent than a human being, such as my friend Timmy, who is of mediocre intelligence. But if I say "Timmy is very smart," what I'm saying is *false*. Judging by the words I've used, the property I've attributed to Timmy is identical with the property I've attributed to Fido. Given that Timmy has that property to a vastly greater degree than Fido, it would seem to follow that, since "Fido is smart" is true, "Timmy is smart" must also be true. And yet "Timmy is smart" is false. How can this be?

Some deal with this by saying that "smart" is ambiguous, like the word "dumb." So "Timmy is smart" and "Fido is smart" have different meanings, like "Timmy is dumb [unintelligent]" and "Timmy is dumb [mute]."

This solution is pretty clearly false. A more plausible one is that the property of being smart *for a dog* is distinct from the property of being smart *for a human*. Fido has the first but not the second. And many humans have the second, but almost all of those lack the first.

A similar, possibly coincident, view is that "smart" is implicitly relational. When, for some object x , you say $\langle x$ is smart \rangle , you are saying that x is smart *relative to some benchmark*, the identity of which the context makes clear. So "Fido is smart" says that Fido is smarter than most dogs, which is true, and "Timmy is smart" says that Timmy is smarter than most human beings, which is false.

I published a paper⁶ arguing that, for any degree-property ϕ , $\langle x$ has ϕ \rangle expresses a proposition of the form: the degree to which x has ϕ exceeds standard S , where S is some standard that, given the context, is clearly the relevant one. (A "degree property" is one that can be had to varying degrees.) But even if this is right, it doesn't follow that such judgments are the literal meanings of such sentences. And there is no independent evidence that $\langle x$ is smart \rangle has the syntactic properties of sentences that clearly *do* have for their literal meanings propositions of the just-described kind. This suggests that, so far as $\langle x$ has ϕ \rangle communicates such a proposition, it isn't because it semantically encodes it.

In any case, it not *obvious* what $\langle x$ is smart \rangle means or, in general, what $\langle x$ has ϕ \rangle means, where ϕ is any degree property. Thus, the literal meanings of such sentences are sufficiently recondite that the intervention of professional semanticists is needed to identify them.

5.0 The Nature of Semantic rules

The English language assigns a certain meaning to the sound "that dog has rabies"; and given the spectacle of a rabid dog, it furnishes one with a sentence with which to describe what one sees. In general, the English language assigns meanings to

sentences and sentences to meanings. This is true of all languages. A language is a systematic way of pairing off sentences with meanings. Any rule that assigns a meaning to an expression is known as a “semantic rule.” Languages are sets of semantic rules.

5.1. An important subtlety

There are a couple of subtle but, in some contexts, important inaccuracies in what I just said. First of all, something isn’t a sentence *until* a meaning has been assigned to it. In a world where there were no animate beings, but in which the forthcoming parenthetical ink deposit ($1 + 1 = 2$) was formed out of twigs, that twig deposit wouldn’t be an expression of any kind. That twig deposit *would* be an expression *if* and *only* if it were endowed with meaning. This shows that something has to have meaning *in order* to be an expression. Thus, *expressions* aren’t assigned meanings. They already have them and don’t need to be assigned them. Therefore a semantic rule can’t be defined as a rule that assigns a meaning to an expression.

5.2 How meaning is assigned to hitherto meaningless and, therefore non-linguistic, entities

Thus, semantic rules assign meanings to non-expressions. But which non-expressions? A story will give us the answer.

You and I want to invent a code that only we two know. We both know a guy whose real name is Larry. We decide that our code name for Larry is to be “Ichabod.” So what’s going on is that we’re creating a semantic rule: one that to the effect that “Ichabod” is to pick out Larry. How exactly is this rule enacted?

In order to implement this rule, I say: Let’s refer to Larry as “Ichabod.” The burst of noise that I produce is a sentence-token. And what you *hear* is some *token* of “Ichabod”—you do not, since one could not, hear the name type itself.

When you hear this burst of noise, along with my proposal concerning our new name for Larry, you know that what I’m saying is to the effect that any other physical object that is similar in the relevant ways to this burst of noise is itself henceforth to refer to Larry. Thus, I am in effect proposing that all and only those bursts of noise that are similar, in the relevant respect, to this particular burst of noise are to refer to Larry. (By implication, I’m proposing the same thing *mutatis mutandis* to hold of all and only ink-deposits that, given certain conventions, are paired off with such bursts of noise.) The thing that, according to my proposal, is henceforth to pick out Larry is the thing of which all and only such bursts of noises (etc.) are *instances*. That thing, like anything else of which there are instances, is a property. It is the property had in common by all and only bursts of noise (etc.) of the relevant type.

That property doesn't (yet) have a meaning; it isn't (yet) an expression. It's a property that existed, and was instantiated, before either or any instances meant anything. So the semantic rule that I'm proposing we adopt assigns a referent to a property that does not itself have a meaning. The same thing *mutatis mutandis* holds of any other semantic rule.

Thus semantic rules assign meanings to properties of physical objects—to morphological or acoustical properties (in other words, to properties that a things has in virtue of having a certain shape or sounding a certain way). A semantic rule is therefore something which assigns a meaning to a property, and language is a set of such rules.

(Technically, this is only an approximation to the truth. The relevant qualifications are found in Section 7.5.)

5.3 What are semantic rules?

Many believe that linguistic meaning is to be understood in function-theoretic terms—that, in other words, semantic rules are mathematical *functions*.

Let us start by defining the word “function.” Given any pair of whole numbers, the expression “+” assigns exactly one whole number to that pair. In general, a function is a rule that, given some class of objects, assigns no more than one object to any given member of that class.

Although the rule expressed by “plus” assigns the number 8 to the pair $\langle 4,4 \rangle$, it doesn't assign that number to that pair in the way in which a person assigns a task to an underlying. In the former case, the word “assigns” has a psychological meaning; in the latter, it has a non-psychological, purely logical meaning. A related point is that the rule that assigns 8 to $\langle 4,4 \rangle$ isn't a social rule, and it therefore isn't something that can be obeyed or disobeyed.

According to the function-theoretic view, semantic rules are rules in the strictly logical sense; that is, they are mathematical functions. The semantic rule for “Socrates” is a function that assigns a certain individual (Socrates) to that word (or to occurrences thereof). The semantic rule for “snow is white” is a function that assigns truth-conditions to that utterance (or to occurrences thereof). And so on.⁷

5.3.1 Why semantic rules are not functions

The just-described view is false. The rule denoted by “+” has always existed and always will. Of course, the *expression* “+” hasn't always existed. But that's irrelevant, since things pre-exist the expressions we use to denote them. “Socrates” is an Anglicization of the name with which Socrates referred to him. Since Socrates lived well before the English language came into existence, “Socrates” (the name, not the person) didn't come into existence until well after its referent went out of it.

The semantic rule that assigns Socrates to ink deposits having certain shapes would exist even if the English language had never come into existence. Like the rule denoted by “+”, that rule has always existed, and always will. So has the rule that assigns Abraham Lincoln to such ink deposits. The mathematical function that assigns the proposition *snow is white* to ink deposits like the italicized one has always existed, as has the mathematical function that assigns the proposition *all horses weigh 18,000 lbs* to those same ink deposits.

But the English language hasn’t always existed. Since the English language *is* a set of semantic rules, those semantic rules haven’t always existed. Therefore, they haven’t always existed. They came into existence quite recently. Therefore, those rules aren’t mathematical functions.

Also, if the semantic rules of English *were* such functions, there would exist a language in which “Socrates” referred to Lincoln and in which “snow is white” meant *all horses weigh 18,000 lbs*, the reason being that the corresponding mathematical functions exist. But there is no such language. Of course, there *could* be such a language. And maybe there will be; maybe somebody will invent a code in which those things have those meanings. But right now they don’t. Such a language is merely possible and, therefore, doesn’t exist. Thus, semantic rules are not rules in the mathematical sense.

5.3.2 The Gricean approach

Understandably, many philosophers of language believe that semantic rules must be understood in psychological, not mathematical, terms. There are different versions of this view. I accept one version of it. But the version I accept bears little resemblance to the versions of it that are usually held, each of which is some variant of the view held by H.P. Grice.

According to Grice (1957), for expression E to have literal meaning M is for it to be the case that, when they utter E, M is what they mean. So “snow is white” has the proposition snow is white for its literal meaning because what people generally mean when they say “snow is white” is that snow is white. People generally mean snow is white in uttering that expression.

Literal meaning is to be understood in terms of speaker’s meaning. That’s the main idea. Neo-Griceans hold that, even though literal meaning cannot in all cases be identified with speaker’s meaning, it is always, ultimately, to be understood in terms of it.

Wittgenstein (1958) advocated a version of this view. “Roughly speaking,” he said, “meaning is use.” Expressions mean what we use them to mean—they mean what *we* mean by them. Wittgenstein nowhere makes it clear what he means by the words “roughly speaking.”

But it doesn't matter, since his statement isn't even roughly true. Literal meaning is isn't identical with speaker's meaning and isn't to be understood in terms of it. It's the other way around. We saw why in Section 1.3.

Also, Grice's theory fails to deal with the fact that the meaning of a subsentential expression isn't something that could be possibly be meant. "Of" has a meaning; so does "skip," "snorkel," "or," "gladly," etc. But whatever it is that "or" means, it cannot, at least not by itself, be what a *person* means. Obviously I can say "snorkel" and mean it—but only if I'm using it as an abbreviation for some whole sentence (e.g., "my favorite activity is to snorkel") and, therefore, to convey something other than its literal meaning.

5.3.3 More problems with the Gricean approach

If Grice were right, the meaning of the sentence:

(SF) *"Smith is now living in France"*

would be fixed by the intentions people have in using it.

But its meaning is *not* fixed by those intentions. It is fixed by the meanings of its parts ("Smith," "France," etc.), together with the way those expressions are ordered in that sentence. The semantic and syntactic rules of English being what they are, SF would have its current meaning even if it had never been used. So whatever the intentions of people using that sentence are, those intentions do nothing in the way of assigning it that meaning. In general, Grice's view is inherently incapable of accommodating the fact meaning is compositional.

5.3.4 How some Griceans deal with the problem just described

Some Griceans respond by saying that, although speaker-meaning doesn't *directly* fix sentence-meaning, it does so indirectly. In their view, it is because of what we mean by sentences of the form <...France...> that such sentences are to the effect that...*France*...and not to the effect that, for example,...*Germany*....

In addition to being an abandonment of Grice's core idea, this move is a failure. Let P be the proposition meant by SF. So far as people utter SF with the intention of affirming P, it's because they believe (correctly, as it happens) that each of the expressions composing it already has a certain meaning.

It's irrelevant that how we use sentences of the form <...France...> *causally* determines what "France" means. There are many ways to *cause* meaning-shifts—many ways to *get* a given expression to have a certain meaning. But the question we're asking isn't "how did 'France' *acquire* its current meaning?", and is instead "whatever it is that 'France' means, what *is* it for it to have that meaning?" And it's no

answer to this question to say that it may have acquired that meaning because of what, at some point in time, people meant by sentences of the form <...France...>.

5.3.5 Why Grice's theory is inconsistent with the normative nature of semantic rules

A billiard ball isn't right to move after being struck; it just does. The relevant scientific laws merely register that fact; they aren't normative—that is, they don't characterize it as good or bad. Unlike scientific laws, semantic rules *are* normative. If, intending to affirm that Smith is female, you say "Smith is male," you've done something wrong. If Grice is right, literal meaning is speaker's meaning. This means that, if Grice were right, speaker's meaning wouldn't be accountable to existing semantic rules. Since it is, Grice is wrong.

5.4 The psychological reality of semantic rules

One view as to the nature of semantic rules is that they are idealized descriptions of the activities of speakers. Proponents of this view seldom if ever identify the facts about the speaker-behavior of which semantic rules are supposedly descriptions. The most natural assumption is that they are idealized descriptions of what people *mean* when they speak and write (etc.). If this assumption is right, then, given the points just made, the view in question is wrong.

But even if this isn't what proponents of this view have in mind, their view is very clearly wrong. If semantic rules are just idealizations of speaker-behavior, then speaker-behavior must pre-exist the semantic rules embodied in it. But if that's the case, then the activity described by semantic rules isn't *guided* by them. An awareness of those rules is no part of what leads people to say the things they do. Those rules are psychologically inert. They have no "psychological reality." This view is held by Nathan Salmon (2007) and also by Scott Soames (2002).

This view is inconsistent with some obvious facts. I know that Smith is now living in France. Wanting to tell you this, I say: "Smith is now living in France." Why do I choose this particular sentence? Because, first of all, I know the relevant semantic rules (viz. that "Smith" refers to Smith, that "living" refers to a certain property) and, secondly, because I believe that, given these facts about semantics, the sentence in question is the right one to express my belief. A knowledge of semantics underlies my speech-act and, by obvious extensions of these points, all non-defective speech-acts.

Another problem with the Salmon-Soames view is that it's inconsistent with the normative character of semantic rules. If semantic rules merely describe existing semantic activity, then that activity isn't answerable to semantic norms. It is; so the Salmon-Soames view is wrong.

5.5 Conceptual role semantics

A little while ago, we discussed Wittgenstein's claim that "meaning is use," i.e., that for an expression to have a given meaning is for it to be used in a certain way. This doctrine is incoherent in many ways. We've already discussed one of those ways; now we'll discuss some of the others.

Expressions have meanings. A meaningless burst of noise isn't an expression. If I cough or guffaw, the burst of noise I've produced doesn't have the sort of meaning had by *bona fide* expressions. It has, at most, meaning in irrelevant, purely evidential sense, e.g., the sense in which a cough may be evidence of a cold. Since anything that is a linguistic expression *ipso facto* has a meaning (in the relevant, linguistic sense), there are no expressions to be used *before* noises, ink-marks, etc., have been assigned meanings. So, since there can be no expression-use until after there is expression-meaning, it makes no sense at all to say that expression-use determines expression-meaning. "Meaningful expression" is a pleonasm.⁸ So, contrary to what Wittgenstein said, meaning isn't use.

Of course, how a given expression is used may well assign it a *new* meaning. But there is all the difference in the world between saying:

(1) *Expression E's having meaning M is causally determined by E's being used in such and such a manner,*

and

(2) *Expression E's having meaning M is identical with E's being used in such and such a manner.*

An expression E's having meaning M cannot possibly be *constituted* by its being used in such and such a manner, since E isn't an expression and, therefore, isn't an expression to be used *until* it has a meaning.

According to a doctrine known as "conceptual role semantics" (CRS), whose exponents include Hartry Field (1977) and Robert Brandom (1994), for an expression to have a given meaning is simply for it to be used in a certain way. But this isn't correct as we just saw.

CRS is incoherent for reasons other than the one just given. According to that doctrine, what a sentence means is determined by what people infer from it and what people infer it from. Whereas commonsense holds that one infers "an even number is less than five" from "two is less than five" *because* the latter already has a given meaning, advocates of CRS say that, on the contrary, it's because people infer "an even number is less than five" and other similar statements from "two is less than five" that the latter has the meaning it has.

This is not a viable view. If we learn tomorrow that, contrary to what we previously thought, Aristotle wrote several plays, we'll infer "somebody who wrote several plays wrote the *Nichomachean Ethics*" from "Aristotle wrote the *Nichomachean Ethics*." But it doesn't follow that "Aristotle wrote the *Nichomachean Ethics*" would have undergone some change in its semantic *meaning*. Changes in what we believe affect what we infer from statements; but they don't categorically change the meanings of those statements. CRS entails that every inference is an analytic inference. Once it's learned that Aristotle wrote plays, it becomes, according to CRS, constitutive of the *meaning* of "Aristotle wrote the *Nichomachean Ethics*" that one can infer from it that a playwright wrote the *Nichomachean Ethics*. But surely the inference from "Aristotle wrote the *Nichomachean Ethics*" to "a playwright wrote the *Nichomachean Ethics*" isn't *analytic*.⁹

What we may infer from a sentence is answerable to its existing meaning. CRS says that a sentence's meaning is answerable to what we infer from it. If correct, that would have the consequence that one couldn't possibly draw a false inference from any sentence. Which, in its turn, would have the consequence that no sentence would mean anything. Which, since nothing meaningless is a sentence, would have the absurd consequence that there neither are, nor could be, sentences.

Consider the sentence "Bill plagiarized his first novel." If one knows that sentence to be true, one can make inferences about Bill's character, his past activities, his ambitions, his values, and so on, that one couldn't make if, other things being equal, one didn't know that sentence to be true. But it's only because of what the sentence *already* means that one can make those inferences. Given any other sentence, the same thing *mutatis mutandis* is true of it. CRS says that what a sentence means is determined by what we infer from it. Since it's the other way around, as we've just seen, CRS is false.

6.0 What is literal meaning?

Where complex expressions are concerned, there is no limit to how much literal and understood meaning may diverge from each other. But where simple expressions are concerned, literal and understood meaning must coalesce. It makes no sense to suppose that people could be systematically wrong as to what "red" meant. If people *thought* that "red" meant what is in fact meant by "blue," then "red" *would* have that meaning. Systematic, widespread error is impossible where semantically simple expressions are concerned. This gives us a way of understanding what literal meaning is.

Even though there can be widespread, systematic misinterpretations of sentences, those misinterpretations do not arise as a result of people failing to know

what the simple parts of sentences mean. They arise as a result of people not knowing how to put those meanings together. *So to the extent* that its meaning is fixed by the fact that it has the form <...Socrates...>, people (English speakers) *do* systematically understand “Socrates was more wise than Plato, but he was less sharp than Aristotle”; and *to the extent* that its meaning was determined by its having the form <...wise...>, people *do* understand that sentence; and so on. So far as that sentence *is* systematically misunderstood, it is because people are having trouble putting the meanings of its constituents together—it is because they’re having trouble figuring out how those meanings ought to be put together.

What a simple expression literally means is determined by what it is that a sentence means by virtue of containing it. Since, where simple expressions are concerned, what people take literal meaning to be coincides with what it really is, a simple expression’s literal meaning coincides with what it is that, in virtue of containing that expression, sentences are *taken* to mean. So a simple expression’s literal meaning is given by a statement saying what it is that, by virtue of containing it, sentences are taken to mean; and a complex expression’s literal meaning is a function, in the mathematical sense, of the meanings of its parts.

Here’s an illustration. “Socrates” is a simple expression. So what people think it means must ultimately coincide with what it actually means. So it refers to Socrates only because people think it refers to Socrates and, therefore, only because people think that “Socrates is intelligent” attributes intelligence to Socrates and, in general, that <...Socrates...> attributes...x...to Socrates.

Of course, people don’t always take utterances of <...Socrates...> to be attempts to *say* that Socrates has...x... It might be clear from the speaker’s tone that what he *really* meant when in saying “Socrates was wise” was that Socrates was *not* wise. But *to the extent* that their belief that “Socrates” refers to Socrates is determinative of what people take the meanings of utterances of the form <...Socrates...> to be saying, what they take it to be saying is that Socrates has...x... That is what it is for them to take “Socrates” to refer to Socrates. And their taking “Socrates” to refer to Socrates *is* for “Socrates” to refer to Socrates, given that “Socrates,” being a simple expression, has the semantics that people think it has. The same thing *mutatis mutandis* is true of every other simple expression.

Bearing these points in mind, let CE be any complex expression, and let $e_1...e_n$ be the simple expressions composing it. How people interpret CE may diverge from its literal meaning. But when this happens, it’s because what it is taken to mean diverges from what, given what people believe its simple parts to literally mean, people are *disposed* to take it to mean.

Since what people take simple expressions to mean *is* what they mean, this is the same as saying the following. A divergence between

(i) *A sentence’s literal meaning*

and

(ii) *That sentence's understood meaning*

is the same thing as a divergence between

(a) *That sentence's literal meaning*

and

(b) *What it is that, given their (correct) beliefs as to what its simple parts literally mean, people are disposed to take that sentence's literal meaning to be.*

The literal meaning of a complex expression is a function of two things: (i) the meanings of its simple parts, and (ii) the order in which those parts are arranged. To say what literal meaning is in general, we need to say what it is for a simple expression to have a given literal meaning. Given the points just made, we can do this. Where simple expressions are concerned, literal and communicated meaning coincide (ultimately)—in other words, such expressions mean what people take them to mean. And where complex expressions are concerned, literal and communicated meaning may diverge, but literal meaning nonetheless coincides with what, given what the simple components of the expression in question literally mean, people are *disposed* to take it to mean.

7.0 Tokens vs. types: some preliminary terminological points

No word is identical with any utterance of it. My utterance of the word “snow” lasts for a fraction of a second. But the word “snow” *itself* endures.

Utterances and inscriptions of expressions are referred to as “expression-tokens” (or just “tokens”). So there are three *tokens* of some one word to the right of the upcoming colon: snow, snow, snow. The things being uttered or inscribed are referred to as “expression-types” (or just “types”).

7.1 Two-dimensional semantics

Some expressions appear to have a two-tiered semantic structure. For example, an occurrence of the pronoun “I” has a referent, this being the person who

uttered it, and it picks out that referent in a certain way (i.e., by way of a certain concept).

Thus, the semantics of “I” is given by the rule that an utterance of “I” refers to a given person if and only if that person has the property of being the one that produced that utterance.

So, if I say “I am tired,” the concept through which my utterance of “I” refers to me is the concept person who produced the utterance in question. This, then, is the concept through which reference is secured; it is, as we’ll henceforth put it, the mediating concept.¹⁰ When Smith says “I am tired,” his utterance of “I” refers to himself, not to me. But the mediating concept remains the same.

If somebody points to me and says “that guy is tall,” the person picked out by “that guy” is me. But in this case, the mediating-concept is different. The semantics rule for “that guy” is given by the rule that, if “that guy” is uttered in a context where there is a unique, salient guy, that utterance refers to that individual.

“I” and “that guy” are context-sensitive expressions: what such an expression refers to depends in a systematic manner on facts about the context of utterance. (We’ll soon refine this vague statement shortly.) Such expressions are known as “indexicals.” Other examples of indexicals are “you,” “he,” “those animals,” “this monkey,” “tomorrow,” “yesterday.” Some indexicals are single words (e.g., “tomorrow,” “he”); others consist of more than one (e.g., “that tall man”). The latter are known as “complex indexicals.”

7.1.1 Demonstratives vs. indexicals

Some indexicals often cannot be successfully used without an accompanying demonstration on the speaker’s part, the purpose of which is to eliminate any doubt as to what the intended referent is. If Jim and Larry are both equally salient in the context in question, an utterance of “that guy” won’t single anyone out. But it will do so if an act of pointing accompanies it. Indexicals that fall into this category are known as “demonstratives.” Not all indexicals are demonstratives. For example, “tomorrow,” “now,” and “here,” aren’t demonstratives. Given an utterance of “tomorrow,” no gesture is needed to make it clear what the intended referent is. If I say “tomorrow I’m going hiking,” it’s clear what the referent of “tomorrow” is; no demonstrative act is necessary, and none could possibly do any good.

7.1.2 Indexicals (continued)

Indexicals have a “two-dimensional” semantic structure. Consider the expression “today.” If I say it right now, that utterance will refer to April, 27, 2009, since that is today’s date. If I say it tomorrow, that utterance will refer to April 28, 2009. But the *rule* that assigns April 27, 2009, to the first utterance is identical with the *rule* that assigns April 28, 2009, to the second utterance. That rule is this: if “today” is uttered on a given day D, that utterance refers to D.

The meaning of an *utterance* “today” is the day it picks out. The meaning of the corresponding word-*type* is the rule just described. Given any indexical, the meaning of a *token* of that indexical is its referent; the meaning of the corresponding indexical-type is the rule that assigns that referent to that token.

7.2 Definite descriptions

In other works of mine, reasons are given for thinking that definite descriptions are not devices of reference. But there are also reasons to think that they are devices of reference.¹¹ (It’s very hard to believe that “the whole number that comes right after one” doesn’t refer to the number two.) And in this section we will suppose them to be just that.

Definite descriptions, like indexicals, have a two-dimensional semantic structure. A given utterance of “the current U.S. President” refers to some individual. Right now such an utterance would refer to Barack Obama. A few years ago, such an utterance referred to Bill Clinton.

But even though an utterance in 2009 of “the U.S. president” doesn’t have the same referent as an utterance in 1999 of that same expression, both utterances are assigned their respective referents by the same semantic rule. The circumstances have changed—hence the change in referent—but the semantic rule has stayed the same.

That rule is: if, at time *t*, *x* uniquely has the property of being a U.S. president, then an utterance at *t* of “the U.S. President” refers to *x*.

7.2.1 Incomplete definite descriptions

Some definite descriptions appear to be “incomplete”—that is, they fail to pick out a single object. So, for example, “the bald guy” could pick out any one of many different people. But, like all expressions, definite descriptions are not used in a vacuum; and the context usually supplies the information needed to enrich the mediating concept enough to enable it to single out a single person. So if you and I are in a room and I say “the bald guy is wealthy,” it’s clear that, so far as I making a

determinate statement, what I mean is that the *contextually salient* bald guy is wealthy.

If phi is a property that obviously has no more than one instance (e.g., the property of being the whole number successor of one or of being the U.S. President), phi's sole instance is contextually salient by default.

7.3 The expressive limitations of indexical-free languages

Indexical-free languages are expressively impoverished; much of what there is to say can't be said in them. The reason is that much information is *perspectival*, and nothing perspectival can be stated in a language that doesn't contain indexicals. Suppose that I spoke a language that didn't contain any indexicals, but was otherwise just like English. In order to express what I believe concerning the weather in my area, I'd have to say "it's 60° in Richmond on Feb. 07, 2009." (And the "is" couldn't be taken as the *present tense* of the verb "to be," since, thus interpreted, it would be an indexical that picked out the time of utterance. The "is" would have to be stripped of any temporal meaning and, thus, be downgraded to an empty grammatical placeholder, like the "it" in "it's raining.") But that utterance wouldn't necessarily express the belief I wanted to express. My believing *it's 60° in Richmond on Feb. 07, 2009*, is different from my thinking *it's now 60° in Richmond*, even though I am in fact in Richmond at the time in question. I could have the one thought and not have the other. I could, after all, not know that I was in Richmond. And my believing either of those things is different from my thinking *it's now 60° here*. One could have any given one of those beliefs without having any of the others. Further, one could *rationally* have any given one of those beliefs without having any of the others. A person who is in Richmond at the time in question can *rationally* believe *it's now 60° here* while rejecting *it's 60° in Richmond on Feb. 07, 2009*. For the data at one's disposal may warrant the first judgment, but not the second.

So nothing perspectival—nothing that embodies any information relating to the speaker's perspective on the world—can be expressed without using an indexical. Nothing could be said about what's going on here, at *this* time; and there would be no me-thoughts (e.g., *I'm thirsty*) could be expressed; one could only express their third-person counterparts (e.g., "JM is thirsty"). Thus, one could not, in an indexical-free language, express the thought *I am JM* or *I am that guy in the mirror*. And such a language would therefore be extremely impoverished.¹²

7.4 Tokens, types, and context-sensitivity

No sentence-*type* containing a context-sensitive component is either true or false; for no such sentence-*type* says anything. It is *tokens* of “I am now tired” that make statements; the corresponding sentence-*type* does not do so. What a token of “I am tired now” affirms depends on facts about the context of utterance. If Smith’s the one who’s speaking, and it’s 3:00 P.M., such an utterance is true just in case Smith is tired at 3:00 P.M. If Jones is the speaker and it’s 4:00 P.M., such an utterance is true just in case Jones is tired at 4:00 P.M.

Thus the semantic rule that assigns a proposition to such a token does so on the basis of the facts about the context of utterance. And the rule in question is clearly this: “If, at time *t*, person *p* tokens the sentence-*type* “I am tired,” then the proposition thereby affirmed is true exactly if *p* is tired at *p*.” The same thing *mutatis mutandis* is true of all context-sensitive expressions. So the rule that assigns a proposition to an utterance of “that man is married to that woman” is: if, in the context of utterance, *x* is a uniquely salient man and *y* is a uniquely salient woman, a token of the sentence-*type* “that man is married to that woman” affirms a proposition that is true if and only if *x* is married to *y*.

So context-sensitive sentence-*types* aren’t true or false; they don’t, in and of themselves, bear propositions. But they have an important semantic role: the identity of the proposition meant by a token of such a type depends on the identity of the type. A token of “I am tired” means one thing; a token of “you are tired,” uttered by the same person at the same time, means something else; and that difference obviously stems from the fact that, because different sentence-*types* were tokened, the semantic rule that assigns a proposition to the one token is different from the rule that assigns a proposition to the other.

This can all be distilled into the following principle: where context-sensitive sentences are concerned, the meaning of a token is a proposition, and the meaning of type is a *rule* that assigns a proposition to one of its tokens on the basis of facts about the context in which that token occurred.

Long story short: the meanings of sentence-tokens are propositions and the meanings of sentence-*types* are rules that assign propositions to their tokens, usually on the basis of facts about the context of tokening.

7.5 Ambiguity and context-sensitivity (revisited) and the type-token distinction (revisited)

Consider the sentence-*type* “that person is a professor.” Is that sentence true or false? No. Some tokens of it are true and some are false. This is because what a given token of that sentence says is a function of the circumstances. If I utter it while pointing at Bob, I’m attributing the property of being a professor to Bob. If I utter it

while pointing to Sally, I'm attributing that property to Sally, not Bob. So what it is that I'm affirming in the one case is different from what it is that I'm affirming in the other case. But that isn't because the sentence "that person is a professor is ambiguous." That sentence is not ambiguous: it has just one meaning. But that meaning is not a proposition; it isn't something that is true or false. That meaning is a rule. That rule in its turn assigns meanings to tokens of that sentence. Those meanings are true or false; those meanings are propositions.

Remember that expression-types are what result when properties *per se*, as opposed to their instances, are assigned meanings. Of course, the rules that assign meanings to properties are semantic rules, since anything that assigns meaning to anything is ipso facto a semantic rule. So the sentence-type of which the following ink deposit—"that guy is a professor"—is a token is what results when some semantic rule assigns a meaning to some property. Let R_1 be the semantic rule in question. Let R_2 be the meaning that R_1 assigns to the just-mentioned property. R_2 is *itself* a semantic rule. But, whereas R_1 assigns a meaning (a rule) to some property, R_2 assigns meanings to instances of that property. So R_2 assigns meanings to particular instances of the morphology had by the following ink deposit ("that guy is a professor"). The meanings that R_2 assigns to those tokens are not themselves rules; they are propositions—they are things that are true or false.

There are thus very different sorts of semantic rules. There are those that assign meanings to properties, and there are those that assign meanings to instances of those properties. The meaning that was assigned to the property we were just discussing is itself a semantic rule. As we'll now see, *every* meaning that is assigned to a property (as opposed to a property-instance) is itself a semantic rule. Consider the following ink deposit: "Barack Obama." There is a semantic rule that assigns a meaning to the *property* of having a morphology similar to that ink deposit. The meaning assigned to that property is *itself* a semantic rule. *That* rule assigns a meaning to each instance of the property in question. It assigns Barack Obama, the person, to any such instance. Thus, any such instance is an expression that picks out Barack Obama.

In the previous section, we said that semantic rule is something which assigns a meaning to a property. That's correct, but incomplete. The right definition is this: a semantic rule is something which assigns a rule to a property that in its turn assigns meanings to instances of that property. Let us now move onto slightly less abstruse material.

7.6 Do any expressions have one-dimensional semantics?

The question arises: What about sentence-types that don't contain context-sensitive components? Where they are concerned, does token-meaning coincide with type-meaning?

First of all, in natural languages, there are no such sentence-types. Every sentence contains a tense-marker. And in virtue of containing a tense-marker, a sentence is such that the proposition expressed by any given one of its tokens is a function of (*inter alia*) when that token occurs. If you say "the U.S. economy is fairly stable," whether you are speaking the truth or not depends on when you say it.

There are some *apparent* counterexamples to this. For all intents and purposes, any two tokens of

(IA¹³) "*The interior angles of a Euclidean triangle add up to 180°*"

express the same proposition. But, from a strictly semantic perspective, IA is context-sensitive, and tokens of it uttered at different times *don't* encode the same proposition. If it were believed that mathematical reality were as volatile as the stock-market, we would without hesitation regard IA as being in the same category as patently context-sensitive sentences such as:

(BH) "*Bill's holdings add up to \$180,000,000.*"

And we'd have no more temptation to regard the tense-marker in IA as inert than we'd have to regard its counterpart in BH as inert. What this shows is that, to the extent that the tense-marker in tokens of IA are doing nothing, it's only because, our beliefs about mathematics being what they are, we choose to see such utterances as expressing atemporal propositions.

Nonetheless, many semanticists hold that, at the level of literal meaning, IA is context-insensitive. They hold, in other words, that IA is what Quine (himself such a semanticist) refers to as an "eternal sentence." (S is an eternal sentence if there is some *one* proposition P such that any two of S's tokens express P.) If only for argument's sake, let's suppose that there are eternal sentences.

It's tempting to hold that an eternal sentence's meaning is identical with those of its tokens. After all, the distinction between type-meaning and token-meaning seems quite hollow *except* where context-sensitive expressions are concerned.

But this reasoning is spurious. The tense-marker on the occurrence of "add" in (IA) has the same semantics that it does in

(BH¹⁴) "*Bill's various holdings add up to \$180,000,000.*"

In BH the tense-marker is obviously doing real work in IA. That's why utterances of BH are true on Monday (before the market crashed) and false on

Wednesday (after the market crashed). Therefore, different tokens of BH express different propositions. Therefore, the meaning of BH—the sentence-type—isn't some proposition, and is instead the rule:

(BHSR) If t is a token of BH that is uttered at time T, t is true iff, at T, Bill's various holdings add up to \$180,000,000.

But, of course, for any time t, the meaning of a token of BH that is produced at t is a proposition that is true just in case.

(BHT) At t, Bill's various holdings add up to \$180,000,000.

Given that occurrence of "add" in any given token of IA has the same literal meaning as its counterpart in any given token of BHT, it follows that, from a narrowly semantic perspective, IA is quite as context-sensitive as BHT. To be sure, there is obviously a sense in which IA is, whereas BHT is not, context-*insensitive*.

But IA's context-insensitivity, we must conclude, is a thoroughly pragmatics-based affair. Obvious extensions of this reasoning show that *all* expression-types have *rules* that assign meanings (or referents) to their tokens and, therefore, that where any expression is concerned, type-meaning diverges from token-meaning. The meaning of the expression-type "that man" is: a token of "that man," uttered in context C, refers to x if, in C, x is uniquely a salient man. If, in a given context, Smith is such a man, then a token in that context of "that man" refers to Smith; if instead Jones is such a man, it refers to Jones; etc. The same is true of "the current President." In one context (the year 1992), tokens of that expression refer to Bill Clinton; in a different context (the year 2009) it refers to Barack Obama.

Some would be tempted to say that, because they have fixed referents, proper names (e.g., "Bill Clinton") don't have a two-dimensional semantic structure. This would be a mistake. Expression-tokens are physical entities—bursts of noise, deposits of ink, etc. Expression-tokens are therefore perceptible entities and are thus capable of transmitting information. Expression-types, on the other hand, are abstract entities and are thus inherently unsuited to be vehicles of communication. Let us develop these points.

Expression-types are properties. Why are they properties? Tokens of a type are instances of it. Anything of which there are instances is ipso facto a property. So expression-types, unlike expression-tokens, are properties and are therefore non-spatiotemporal entities. And, as we just noted, it makes no sense to suppose a nonspatiotemporal entity could mediate information or, therefore, could in any significant sense be a symbol. Also, given the profound metaphysical differences between tokens and types, it would be theoretical arbitrariness of the worst kind to suppose that types could discharge the same semantic functions as their tokens. We

must therefore assume that, whereas tokens of “Bill Clinton” refer to Bill Clinton, the corresponding type does not refer to Bill Clinton; and we must also assume that “Bill Clinton,” the expression-type, has for its meaning a (constant) function that assigns Bill Clinton to any given one of its tokens. In general, proper names, no less than indexicals and definite descriptions have two-dimensional semantics.

The semantic rule corresponding to a proper name is a constant function, whereas the semantic rule corresponding to an indexical or definite description is not a constant function. All utterances of “Bill Clinton” refer to Bill Clinton, but not every utterance of “the current U.S. president,” or of “that guy over there,” so refer. This has encouraged the erroneous view that “Bill Clinton” itself, the expression-type, refers to Bill Clinton.¹⁵

7.7 Ambiguity and context-sensitivity (re-revisited) and the type-token distinction (re-revisited)

Remember that expression-types are what result when properties *per se*, as opposed to their instances, are assigned meanings. Of course, the rules that assign meanings to properties are semantic rules, since anything that assigns meaning to anything is ipso facto a semantic rule. So the sentence-*type* of which the following ink deposit—“that guy is a professor”—is a token is what results when some semantic rule assigns a meaning to some property. Let R_1 be the semantic rule in question. Let R_2 be the meaning that R_1 assigns to the just mentioned property. R_2 is *itself* a semantic rule. But, whereas R_1 assigns a meaning (a rule) to some property, R_2 assigns meanings to instances of that property. So R_2 assigns meanings to particular instances of the morphology had by the ink deposit (“that guy is a professor”). The meanings that R_2 assigns to those tokens are not themselves rules; they are propositions—they are things that are true or false.

There are thus very different sorts of semantic rules. There are those that assign meanings to properties, and there are those that assign meanings to instances of those properties. Consider the following ink deposit: “Barack Obama.” There is a semantic rule that assigns a meaning to the *property* of having a morphology similar to that ink deposit. Let PBO be that property. The meaning assigned to PBO is *itself* a semantic rule. In other words, the meaning of PBO is given by the statement that:

B₁: Any given token of PBO refers to Barack Obama.

Given some specific token *t* of “Barack Obama”, the semantic rule for *t* is:

B₂: t refers to Barack Obama.

B_1 and B_2 are very different rules. Unlike B_2 , B_1 doesn't say anything about any *specific* token of "Barack Obama" or any other expression. The need for a two-dimensionalist approach to semantics is embedded in the very concept of what an expression is.

8.0 Logical form¹⁶

Some statements that, given their grammatical forms, appear to be about objects are in fact about properties. Frege was the first to see clearly that logical and grammatical may diverge—he was the first to grasp the very idea of such a divergence. This insight of his is embodied in his statement that the sentence

(WM) "*Whales are mammals*"

isn't about whales. As paradoxical as it may seem, he was right. WM says that *if* an object is a whale, *then* it's a mammal. But there is no specific object x such that WM says that x is a mammal. *A fortiori* there is no specific whale x such that WM says that x is a mammal; and for any number n , no matter how high, there are no whales $x_1 \dots x_n$ such that WM says that x_i ($1 \leq i \leq n$) is a whale. So just as Frege said, WM isn't about whales.

WM makes a statement, not about whales, but about the property of being a whale. It says that

(WM*) *the property of being a whale has the property of being instantiated only by mammals.*

WM* perspicuously represents WM's meaning. In other words, WM* represents WM's logical form. WM*'s grammatical form diverges from WM's grammatical form.

The reason is that WM's logical and grammatical forms pull apart is that WM contains a *quantifier*. (Examples of quantifiers are "all birds," "no man," "most whales," "three birds," and "some individuals." We'll define the term "quantifier" in a moment.) Given *any* statement of English, or any other non-artificial language, that contains quantifiers, the logical and grammatical forms of that sentence diverge. A quantifier is an expression having the property that, if a sentence contains it, that sentence is ipso facto to the effect that the *extension* of one property has a certain degree of overlap, ranging anywhere from no overlap to total overlap, with the extension of some other property. "All birds have beaks" says that the extension of the property of being a bird is a subset of the extension of the property of having a beak.

“Only birds have beaks” says (falsely) that the extension of the property of having a beak is a subset of the extension of the property of being a bird.

In virtue of a containing a quantifier, a sentence about properties, not about specific objects. Of course, a quantified (quantifier-containing) statement can *also* be about individuals. But it isn’t in virtue of containing a quantifier that a sentence concerns individuals; and it *is* in virtue of containing a quantifier that a sentence concerns properties.

In natural language, quantified sentences are pseudo-objectual statements. They appear to be about objects, but are really about properties. That is why they must be reparsed if their logical forms are to be exposed. It immediately follows that a statement’s grammatical form may diverge from its logical form. Frege discovered both facts and, therewith, created modern analytic philosophy.

8.1 Contextual definition

The fact that grammatical and logical form sometimes diverge is related to the fact some expressions are to be defined contextually. To define an expression contextually is to how, in virtue of containing it, a sentence’s meaning is affected. So, for example, “someone” doesn’t refer to anyone. It doesn’t refer to John or Sally or Jane. That is why, given any proper name N, the statement

(NS) $\langle N \text{ doesn't snore} \rangle$

is compatible with¹⁷

(SS) “Someone snores.”

Thus, the meaning of “someone” isn’t given by some rule that pairs it off with this or that individual. In other words, there is no individual N such that the semantic rule for “someone” is:

(WRS¹⁸) $\langle \text{Someone has } \psi \rangle$ is true just in case N has ψ .

Rather, the meaning of “someone” is given by the statement that:

(RS) $\langle \text{Someone has } \psi \rangle$ is true just in case the property of being a ψ is instantiated.¹⁹

Thus, the logical form SS is clearly displayed by the sentence:

(PS²⁰) “The property of being a snorer is instantiated.”

PS’s logical form therefore coincides with its grammatical form. Since SS has a different grammatical form from PS, it follows that SS’s logical form diverges from its logical form. The divergence between SS’s grammatical and logical forms is obviously a consequence of the fact that “someone” must be defined contextually.

8.2 Contextual definition: its scope and limits

Nonetheless, it would be an overstatement to say that *whenever* a sentence contains an expression that must be defined contextually, its grammatical form pulls apart from its logical form. As we’re about to see, *every* expression is to be defined contextually. Obviously grammatical form doesn’t *always* pull apart from logical form. Therefore, it isn’t *always* the case that, in virtue of containing an expression that must be defined contextually, a sentence’s logical and grammatical forms diverge.

How can it be said that all definitions are contextual definitions? Aren’t there also *denotative definitions*? (A denotative definition of an expression E says what E means by saying what it denotes.) Isn’t “Socrates” defined denotatively? Yes, it is. There is some object x such that one says what “Socrates” means if, and only if, one says that “Socrates” refers to x.

But in saying of some object x that “Socrates” picks out x, one is saying that, in virtue of having the form $\langle \text{Socrates has } \psi \rangle$, a sentence S is to the effect that x has ψ . “Socrates” refers to Socrates because, the remaining semantic rules of English being what they are, if you wish to attribute the property of being wise to Socrates, you can do so by saying “Socrates is wise”—because, in general, for any property ψ , if you wish to attribute ψ to Socrates you can do so saying $\langle \text{Socrates has } \psi \rangle$. If, in saying “Socrates was wise,” you were really saying that it was Aristotle, not Socrates, who was wise, then “Socrates” wouldn’t refer to Socrates, at least not in that context. In general, to say that E refers to O is to make a statement about effect a sentence’s containing E has on its truth-conditions. More precisely, it is to say that, in virtue of having the form $\langle \text{E has } \psi \rangle$, a sentence attributes ψ to O, for any property ψ .

But the logical forms of sentences of the form $\langle \text{E has } \psi \rangle$ don’t necessarily pull apart from their logical forms. Since there is some object x (namely, Socrates) such that ‘Socrates is tall’ says that x is tall, and since $\langle \text{x is tall} \rangle$ has the same form as “Socrates is tall,” the fact that “Socrates” is to be defined contextually does *not* entail that logical and grammatical form ever diverge.

So it is only because *certain* expressions are to be defined contextually that such divergences occur. But which expressions? *Those that cannot also be defined*

denotatively. Whenever an expression can be defined denotatively, a sentence's logical form will not, by virtue of that sentence's containing that expression, diverge from its logical form. (But, of course, that sentence's logical form may diverge from its logical form for some other reason. Thus, the logical form of "Socrates saw someone" diverges from its grammatical form; but the reason for this is that it contains the word "someone." The occurrence of "Socrates" in that sentence isn't what induces that divergence.)

Let E be an arbitrary expression that can be defined denotatively. In other words, suppose there to be some object x such that E is defined by saying that it refers to x . The logical form of $\langle E \text{ has } \psi \rangle$ is: $x \text{ has } \psi$. In general, *to the extent* that the logical form a sentence containing E is determined by its containing that expression, that sentence's logical and grammatical forms coincide. But, since such a sentence's logical form isn't determined *only* by its containing E, and is also a function of the semantics of the other expressions occurring in it, its grammatical form may still diverge from its logical form. After all such a sentence may contain an occurrence of "someone" or "nobody" or some other expression that induces such a divergence.

8.3 Frege's generalization of the concept of a function

What made it possible for Frege to revolutionize logic was his insight that grammatical and logical form diverge; and what made the latter insight possible was his generalization of the concept of a mathematical *function*.

A mathematical function is a rule that assigns no more than one object to each object falling in a given class. (So "+1" can be thought of as expressing a function or rule that assigns 2 to 1, 3 to 2, etc.)

According to Frege, the occurrence of "snores" in

(PS) "*Plato snores*"

is best represented as the open sentence $\langle x \text{ snores} \rangle$; and that open sentence is best thought of as expressing a function that assigns the truth-value *true* to each snorer and the truth-value *false* to each thing that doesn't snore. (For brevity's sake, I'll henceforth use the words "truth" and "falsehood" instead of, respectively, "the truth-value *true*" and "the truth-value *false*.")

Here's the idea. A true sentence results if the variable in $\langle x \text{ is even} \rangle$ is replaced with an expression denoting an even number and false if it's replaced with a number that doesn't denote such a number. ("2 is even" is true and "3 is even" is false.) So we can think of $\langle x \text{ is even} \rangle$ as assigning truth to two, four, etc., and falsity to one, three,

etc. For similar reasons, we can see $\langle x \text{ snores} \rangle$ as assigning truth to snorer Bob and falsehood to non-snorer Wilma, and so on.

Frege sees the occurrence of “is taller than” in

(BTM) “*Bill is taller than Mary*”

as being identical with the open sentence $\langle x \text{ is taller than } y \rangle$, and he sees that open sentences as expressing a function that assigns truth-values to *ordered pairs* of objects. So $\langle x \text{ is taller than } y \rangle$ assigns truth to the ordered pair $\langle x, y \rangle$ if x is taller than y and otherwise assigns falsehood to that pair.

Frege treats:

(BJMP) “*Bill is standing between Mary and John*”

as comprising the open-sentence $\langle x \text{ is standing between } y \text{ and } z \rangle$; and he sees that open sentence as expression a function that assigns truth to ordered triples of objects—as assigning that truth value to $\langle \text{Bill, Mary, John} \rangle$ just in case Bill is standing between Mary and John.

In this way, Frege was able to assign a *single* form to all *atomic* sentences. An atomic sentence is one that, unlike “Bill is tall and Sally is smart,” doesn’t consist of other sentences and that, unlike “someone snores,” doesn’t contain quantifiers. Non-atomic sentences are *molecular*. A molecular sentence is one that either consists of other sentences or contains a quantifier. Frege was able to show that, ultimately, all atomic sentences have the form $\langle O \text{ has } \psi \rangle$. Contrary to first appearances, BTM has the form: *the ordered pair $\langle \text{Bill, Mary} \rangle$ has ψ* , where ψ is the property had by such a pair just in case its first member is taller than its second.

Pre-Fregean logicians saw each of PS, BTM, and BJMP as having a different form from each of the other two, and this made it impossible for them to do anything meaningful in the way of formalizing inferences involving atomic sentences. But Frege didn’t have this problem, since he, unlike them, wasn’t made blind by grammatical surface structure to the underlying structural similarities.

-
1. See Grice’s classic (1957) article “Meaning.”
 2. In this context, the word “people” is meant to refer to sentient creatures of any kind—Martians, dolphins, etc. I wish to leave it open whether non-human creatures have conventions or, therefore, languages.
 3. I say “typically” because not all linguists or philosophers of language consider this division to be a legitimate one. Wittgenstein (1958), for example, thought that semantics collapsed into pragmatics.

4. Given that I'm only indicating in broad terms what the discipline of semantics is about, I've omitted some steps from this argument. See Kaplan (1989) for a more thorough analysis of sentence-utterances like the one being discussed (viz. sentence-utterances that contain **demonstratives**—words like “this” and “that.”) There's a second reason why my adding the missing steps unnecessary. The point I'm trying to make (viz. that it isn't always obvious what is literally meant by an expression) is *itself* illustrated by the fact that it may be debated whether U1 and U2 have the same meaning.
5. Although he does an, in some ways, credible job of defending this analysis of utterances containing demonstratives, Kaplan doesn't answer this very question. And this question, it seems to me, is much more important than the narrowly semantic question of what the literal meanings of utterances of “that man,” “this world,” etc., are. Kaplan notes the difficult (viz. that sentence-tokens that have the very same literal meaning can differ dramatically in what they communicate), and responds by saying (I quote, without omitting relevant contextual material): “Uh oh!” (See Kaplan (1989).) Which pretty much exhausts what he has to say about the matter. It also exhausts what Kaplan-disciples, such as Nathan Salmon, have to say about the matter. (See Salmon (1986), (2005), (2007).)
6. In 2004, in the journal *History and Philosophy of Logic*.
7. Lewis (1975).
8. One of the problems with logical positivism was that it assumed, incoherently, that *expressions* could be meaningless. There can be meaningless expression-like entities—but no meaningless expressions. And if, as the logical positivists contended, “there is no free will in a deterministic universe” is truly meaningless, then it isn't a sentence at all; it is, at best, a sentence-like entity. But it plainly is a sentence. It may be (though I do not myself believe this) that “there is no free will in a deterministic universe” is in some way or other an *incoherent* statement. But an incoherent sentence is one that bears a self-undermining meaning. And there's a big difference between bearing a self-undermining meaning and bearing *no* meaning. The logical positivists didn't make this distinction.
9. Jerry Fodor and Ernie Lepore do a beautiful job of stating this argument and others like it in *The compositionality Papers*.
10. Mediating concepts are obviously similar to what Frege refers to as ‘senses.’ But, not wanting to prejudge the delicate question of whether they're identical with Fregean senses, I've decided to give them a different name.
11. My personal view is that *ultimately* definite descriptions are devices of reference. But they are not to be understood in remotely the way that Frege thought. My official policy in this book is that they are to be understood in the way that Russell proposed. Russell's analysis, though not (in my view) a correct answer to the narrowly *semantic* question ‘are definite descriptions referring terms?’ makes it much easier to absorb the larger, more important epistemological, logical, and

metaphysical points that it was his real purpose in putting forth the Theory of Descriptions to defend.

12. See Barwise and Perry (1983) and Kuczynski (2006).
13. "IA" is short for "interior angles."
14. "BH" is short for "Bill's holdings."
15. Barwise and Perry (1983) make similar points. See Kuczynski (2007) for a thorough defense of this position. P.F. Strawson (1950) was the first author to put forth a viewpoint of this general kind.
16. Frege (1954).
17. Two statements are compatible if neither entails the negation of the other.
18. "WRS" stands for "wrong rule for 'someone.' "
19. Arguably, not all sentences containing 'someone' have the form <someone has psi.> For example, "Plato kicked someone" seems not to have that form. For this reason, RFS, at least arguably, isn't adequate. A more accurate, because more general, definition of "someone" is:
RFS*) In virtue of having the form <...someone...>, a sentence says that the property ...x... is instantiated. Thus, "Plato kicked someone" says that the property being a thing that is kicked by Plato is instantiated. RFS* collapses into RFS, the reason being that any sentence of the form <...someone...> can be represented as having the form <someone has psi.> For example, "Plato kicked someone" says the same thing as "someone is such that Plato kicked that person."
What we just said about 'someone' is true *mutatis mutandis* of all other expressions that are to be defined contextually, which, given the fact that all expressions are to be defined contextually, means that it's true of all expressions.
20. 'PS' stands for "the property of snoring."

