
Minding the Gap

KENT BACH

1 The Semantics-Pragmatics Gap

I'm all for pragmatics. But I don't think it helps the cause to blur the boundary between semantics and pragmatics. The purpose of this chapter is to explain why.

The basic reason is this. Even though, as people have been pointing out for some years now, the linguistic meaning of a given sentence generally underdetermines what a speaker means in uttering it, it does not follow that linguistic meaning is infected or infested by what some of these same people call 'pragmatic meaning'. There is no such thing as pragmatic meaning, at least nothing that is commensurate with linguistic meaning. There is what the sentence means and what the speaker means in uttering it.

The semantic-pragmatic distinction is not fit to be blurred. What lies on either side of the distinction, the semantic and the pragmatic, may each be messy in various ways, but that doesn't blur the distinction itself. Taken as properties of sentences, semantic properties are on a par with syntactic and phonological properties: they are linguistic properties. Pragmatic properties, on the other hand, belong to acts of uttering sentences in the course of communicating. Sentences have the properties they have independently of anybody's act of uttering them. Speakers' intentions do not endow them with new semantic properties (here I mean sentence types, not tokens). Acts of uttering sentence types (producing sentence tokens) have pragmatic properties. The fact that a given sentence means what it does entails nothing about what a speaker means in uttering it. A speaker could mean precisely

what it means, no more and no less, but nothing about its meaning guarantees this. The speaker might mean something else, something more, or nothing at all.

In what follows I will sketch a picture different aspects of which I have addressed in more detail in other places. In those places I have used illustrative examples and discussed alternative approaches to specific issues.¹ Here I will simply state my position and summarize the arguments for it. This will make the presentation more abstract (and dogmatic) but also more compact, so that the picture can be viewed more clearly as a whole.

2 Why There is a Gap

There are many phenomena that are thought to show that the semantic-pragmatic distinction is blurry, arbitrary, or even nonexistent. Here I will not review these phenomena or rebut these arguments but will instead suggest that such phenomena fit in with the distinction and that the distinction itself is immune to such arguments – at least if it is properly understood as the innocuous (but important) distinction that it is. As will be explained in subsequent sections, these arguments rest on blurring or disregarding certain other distinctions which, at least when pointed out, are too obvious even to need defending. Once we are clear on those other distinctions, we can stop worrying about the semantic-pragmatic distinction itself. This is not to say that there are not plenty of pertinent linguistic phenomena to worry about, but once the line between semantics and pragmatics is clear, there should generally be no question on which side of the line a given phenomenon falls. There may be some residual hard cases, linguistic phenomena that are not well-understood, and in some cases, there may be both a semantic and a pragmatic side to the phenomenon, but these facts do not create a problem with the semantic-pragmatic distinction itself.²

The reason there is a gap and a clear-cut one at that is that semantics and pragmatics have distinct subject matters, sentences and utterances, respectively. Semantics is the part of grammar that pairs forms with meanings. Presumably the meaning of a sentence is determined compositionally by the meanings of its constituents as a function of its syntactic structure. Pragmatics is concerned with what speakers do in uttering sentences. So if semantics concerns sentences, in particular their linguistic meanings, and pragmatics concerns uses that speakers make of sentences (and, correla-

¹These other places will be cited in subsequent footnotes. The semantic-pragmatic distinction itself is discussed in Bach 1999a.

²For discussion of some seemingly problematic phenomena, see Bach 1999a: 75-80, and for discussion of some genuinely problematic phenomena see Bach 2001a: 31-40.

tively, hearers' understanding of speakers' uses), then, it would seem, the semantic-pragmatic distinction should hold up without incident.

Even so, there are certain linguistic phenomena that may seem to threaten the gap. Clefts, stress, and other devices for packaging or structuring information (see Lambrecht 1994), grammatical mood (imperative, interrogative, etc.), and utterance modifiers are examples of linguistic expressions and devices which, though they have no bearing on truth conditions, nonetheless semantically encode information that bears on use.³ Is such information semantic or pragmatic? If such phenomena seem to threaten the gap, that is only because there is a certain ambiguity involving the term *semantic*. It can pertain specifically to linguistic information relevant to the truth conditions of a sentence, or it can be used more inclusively and concern any matter of linguistic meaning. The examples just mentioned have semantic content in the latter, more inclusive sense. Indeed, they pertain to use. Does that blur the semantic-pragmatic distinction? Not really. If a speaker uses the imperative form or a cleft construction, for example, there is still the distinction between the information encoded by the form or construction and the fact that the speaker uses it. Such semantic information is in a sense procedural, but that does not make it pragmatic in the sense of being a property of the speaker's act of utterance.

Even if what a speaker means consists precisely in the semantic content of the sentence he utters, this fact is not determined by that semantic content. The reason is very simple: no sentence has to be used in accordance with its semantic content. *Any* sentence can be used in a nonliteral or indirect way. That a speaker is attempting to communicate something, and what that is, is a matter of his communicative intention (if indeed he has one). If he is speaking literally and means precisely what his words mean, even that fact depends on his communicative intention. But what he utters, if a grammatical sentence, has the semantic content that it has whether the speaker is using it in a strictly literal way or not.

In any case, it is not the business of semantics to account for the contents of utterances that are not literal, since in those cases the speaker is trying to convey something that is not predictable from the meaning of the uttered sentence (or, if it is ambiguous, from its operative meaning).⁴ Obviously not just anything a speaker means, no matter how far removed it is from what the sentence means, counts as semantic content, and the semantic

³Utterance modifiers, such as *by the way* and *frankly*, are discussed in Bach 1999b: 356-60.

⁴Disambiguation is a pragmatic matter, part of the process whereby the hearer figures out what the speaker means in uttering the sentence. The fact that a sentence is ambiguous and that someone who utters it intends it to be read in one way rather than another does not imply that in the context the sentence loses its other meaning(s).

content of the sentence is the same whether or not an utterance of it is literal. Semantic content is a property of the sentence, not the utterance.

The term *utterance* is often applied to sentence tokens, but there is no need to attribute semantic content to them (apart from that of the sentence types of which they are tokens). It is as an intentional act performed by a speaker that an utterance has a content distinct from that of the uttered sentence. But in this sense the content of an utterance is really the content of the speaker's communicative intention in making the utterance. So the only linguistic content relevant here is the semantic content of the sentence, and the only other relevant content is the content of the speaker's intention. Focusing on the normal case, where communication is successful (the listener succeeds in identifying the speaker's communicative intention), can make it seem as though an utterance has content in its own right, independently of the speaker's intention. But this is illusory, as is evident whenever communication fails. In that case, where the speaker means one thing and his audience thinks he means something else, there is what the speaker means and what his listener takes him to mean, but there is no independent utterance content.

In practice, we generally don't say what precisely we mean, because we leave much of what we mean to inference. Even when we are using words or phrases literally, we generally do not use the entire sentence literally. That is, what the sentence means, and what we thereby say in uttering it, comprises merely a skeletal version of what we mean. For example, if at a McDonald's you say, 'I want a hamburger with everything', presumably you mean that you want to be given, within a short time, a cooked, ready to eat, uncontaminated hamburger of normal size with all the trimmings. But you do not say most of this, and the sentence you use is not an elliptical version of some more elaborate sentence that spells these things out.⁵

This phenomenon of *sentence nonliterality* is so pervasive that we tend not to notice it, not just when engaged in ordinary conversation but even when theorizing about language and communication. We generally do not make fully explicit what we mean in uttering a sentence, even when we are using the individual words in them literally. Although for every sentence we do utter, there is a more elaborate, qualified version we could utter that would make what we mean more explicit, these are not the sentences we do utter. Indeed, they are not ones we even think to utter. We do not form a thought to express, think of a convoluted sentence that would express it fully explicitly, and then, in the interests of conversational efficiency, work out a stripped down version of the sentence to utter instead. Yet somehow

⁵ For a variety of further examples and further discussion, see Bach 1994 and 2001b.

we often manage to make evident our relatively complex communicative intentions by uttering relatively simple sentences.

3 Communicative Intentions and Inferences

Austin (1962: 94-107) drew a fundamental distinction between *locutionary* and *illocutionary* acts. This distinction is easy to disregard, partly because the word ‘say’ plays a dual role as a locutionary and as an illocutionary verb, roughly synonymous with ‘state’ (or ‘assert’). A locutionary act is the act of saying something in uttering a sentence. In this sense, one can say something without stating it. One might not be stating anything, or one might be speaking figuratively and be stating something else, though not expressly. The illocutionary act a speaker performs in saying something is a matter of his communicative intention.⁶ What he means in saying what he says (if he means anything at all) may be just what he says, what he says and more, or something else entirely. What he says provides the linguistic contribution to the audience’s inference to his communicative intention. Given what one intends to communicate, succeeding at communicating requires choosing a sentence whose utterance will make one’s intention evident. One must utter a sentence which, given the mutually salient information that comprises the extralinguistic cognitive context of utterance, enables one’s audience to recognize one’s communicative intention. There is something distinctive here about a communicative intention. Unlike intentions of other sorts, its fulfillment consists in its recognition. Whereas you can’t do a pirouette by way of anyone’s recognizing your intention to do so, you can succeed in communicating something simply by enabling your listeners to recognize your intention to communicate it.

In uttering a sentence, a speaker normally expresses an attitude toward a proposition. For example, asserting is to express a belief, requesting is to express a desire, and promising is to express an intention. One does not just express a proposition. Here is a Gricean conception of what is it to express an attitude:

Expressing an attitude: In making an utterance S expresses A(p) iff S intends H, partly by way of recognizing that this (entire) intention is to be recognized, to take the utterance as a reason to think S has A(p).

This definition, a slightly modified version of the formulation given by Bach and Harnish (1979: 15), allows for the possibility that the speaker

⁶Here I depart from Austin (1962), who assimilated all illocutionary acts to the regularized ones performed in special institutional or social settings. I follow Strawson (1964), who argued that most ordinary illocutionary acts are not conventional but communicative in character and succeed not by conformity to convention but, in Gricean fashion, by recognition of intention.

does not actually have the attitude he is expressing. After all, he could be lying (in the case of an assertion) or be otherwise insincere. This definition incorporates Grice's (1957) idea that the hearer is to rely partly on the fact that the speaker intends him to recognize the intention. This does not tell her what the intention is (that would introduce not just reflexivity but circularity into Grice's picture), but it does crucially constrain the inference involved in recognizing that intention.

Speakers and listeners rely on certain presumptions, speakers in order to make their communicative intentions evident and listeners to recognize those intentions. The listener presumes, and the speaker expects him to presume, that the speaker is being cooperative and is speaking truthfully, informatively, relevantly, perspicuously, and otherwise appropriately.⁷ Because of their potential clashes, these presumptions should not be viewed as comprising a decision procedure. Rather, they provide different dimensions of considerations for the hearer to take into account in figuring out the speaker's communicative intention. They ground strategies for a speaker, on the basis of what he says and the fact that he says it, to make what he means evident to the hearer, and correlative strategies for the hearer to figure out what the speaker means. If an utterance superficially appears not to conform to these presumptions, the listener looks for a way to take the utterance so that it does conform. She does so partly on the supposition that she is intended to do so. However, contrary to the popular misconception that the maxims or presumptions play a role only in implicature (and in oblique and figurative speech generally), they are operative even when one means just what one says.

It is easy to exaggerate how well linguistic communication works. No doubt many failures of communication do not get noticed and, when they do, do not get mentioned. Still, in order for it to work as well as it does, we as speakers must be very good at selecting sentences whose utterance makes evident to our listeners what we mean, and as listeners we must be very good at figuring out what speakers mean in uttering them. To the extent that we leave much of what we mean to inference, we rely on our listener's ability to figure out what we mean on the basis of what we say, given the circumstances in which we say it, the fact that we said that rather than something else, and the presumption that we said it with a recognizable intention. How we manage to make ourselves understood and how we manage to understand others are very complex processes which, like most cognitive processes, are far beyond the reach of contemporary psychology to

⁷These presumptions correspond to Grice's maxims (1989: 26-7). Construing them as maxims treats them as guidelines for how to communicate successfully, but I think it is better to construe them as presumptions about speaker's intentions.

explain. As theorists, the best we can do is speculate on some of the features of these processes, most notably what it takes to implement Grice's discovery that communication involves a distinctively reflexive intention (and its recognition). Specifically, the intention includes, as part of its content, that the audience recognize this very intention by taking into account the fact that they are intended to recognize it.⁸

4 Muddling the Gap: Neglected Ambiguities

There are various key terms which, if not literally ambiguous, are at least used in two distinct ways. In each pair in the list below, the first use is semantic in character and the second is pragmatic. It is important to be cognizant of the two contrasting uses of each term.

meaning

- linguistic meaning: sense of an expression (word, phrase, or sentence)
- speaker's meaning: what a speaker means

speaker's meaning

- what a speaker means by a sentence (or phrase) when using it
- what a speaker means (tries to communicate) in uttering a sentence

reference

- by an expression to an object
- by a speaker with an expression to an object

utterance

- what is uttered
- act of uttering

utterance meaning

- meaning of an uttered sentence
- speaker's meaning in uttering a sentence

say

- perform a locutionary act
- state or assert, especially in using a declarative sentence without using any of its constituent expressions nonliterally

what is said

- the content of a locutionary act (or equivalently, the semantic content of sentence, relative to a context of utterance)

⁸The concepts in this section are explained much more fully in Bach forthcoming b. They were developed originally in Bach and Hamish 1979.

- the content of the assertion made in using an declarative sentence without using any of its constituent expressions nonliterally

context

- set of parameters whose values fix or delimit the semantic values of expressions with variable references
- set of salient mutual beliefs and presumptions among participants at a stage in a conversation

determine

- make the case (constitutive determination)
- ascertain (epistemic determination)

interpretation

- assignment of semantic values
- inference to speaker's communicative intention

demonstrative reference

- reference by a demonstrative
- speaker's reference by means by demonstrating

use (a term) to refer

- use a term that refers
- use a term and thereby refer

Most of these distinctions are self-explanatory, but a few require explanation and justification. As we will see in the next two sections, blurring these distinctions is essential to the various closely related doctrines that go by such names as *contextualism*, *truth-conditional pragmatics*, and *linguistic pragmatism*.

As mentioned above, Austin's distinction between locutionary and illocutionary acts, between saying something and doing something in saying it, is commonly neglected these days, perhaps because it is so easy to use 'say' interchangeably with 'state' or 'assert'. But stating or asserting is to perform an illocutionary act, of meaning and trying to communicate something, and that goes beyond mere saying (in the locutionary sense). Why is the locutionary notion of saying needed, along with the correlative, strictly semantic notion of what is said? It is needed to account for the each of the following cases, situations in which the speaker:

- says something but doesn't mean anything at all (by 'mean' here I mean 'intend to communicate')
- does not say what he intends to say, as in the misuse of a word or a slip of the tongue

- means what he says and something else as well (cases of implicature and of indirect speech acts in general)
- (intentionally) says one thing and means something else instead (non-literal utterances)

These are all cases in which what the speaker means, if indeed he means (intends to communicate) anything, is not identical to what he says, but cases in which he still says something.

What is loosely called ‘context’ is the conversational setting broadly construed. It is the mutual cognitive context, or salient common ground. It includes the current state of the conversation (what has just been said, what has just been referred to, etc.), the physical setting (if the conversants are face-to-face), salient mutual knowledge between the conversants, and relevant broader common knowledge. As I will explain more fully in the next section, context does not determine (in the sense of *constitute*), but merely enables the hearer to determine (in the sense of *ascertain*), what the speaker means. Because it can constrain what a hearer could reasonably take a speaker to mean in saying what he says, it can constrain what the speaker could *reasonably* mean in saying what he says. But it is incapable of determining what the speaker actually does mean. That is a matter of the speaker’s communicative intention, however reasonable or unreasonable it may be.

The ambiguity of the phrase *utterance interpretation* is especially unfortunate. It already contains the ambiguous term *utterance*, which can mean what is uttered or the act of uttering it. Considered as the assignment of meanings to sentences, utterance interpretation is semantic interpretation, a mapping of sentence forms onto their linguistic meanings. In that sense, utterance interpretation is not a psychological process but something more abstract. However, this phrase is also used to mean the psychological process whereby listeners figure out what speakers are trying to communicate. This use would be unexceptionable but for the fact that this psychological process is often treated as if it were something more abstract, more akin to semantic interpretation. That is, some philosophers and linguistics treat utterance interpretation in this sense as if it were a mapping from syntactic structure to utterance contents, except that the mapping is sensitive to broadly contextual factors. In so doing, they seem to think that an utterance (as opposed to a sentence) can express things independently of what the speaker means in making it, just because of the context in which it is made. It is as if meanings could somehow be read off of utterances independently of inferring the speaker’s intention, in a way analogous to semantic interpretations of sentences but without the constraint of being a projection of syntactic structure.

5 Putting Context in Context

It is often casually remarked that what a speaker says or means in uttering a given sentence ‘depends on context’, is ‘determined’ or ‘provided’ by context, or is otherwise a ‘matter of context’. That is not literally true. Assume that by context we mean something like the mutually salient features of the conversational situation. Does context determine what a speaker says? Suppose a dinner host utters the ambiguous sentence *The chicken is ready to eat*. Presumably she is not saying that a certain chicken is hungry. Even so, given the ambiguity of the sentence, she could be saying that, however bizarrely. Context doesn’t make it the case that she does not. Of course, she could not reasonably expect such a communicative intention to be recognized.⁹

Demonstratives and indexicals are often casually described as ‘context-sensitive’ or ‘context-dependent’. These above point about context applies to these terms, although so-called pure or automatic indexicals, such as ‘I’ and ‘today’, are an exception. The reference of a pure indexical is determined by its linguistic meaning as a function of a contextual variable (call that the *semantic* context). However, the reference of other indexicals and of demonstratives is not so determined. They suffer from a *character deficiency*. For them the role of context is not to determine reference, in the sense of constituting it, but to enable the listener to determine it, in the sense of ascertaining it. But this is cognitive, not semantic context. Their reference is a matter of the speaker’s communicative intention, which is not just another contextual variable. In fact, rather than describing them as referring, it would be better to describe speakers as using them to refer.¹⁰

It might be supposed that to the extent that the meaning of a sentence does not determine the content of an utterance of the sentence, the speaker’s communicative intention does, and that this intention is part of the context. However, there is a problem with supposing that. If context were defined so broadly as to include anything other than linguistic meaning that is relevant to determining what a speaker means, then of course the speaker’s intention would be part of the context. But if the context is to play the explanatory role claimed of it, it must be something that is the same for the speaker as it is for his audience, and obviously the role of the speaker’s intention is not the same for both. Context can constrain what the speaker can succeed in communicating given what he says, but it cannot constrain what he intends to communicate in choosing what to say. Of course, in implementing his

⁹The uses and abuses of context are discussed more fully in Bach forthcoming a.

¹⁰Never mind that demonstratives and impure indexicals also have clearly non-referential uses, e.g., as something like bound variables.

intention (in order to get his communicative intention recognized), the speaker needs to select words whose utterance in the context will enable the hearer to figure out what he is trying to communicate.

6 Fudging the Gap: Semantic Incompleteness

Philosophers and linguists are coming to recognize that many if not most sentences do not express complete propositions, even when the sentence is devoid of ambiguity, vagueness, or indexicality. This is a very interesting fact, but it does not have the radical consequences some people think it has.

A sentence can fail to express a complete proposition either because it is missing an argument and is therefore *semantically incomplete* or because it is *semantically underdeterminate* and contains a word or phrase whose meaning needs to be tightened in one way or another.¹¹ Regardless of the reason that a given sentence fails to express a complete and determinate proposition, a speaker who utters such a sentence cannot mean merely what the sentence means. Presumably, what he means must be a complete proposition (something capable of being true or false).

This is an interesting fact, but it does not have radical consequences, either for semantics or pragmatics. Assume that the meaning of a sentence is determined compositionally, by the meanings of its constituents in a way that is predictable from how its constituents fit together syntactically. Even so, there is no guarantee that what is thus determined is a complete proposition. This may undermine the naive assumption that the output of a semantic theory for a language is a (recursive) specification of the truth conditions for all its (declarative) sentences, but it does not undermine semantic compositionality. It is just a brute fact about language that some syntactically complete sentences are not semantically complete.

Also, it might seem that if the semantics of sentences often fails to generate complete propositions, then pragmatics must intrude into semantics. The idea is that if a speaker utters a semantically incomplete or underdeterminate sentence and what he means must be a complete and determinate proposition, then the additional elements or tightening of meaning that is required to yield a proposition must automatically be incorporated in the semantic content of the utterance. This is the idea behind so-called truth-conditional pragmatics. It is also supposed that such sentences are in some sense context-sensitive or context-dependent. That leads to so-called contextualism in semantics, which not only supposes that semantics must deliver truth conditions but that context somehow comes to the rescue of semantics.

¹¹So-called scope ambiguities may in some cases really be instances of underdetermination.

The trouble with these ideas is that they needlessly confuse the semantic contents of sentences with the intentional contents of utterances, that is, of the attitudes people express in making utterances. The fact that a sentence is semantically incomplete or underdeterminate does not mean that it is semantically defective. Yes, it does have a character deficiency, but still it has a definite content (or contents, if it is ambiguous), as given compositionally by the semantics. There is no empirical or theoretical reason why, at the level of *semantics*, the content of the sentence has to be supplemented or refined.¹² The fact that a speaker cannot mean merely what the sentence means comes into play when the speaker utters the sentence, insofar as the speaker has some communicative intention in uttering it. And the hearer, taking into account that he is to recognize that intention, must recognize that there is a gap between the semantic content of the sentence and the intentional content of the utterance. But there is no need for a truth-conditional pragmatics to fill out the semantic content of the sentence by treating it as a ‘function’ of ‘context’. Pragmatic interpretation, as a cognitive process, cannot be understood as an abstract ‘process’, on the model of semantic interpretation. Rather, mutually evident facts enable the speaker, at least when communication goes well, to make his intention evident and enable the hearer to identify it.

Pragmatically speaking, when a speaker does mean merely what the sentence means, it makes little difference whether or not the sentence he utters expresses a complete and determinate proposition. Obviously, if the sentence is semantically incomplete or underdeterminate, the speaker *cannot* mean merely what it means, but it is generally as obvious that he *does* not mean what the sentence means even when the sentence does express a complete and determinate proposition. For it is generally true that the sentences we use do not make fully explicit what we mean. This is the norm, not the exception.

¹²I am not denying that with certain lexical items there may be syntactic or lexical reasons for supposing that they have an implicit argument position associated with them. However, such items are a source of hidden indexicality, not semantic incompleteness. The view that I am objecting to supposes that even in the absence of such an item, some sort of unarticulated constituent is present, mandated not by the grammar but by the need for a complete proposition to be expressed. In my view, this need is purely pragmatic. It cannot be assumed that since a sentence would not express a complete proposition unless it contained some hidden argument slot whose value is somehow provided contextually (or else the slot is quantified over), the sentence actually contains such an argument slot.

7 Missing the Gap: Faulty Intuitions

It is important to keep in mind that communication can fail and often does. We cannot just take a sample sentence and read off the likely communicative intention of someone who utters it, if only because they might be using it in some strange, perhaps even incomprehensible way. Displaying a sentence and putting a number in front of it does not make evident what a person means in uttering it. Moreover and more importantly, for a great many sentences it is not true that the most likely use of that sentence is a literal one. A great many sentences are very hard to use literally, even if one uses the individual words and phrases in them literally.¹³

In fact, there is nothing particularly normal about speaking literally. That is why it is naive to suppose that the semantic content of a declarative sentence (relative to a context) can be defined as what a speaker would normally assert in uttering the sentence. To suppose that for most sentences a likely use of it is a literal one can lead to erroneous semantic judgments. Relying on such ‘intuitions’ implicitly assumes that if a sentence would normally be used in a certain way, that use is a literal use. However, intuitions that purport to be about the semantic contents of sentences might in fact be responsive to what sentences are typically used to assert. Indeed, it is to be expected that these intuitions should be faulty, given our ability to say one thing and successfully convey something much richer and our correlative ability to recognize what others are doing when they are exercising that first ability. These abilities of ours are so fluent as to distort our reflective semantic judgments. What we think is expressed by a given sentence is colored by what we think a speaker is likely to mean in uttering it, especially if there is a typical sort of context which we tend to imagine it being used in. To keep one’s semantic judgments from being pragmatically contaminated, it is always a good idea to imagine a variety of contexts of use, even wildly improbable ones.¹⁴

Intuitions are tainted by the fact that when a sentence is considered in isolation, certain default assumptions are made about the circumstances of utterance. We read things into the meaning of a sentence or into what a speaker says in uttering it that are really consequences of its being uttered under normal circumstances. These default assumptions depend on our knowledge of the world and of people’s typical communicative purposes. So we tend not to discriminate between the semantic content of a sentence and the likely force of uttering a sentence with that content.

¹³As I have argued (Bach 2001b), even when one is using all the words in a sentence literally, one can still be speaking loosely, hence not strictly literally.

¹⁴For more on why seemingly semantic intuitions cannot be trusted, see Bach 2002.

Not only are these intuitions often responsive to nonsemantic information, they play no direct role in ordinary communication. In the course of speaking and listening to one another, we generally do not consciously reflect on the semantic content of the sentences we hear. We are focused on what we are communicating and on what is being communicated to us. Not only that, we do not have to be able to make accurate judgments about what information is semantic and what is not in order to be sensitive to semantic information.

Moreover, our seemingly semantic intuitions are responsive to *pragmatic regularities*. These include regularized uses of specific expressions and constructions that go beyond conventional meaning, as well as general patterns of efficient communication, which involve streamlining stratagems on the part of speakers and inferential heuristics on the parts of listeners. These regularities are pragmatic because it is the *use* of a given sentence, not the sentence itself, that carries the additional element of information. *That* a speaker says what he says rather than something else can contribute to what a speaker is likely to be taken to mean. Indeed, that he says it one way, by using certain words rather than certain others, can also contribute to what he is likely to be taken to mean in using those words.

Also, it is statistically unsound to focus on sentences representative of those we use or might use. That would be to commit a massive sampling error. English (or any other language) is just as capable of generating unusable sentences as usable ones. Even unusable sentences have perfectly good meanings, no matter how hard to parse or how absurd it would be to assert their semantic contents. Imagine that most of the (English) sentences that philosophers and linguists use as examples are representative of the vast variety of sentences that (English) speakers have used or might use. Even so, sentences actually used or even potentially usable are not representative of (English) sentences in general. That is not just because most sentences are far too long to be used in real life. Most sentences, even ones of fairly modest length, express things that are too bizarre ever to say, much less mean (the really long ones express things that are just too complex to understand).

This sampling error compounds the mistake of supposing that intuitive contents are semantic contents. Recognizing this error leads to a simple argument that encapsulates the disconnect between seemingly semantic intuitions and the semantics of sentences:

1. Most of the syntactically well-formed sentences of a natural language, even those of a reasonable length, do not express things we are ever likely to mean.

2. A natural language is not so designed that most of its sentences, even among those of utterable length, are sentences anyone would ever have occasion to use.
 3. But if the semantic content of a sentence is its 'intuitive content', then if it doesn't have an intuitive content, it doesn't have a semantic content.
 4. The contents of most sentences are just too weird to be their intuitive semantic contents.
 5. So, either most sentences either don't have semantic contents or their semantic contents are not their intuitive contents.
 6. All well-formed sentences, even useless ones, have semantic contents.
- ∴ The semantic contents of most sentences are not their intuitive contents.

8 Semantic Illusions

I have suggested that seemingly semantic intuitions are often illusory because they tend to be insensitive to the difference between semantic facts and pragmatic regularities. So we cannot rely on such intuitions too heavily and should not draw hasty conclusions about the semantic content of a given sentence just because we have imagined the likely import of a typical utterance of it. Also, seemingly semantic intuitions can combine with simplistic theoretical assumptions about semantics to yield a variety of semantic illusions:

The conventionalization illusion: that a regularized use of an expression or construction can only be a matter of meaning (or semantic convention). This illusion assumes that there cannot be pragmatic regularities, cases of standardization without conventionalization of use.

The assertion illusion: that the semantic content of a declarative sentence is what a speaker would be asserting (the content of the belief he would be expressing) if he sincerely uttered it and used all of its constituent expressions literally. This illusion conflates minimal assertions with locutionary acts.

The processing illusion: that a proposition does not qualify as what is said in the utterance of a sentence if it is not represented in the process of understanding the utterance. This illusion conflates information which is represented in a cognitive process with information which is accessible to the process and to which the process is sensitive.

The illusion of utterance (or token) semantics: that because sentences (generally) are neither true nor false independently of contexts of utterances, the proper subject of semantics is utterances, not sentences. This illusion trades on the ambiguity of 'utterance', which can mean either an uttered sentence or the act of uttering it. The problem is that the semantics of utterances in

the first sense is sentence semantics and that the study of contents of acts of utterance belongs to pragmatics. So there is no role for utterance semantics.

The said-if-not-implicated illusion: that if a speaker means something that he doesn't merely implicate, he must be saying it. This illusion treats saying and implicating as an exclusive dichotomy. It overlooks the fact that part of what a speaker means can be implicit in what he says or, rather, in his saying it. This illusion is aided and abetted by the rampant use in certain circles of the neologism *explicature*, which is a cognate of *explicate* but is nevertheless used as if it meant explicit content. An explicature spells out something that is not explicit in what is said. It makes explicit what isn't. A better term is *impliciture*.¹⁵

The illusion of context dependence: that when content varies from one context to another, the variation in content depends on features of the context. This illusion conflates contextual variability with context dependence.

The one sentence, one proposition illusion: that a declarative sentence always expresses (modulo ambiguity and indexicality) one and only one proposition. In fact, many sentences fail to express a complete proposition, and many sentences express more than one proposition.¹⁶

9 Bridging the Gap

The semantic-pragmatic distinction is a well-defined and theoretically warranted distinction. Maintaining it requires recognizing the limitations of semantics and the reach of pragmatics. Semantics concerns the meanings of sentences, but these often fail to determine complete propositions, even modulo ambiguity, vagueness, and indexicality. As long as it is not assumed that the job of semantics is to give truth conditions of (declarative) sentences, there is no reason to suppose that pragmatics needs to intrude on semantics. This is to be expected, since pragmatics is concerned with utterances of sentences, not with sentences themselves. The fact that a speaker utters a sentence plays a key role in what he can reasonably expect to communicate in uttering it and in what the listener can reasonably take him to be communicating. When a sentence is uttered, it does not encode the fact that it is uttered. This is essentially a pragmatic fact.

References

Austin, J. L. 1962. *How to do Things with Words*. Oxford: Oxford University Press.

¹⁵The notion of conversational impliciture was introduced and explained in Bach 1994.

¹⁶For a detailed diagnosis and discussion of the second part of this illusion, see Bach 1999b: 343-55.

- Bach, K. 1994. Conversational Implicature. *Mind & Language* 9: 124-62.
- Bach, K. 1999a. The Semantics-Pragmatics Distinction: What It Is and Why It Matters. *The Semantics-Pragmatics Interface from Different Points of View*, ed. K. Turner, 65-84. Oxford: Elsevier.
- Bach, K. 1999b. The Myth of Conventional Implicature. *Linguistics and Philosophy* 22: 327-66.
- Bach, K. 2001a. You Don't Say? *Synthèse* 128: 15-44.
- Bach, K. 2001b. Speaking loosely: Sentence Nonliterality. *Midwest Studies in Philosophy*, vol. 25: *Figurative Language*, ed. P. French and H. Wettstein, 249-63. Oxford: Blackwell.
- Bach, K. 2002. Seemingly Semantic Intuitions. *Meaning and Truth*, ed. J. K. Campbell, M. O'Rourke, and D. Shier, 21-33. New York: Seven Bridges Press.
- Bach, K. forthcoming a. Context ex Machina. *Semantics vs. Pragmatics*, ed. Z. Szabó. Oxford: Oxford University Press.
- Bach, K. forthcoming b. Speech Acts and Pragmatics. *The Blackwell Guide to the Philosophy of Language*, ed. M. Devitt and R. Hanley. Oxford: Blackwell.
- Bach, K. and R. M. Harnish 1979. *Linguistic Communication and Speech Acts*. Cambridge, MA: MIT Press.
- Grice, H. P. 1957. Meaning. *Philosophical Review* 66: 377-88.
- Grice, H. P. 1989. *Studies in the Way of Words*, Cambridge, MA: Harvard University Press.
- Lambrecht, K. 1994. *Information Structure and Sentence Form*. Cambridge: Cambridge University Press.
- Strawson, P. F. 1964. Intention and Convention in Speech Acts. *Philosophical Review* 73: 439-60.