
The Lure of Linguistification

KENT BACH

*Referring is not something an expression does;
it is something that someone can use an ex-
pression to do.* - P. F. Strawson ('On Referring')

1 What is linguistification?

Strawson's well-known dictum identifies a classic example of what I call "linguistification." Although I don't think he was quite right, I do think he was on to something.¹ I've noticed numerous cases in which philosophers have attributed to linguistic expressions properties that are actually characteristic of people in using these expressions. Language itself can be misleading in this regard. Eating utensils don't eat, and dining rooms don't dine. So perhaps so-called referring expressions don't refer (in at least some cases). There are lots of other things that linguistic expressions are often said to do but generally don't do. I will identify a variety of such things and try to explain the lure of linguistification and why it should be resisted.

Think of linguistification by analogy with personification: attributing linguistic properties to nonlinguistic phenomena. For my purposes, it also includes attributing nonlinguistic properties to linguistic items, treating these properties as if they were linguistic. Linguistification is widespread. It has reached epidemic proportions and needs to be eradicated.

¹ I think some expressions actually do refer, but this is a much narrower range of expressions than those that philosophers commonly call referring expressions. See the 'instrumental fallacy' at the beginning of Section 3.

This is important for two related reasons. First, attributing properties to linguistic expressions they don't have needlessly increases the explanatory burden on linguistics. The task of linguistics is demanding enough. Not only that, linguistification gives rise to a fundamental misconception of the process of ordinary communication. This process is not simply a matter of one person literally putting a thought into words and the other person decodes back into that thought. Much of what a speaker means in uttering something cannot be traced back to what the uttered words mean or even to their possibly context-sensitive semantic contents. Words carry information all right, but that information hardly determines what we mean in uttering them. Even when we aren't implicating anything entirely separate from what our words mean, we leave plenty of pragmatic gaps to be filled and sometimes even semantic slack to be tightened. It is important to recognize the extent to which we as speakers and hearers exploit extralinguistic information in communicating to and in understanding one another, and not to confuse this process with semantic context sensitivity.

Linguistification arises partly from overlooking or even trading on certain dangerous semantic-pragmatic ambiguities. Failing to keep these ambiguities in mind leads to classifying various pragmatic phenomena as semantic. Since these ambiguities, even if not familiar, are easy to see, my explanation of them will be brief. More interestingly, linguistification is also generated by way of various common fallacies. The bulk of this paper will be devoted to identifying these fallacies, diagnosing their sources, and explaining how they give rise to linguistification.

2 Linguistifying Ambiguities

There are many linguistically relevant terms that can apply either to language or to things speakers do in using it. As glaring as these ambiguities are, they are often neglected or conflated. It would be churlish to trot out illustrations. Instead, I will identify the main ones and hope that it is clear why they should not be ignored. Ignoring them can lead one to think that if a speaker does something in using a bit of language, then the linguistic item itself does the same thing, or at least something like it.

The most familiar such ambiguities involve the words 'meaning' and 'reference'. There is linguistic meaning, the meaning of an expression or an entire sentence, and there is speaker meaning, what a speaker means in using a particular expression or in uttering a sentence.² Needless to say, al-

² A speaker can also mean something in uttering several sentences, something over and above what she means in uttering the individual sentences. By the way, I will generally use the pronoun 'she' for the speaker and 'he' for the hearer.

though a speaker often means what the sentence she utters means (or, when there is ambiguity, one of the things it means), she might mean something else, or something more. As to reference, certain expressions refer, and speakers use expressions to refer. It is fairly rare for an expression that refers to one thing to be used by a speaker to refer to something else. Much more common is the case in which the expression does not refer at all, as in the case of definite descriptions, and yet the speaker uses the expression to refer.³ The phrase ‘referring expression’ can be misleading in this regard. Just as an eating utensil does not eat, so many so-called referring expressions do not refer.

The verb ‘say’ is notoriously ambiguous. Sometimes it is used such that to say something entails meaning it – in this sense it is roughly synonymous with ‘state’ and ‘assert’ – but in another sense it is used such that one can say something without meaning it (one might mean something else or perhaps nothing at all). Not only that, sometimes we describe sentences themselves, like billboards, clocks, and thermometers, as saying things.

Now consider the word ‘utterance’ and the phrase ‘utterance content’. An utterance can be the act of uttering some linguistic material, paradigmatically a sentence, or it can be the uttered sentence. But it is common to describe both the contents of speakers’ acts of uttering sentences and the semantic contents of sentence themselves as ‘utterance contents’. Clearly the fact that the same phrase is used to describe both does not mean that what they describe comes to the same thing. For example, if you utter the sentence ‘I’m sorry’ and mean that you are sorry for breaking your friend’s priceless vase, that proposition is the content of your apologetic act of uttering that sentence, not the content of the sentence you uttered, even as you are using it. It is important to appreciate that the sentence does not acquire an additional semantic property, thereby having richer propositional content, because of what you mean in uttering it. Using the phrase ‘utterance content’ for both the semantic content of the sentence and what the speaker means in uttering it can only lead to confusion by fostering the impression that sentence as uttered has a richer content than what is predictable from its meaning. Of course, who the speaker is and when the sentence is uttered are part of its semantic content as used in a context, but this is predictable from its meaning (from the meanings of its constituents ‘I’ and ‘now’ in particular). That breaking the vase is what the speaker is apologizing for is not predictable in this way. It is something the speaker means, not something the sentence means. After all, the speaker did not say, ‘I am sorry *for breaking your priceless vase*’.

³ I am assuming that Russell was basically right about definite descriptions, i.e., that what they denote are not constituents of the propositions expressed by sentences in which they occur.

The word ‘interpretation’ involves an allied ambiguity, whose perniciousness is compounded in the phrase ‘utterance interpretation’. This phrase can refer to what a speaker means in uttering a sentence or, indeed, to the psychological process whereby a hearer tries to figure out what the speaker means. However, this phrase is sometimes also used as if it meant something more abstract, something akin to semantic interpretation. It makes perfectly good sense to ascribe semantic interpretation, like phonological form and syntactic structure, to sentences, but not if this is construed as something richer than a projection of syntactic structure, something that incorporates what a speaker means in uttering the sentence on a particular occasion. That is not semantic interpretation, the output of a mapping from syntactic structure. This mapping is a ‘process’ only in an abstract sense. It is not a psychological process whereby a hearer tries to figure out what the speaker means, and it is not the output of that process. A sentence, as used by a speaker has a certain semantic interpretation, but that should not be identified with what the speaker means in uttering it or with the process of identifying that.⁴

3 Linguistifying Fallacies

The most instructive way to appreciate the perils of linguistification is to take note of the various fallacies associated with it. I will present and discuss the most important ones I’ve noticed, and illustrate the problems they give rise to. Each is formulated as a conditional, and the fallacy is to suppose that the antecedent implies the consequent. This leaves open the possibility that the consequent is true anyway, but showing that requires arguing for it on some additional or distinct basis. In my view, the consequents are all false, although some special cases of some of them may well be true.

Instrumental fallacy: If speakers use an expression to do something, then the expression does that thing (or something similar).

Some of us use chopsticks to eat sushi, but chopsticks don’t eat sushi. So if we use an expression to refer something, we shouldn’t infer that the expression itself refers to it. The expression might not refer at all. Call this instance of instrumental fallacy the ‘referential fallacy’. Somewhat similarly, if we use an expression to mean something, we shouldn’t infer that the expression itself means what we mean in using it. However, this ‘meaning

⁴ Akin to the ambiguity of the word ‘interpretation’ is that of the word ‘determination’, which, as we will see in the discussion of the ‘evidential fallacy’ below, can mean either what fixes content (whether that of a linguistic expression or of a speaker’s communicative intention in using the expression) or the process of figuring out what is meant.

fallacy' is not entirely analogous to the referential fallacy, since expressions in general do have meanings.

One common case of the referential fallacy deserves special mention, since it trades on a particularly pernicious ambiguity. The phrase 'demonstrative reference' can mean either reference by a demonstrative or speaker's reference by means of demonstrating while using a demonstrative. It is commonly supposed that if a speaker uses a demonstrative to refer to something, specifically something that she is demonstrating, then the demonstrative itself refers to that thing. This has to be argued for, not just assumed (unless it is construed as just trivially true, in the way that if you use a hammer to hammer the nail, the hammer hammers the nail).

Context-sensitivity fallacy: If something is done in context, it is done by context.

Expressions of many sorts have been claimed to be context-sensitive. These include gradable adjectives ('tall', 'old', 'smart'), relational adjectives ('local', 'legal', 'qualified'), relational nouns ('neighbor', 'mentor', 'enemy'), predicates of personal taste ('fun', 'funny', 'delicious'), quantificational words and phrases ('everyone', 'few students', 'the culprit'), environmental reports ('It is raining', 'It is noisy'), and terms that can be used without complements ('ready', 'late', 'finish'). However, proponents of such claims generally neglect to explain how the meaning of an expression alleged to be context-sensitive determines the expression's semantic content in a given context. They tend just to assume that if sentences containing such an expression can be used to assert different things in different contexts, it must be the presence of the expression (and presumably its meaning) that explains this contextual variation. The story is pretty straightforward with indexicals like *I* and *now*, but it is not obvious, to say the least, how that story generalizes. Instead, philosophers just assume that the operative contextual factor is what is relevant or salient in the context, or is perhaps what the speaker intends. They do not explain how the linguistic meaning of the expression in question determines this.

Proposition fallacy: If a sentence does not semantically express a proposition independently of context, then it expresses a proposition relative to a context, in which case it is context-sensitive.

A great many sentences, such as 'Jerry is ready', 'Tom is tall', and 'Leaves are edible', do not express a proposition independently of context. It does not follow that such a sentence expresses a proposition relative to a context, for it may not express a proposition at all. Many supposed cases of context sensitivity are really instances of something else: *semantic incompleteness*. A semantically incomplete sentence falls short of fully expressing

a proposition, and is therefore incapable of being true or false, even relative to a context.⁵

The proposition fallacy arises from the false assumption, perhaps stemming from the grammar school dictum that every complete sentence expresses a complete thought, that every declarative sentence must semantically express a proposition, if not independently of then relative to context. There is a truth underlying this assumption, but it pertains to speakers, not to sentences. If a speaker assertively utters a declarative sentence, presumably what the speaker means is a proposition. But this does not require that the sentence the speaker utters expresses that or any other proposition, even relative to the context.

The proposition fallacy and the context-sensitivity fallacy obviously go hand in hand. If one thinks that all well-formed declarative sentences express propositions but that only some do so independently of context, it is natural to suppose that the others do so in a way that depends on the context, and to leave the matter at that. Of course, what a speaker means in uttering a given semantically incomplete sentence must be a proposition and which proposition that is can vary from one context to another, but this does not mean that the context somehow determines, along with the meaning of the sentence, the proposition that the speaker means in uttering that sentence. That happens in context all right, but it is not done by context.

Evidential fallacy: If some contextual fact or consideration provides evidence that helps a listener determine an aspect of what is said or meant, then it helps determine that aspect of what is said or meant.

This remarkably widespread fallacy is illustrated by the common misinterpretation of Grice's theory of conversation according to which the Cooperative Principle and its attendant conversational maxims determine what a speaker implicates. Supposing that they do blatantly trades on the ambiguity of the word 'determine', which can mean either *ascertain* or *fix*. No one would fall for it if it read thus: 'If some contextual fact or consideration provides evidence that helps a listener ascertain an aspect of what is said or meant, then it helps fix that aspect of what is said or meant'. No one would

⁵ The idea of semantic incompleteness is straightforward if you think in terms of structured propositions rather than truth conditions, as built up out of objects, properties, and relations. Since these are made up of building blocks assembled in a particular way, it makes sense to suppose that in some cases such an assemblage, put together compositionally from a sentence's constituents according to its syntactic structure, might fail to comprise a proposition. That is because, although this comprises the entire semantic content of the sentence, it lacks at least one constituent needed for it to be true or false and to be the content of a thought or a statement. Even so, a semantically incomplete sentence, like a mere phrase, can be used to *assert* a proposition.

suppose that, for example, the maxim of relevance makes it the case that a superficially irrelevant remark is in fact relevant (if construed in a certain way). Implicatures are things speakers mean, and what a speaker means is a matter of her communicative intention, not a matter of anything imposed by the maxims. Their role is different.

Similarly, it is often supposed that when a speaker uses a demonstrative to refer to certain relevant or salient object, the relevance or salience of this object makes it the case that the demonstrative (as used by the speaker) refers to that object. If I say, for example, ‘That must be a Ferrari’, while looking in the general direction of a bright red sports car parked on the side of a road along with many other cars, its relevance or salience may help you figure out which car I am using ‘that’ to refer to, but that is not what makes that car the referent. What does that is my intention to refer to that car. To be sure, your recognition of my intention depends on that car being the salient and obviously relevant object (I didn’t bother saying ‘That *bright red car* must be a Ferrari’). You rely on this fact to recognize my intention. Indeed, the rationality of my referential intention requires that I can reasonably expect it to be recognized, and that depends on my having reason to think that what I am intending to refer to be distinguishable by you in some way, such that you can reasonably take me to be talking about it rather than anything else (I didn’t even say ‘that car’). But it is not part of the meaning of ‘that’ that it be used to refer to something relevant or salient. There is no need to suppose that its meaning imposes this constraint on my rational use of it, since in using it I could not reasonably intend to refer to a particular object without expecting you to recognize which thing I am referring to. And that expectation is reasonable only if I take that object to be relevant or salient.

Intuition fallacy: If a certain proposition is the ‘intuitive truth-conditional content’ of an utterance, then it is the semantic content of the uttered sentence.

Intuitions should be respected but not worshipped. The fallacy here is the result of assuming that the intuitions in question are responsive to the semantic content of a given sentence rather than to the truth-conditional content of a likely utterance of that sentence, i.e., to what the speaker means in uttering it. The fact that the intuition is dubbed ‘semantic’, just because it concerns the truth-condition of something, does not mean that it is in fact responsive solely to the semantic content of the uttered sentence.

Take the sentence ‘Gene hasn’t gone to work’. We imagine what we would take someone to mean, or perhaps what we would mean, in uttering the sentence. Then we ‘intuit’ its semantic content as the proposition that Gene hasn’t gone to work today, as if the time restriction is part of the se-

semantic content. This can't be right, since there is nothing in the sentence to impose that restriction. What happens when we consider a sentence in isolation is that, relying on our knowledge of the world and of people's typical communicative aims, we tend to make certain default assumptions about the circumstances of utterance. So our supposedly semantic intuition is insensitive to the difference between the semantic content of the sentence and what it is most likely to be used to convey. This is just what you'd expect insofar as speakers typically don't make fully explicit what they mean and have good reason to expect it to be obvious what they leave out. Indeed, making an obvious aspect of what one means would be misleading.⁶

Hidden indexical/covert variable fallacy: if an expression's use needs to be more specific than its meaning provides, or if its scope needs to be restricted somehow, it has an hidden indexical or a covert variable associated with it, whose value is somehow 'supplied by context'.

Here the fallacy is based on the assumption that the intuitive truth-conditional content of an utterance must be fully accounted for by that of the sentence being uttered. And that is tantamount to assuming that unarticulated constituents of the proposition that the speaker is trying to convey can only be semantic values of unpronounced ('phonologically null') constituents of the uttered sentence. Hidden indexical theories of various linguistic phenomena, such as propositional attitude ascriptions, illustrate this fallacy, at least when based on nothing more than appeal to intuitions about truth-conditional contents. So-called 'quantifier domain restriction' is another example of this fallacy. If you say 'We can start the meeting – everyone is here' and mean that everyone in the department is there, the fallacy is to assume that this restriction must be syntactically marked. A further part of the fallacy is to assume that context somehow 'supplies' or 'provides' this restriction (by determining the value of a hidden variable), when in fact it is the speaker's intention that restricts how the quantifier is being used.

It would be a mistake to insist that there is never any basis for positing phonologically null constituents. The history of modern syntax provides plenty of good reasons for doing so in certain cases. The fallacy is to do it without any independently motivated syntactic basis.

Lexical fallacy: If to be F is to be F relative to something (of a certain type), then the lexical entry for 'F' must associate a variable (or slot) with 'F'.

Consider adjectives like 'relevant', 'qualified', and 'legal'. A topic can't be just plain relevant, a person can't be just plain qualified, and an action can't be just plain legal. This does not show that the lexical entries for these

⁶ It would invite an inference as to why the speaker bothered to do so, and thereby suggest that something was meant that would normally not be.

adjectives must associate them with a variable (or slot) which, when given a value (or filled), specifies that for which the topic, person, or action is relevant, qualified, or legal. It is a metaphysical fact about being relevant, being qualified, or being legal that a topic can't be just plain relevant, a person can't be just plain qualified, and an action can't be just plain legal. That is, being relevant, being qualified, and being legal are not properties but what might we might call *property functions*. Again, these are metaphysical facts, not lexical facts.

The lexical fallacy is easy to fall into if you assume that the lexicon has to incorporate metaphysics. But why should it? Whether 'legal' expresses a property or a property function is a matter of what being legal is, not a matter of what 'legal' means.

The next three fallacies are closely related to hidden indexical and the lexical fallacies (and are similar to one another), but do not involve attributing gratuitous syntactic or lexical properties to the expressions in question.

Self-explication fallacy: If an expression 'F' is used to mean 'F + C', then it means (or has as its content) F+ C in that context.

Some examples will illustrate how easy it is to fall into this fallacy. Suppose I am talking about my 8-yr. old grandson and say, 'Billy is tall'. I am likely to mean that Billy is tall for an 8-yr. old, or perhaps that he is tall for his age. But surely the word 'tall' does not mean anything like *tall for an 8-yr. old*, not even in that context. If it meant that in that context, then what would 'tall for an 8-yr. old' mean in that context? *Tall for an 8-yr. old*? The additional element of what I mean, *for an 8-yr. old*, is not part of the content of 'tall', not even as used in that context. If it did, then if I had instead uttered 'Billy is tall for an 8-yr. old', the whole sentence would mean *Billy is tall for an 8-yr. old for an 8-yr. old* in that context. So the word 'tall' is not the locus of additional meaning *for an 8-yr. old*.⁷

Similar points apply to terms like 'edible', 'terrifying', and 'obvious'. What is edible to a goat or to a pig might not be edible to a human, but that doesn't mean that 'edible' has a different meaning depending on whether, in saying 'Leaves are edible', a speaker means that leaves are edible to goats, that leaves are edible to pigs, or that leaves are edible to humans. The word 'edible' means the same thing in each case. What varies is what the speaker means in uttering the sentence. If, as used in a certain context, 'edible' really meant, say, 'edible to goats', then if the speaker had uttered 'Leaves

⁷ The word 'tall' is also vague, but whether it is context-sensitive because it is vague is another question. If we were to address this question, we would have to keep in mind the observation that variation over contexts does not automatically indicate context sensitivity. The context itself, not what happens in context, has to be what makes the difference.

are edible to goats' in that context, the whole sentence would nonsensically mean *Leaves are edible to goats to goats*.

Property fallacy: If what a relative predicate 'F' applies to can vary from one context to another, then something can be F relative to one context and not-F relative to another, in which case what it is to be F can vary from one context to another.

Claims of this form are commonly made about what is expressed by relative adjectives like 'rich', 'old', and 'boring'. It is said that what 'counts as' being old, for example, depends on the context. A 39-yr old NBA basketball player counts as old, it is said, whereas a 39-yr old CEO does not count as old. Obviously there is a difference between being old for an NBA player and being old for a CEO, but it is misleading to describe either one as something that counts as being old. The phrase 'is old' simply does not express a property, much less a property that one person can possess and another person of the same age does not possess. To be sure, the NBA player has the property of being old for his profession and the CEO lacks the property of being old for his profession, but being old for one's profession is not a property; or, if it is, it is not the property that 'is old' expresses.

Relativist fallacy: If something cannot be F full stop but F only relative to something, such as a standard, a person, or a perspective, then the proposition that a sentence of the form 'a is F' expresses cannot be true or false full stop but true or false only relative to a contextually provided standard, person, or perspective.

Examples of things that something cannot be full stop include being rich, being tasty, and being to the right. These putative properties are ascribed in sentences like 'Gingrich is rich', 'Spinach is tasty', and 'The bathroom is on the right'. Now according to a contextualist account, such a sentence expresses a different proposition in different contexts, and this requires claiming that the property being ascribed depends somehow on the context (this claim relies on the context-sensitivity fallacy, among others). A relativist account of such a sentence claims that it expresses the same proposition regardless of the context – it's not the propositional content that varies – but that this proposition's truth-value depends partly on some factor determined by the context (this could be the context of assessment rather than the context of utterance).

Why this is a fallacy should come as no surprise. In fact, the relativist fallacy can be construed as a version of either the proposition fallacy or the property fallacy. On one interpretation, relativism assumes that the relevant sentences express propositions even though they omit mention of a standard, a person, or a perspective. But these are not ordinary propositions,

since their truth or falsity does not depend solely on what the world is like. On the other interpretation, what varies with the context (of utterance or of assessment, as the case may be) is what counts as having the relevant property, and that depends on the standard, person, or perspective that is somehow contextually determined. But this is tantamount to treating a property function as if it were a property. The way to avoid the relativist fallacy is to notice that sentences like ‘Gingrich is rich’, ‘Spinach is tasty’, and ‘The bathroom is on the right’, which do not express propositions or ascribe properties, can be turned into more expansive sentences that do specify the relevant standard, person, or perspective. A more expansive sentence, such as ‘Spinach is tasty to me’, does express a proposition and does ascribe a property.

4 Summing Up

I do not wish to exaggerate what has been shown here. The various ways in which one can be lured into one sort of linguistification or another were identified by means of a fallacy. In each case a natural but dubious conclusion about language was drawn from a true if not platitudinous premise, in most cases pertaining to language use. These conclusions were called into question but not decisively shown to be false. But that is enough to suggest at least that anyone wishing to endorse one or another of these conclusions should go to the trouble of arguing for it. One shouldn’t blandly (or blindly) assume that expressions do what speakers use them to do, that every well-formed declarative sentence expresses a proposition at least relative to a context, that variations in context or in values of covert variables account for apparent variations in content, that one-place predicates invariably express properties, etc. It’s a good idea to keep in mind that for lack of a constituent of one sort or another, a great many ordinary sentences do not make fully explicit what we do mean or could mean in using them. This does not show that they contain more semantic information than meets the eye (or ear) but only that we use them, and in many cases must use them, to convey more than the information they contain.