

# On Failing to Capture Some (or Even All) of What is Communicated

Kent Johnson

**Abstract** This paper examines a methodological argument launched against Cappelen and Lepore’s “minimalist” semantics. The charge is that this semantic theory – and by implication a great many other ones – cannot be correct, because they fail to capture some of the “intuitive truth conditions” of the relevant sentences. In response, I argue that this charge rests on the claim that an acceptable scientific theory must (at least sometimes) capture all of the overt phenomena under study. But this claim, I contend, is false. In actual practice, scientific models will often never capture all of the behavior of the relevant phenomenon, and this feature does not undermine them as such. I maintain that semantic theorizing is just an instance of this more general aspect of scientific methodology.

**Keywords** Linguistics · semantics · model · error · methodology

## 1 Introduction

Recently, Herman Cappelen and Ernie Lepore (hereafter CL) have developed and defended a particular view, semantic minimalism, regarding the general structure and scope of semantic theories for natural languages (CL, 2005). According to semantic minimalism, semantic context-sensitivity in our expressions is limited to a handful of expressions (e.g., *I*, *you*, *she*, *this*, *that*, *tomorrow*, etc.). In particular, there do not exist unobserved (i.e., phonologically and orthographically unrealized) context-sensitive elements in either the syntactic or semantic structure of an expression. As CL show in great detail, semantic minimalism contrasts sharply with virtually every other existing view about semantic theory. In particular, it is extremely common for semanticists and philosophers of language to assume that a correct semantics for (1a) and (2a) assigns the additional structure in (1b) and (2b) as part of the correct meaning of these sentences:

---

K. Johnson  
Department of Logic and Philosophy of Science, University of California, Irvine, CA 92697-5100,  
USA

- (1) a. Mary is ready;  
b. Mary is ready *to X*.
- (2) a. It is raining;  
b. It is raining *in location X*.

According to the standard view, *Mary is ready* means something like Mary is ready to do some particular salient activity, or is ready for something (in particular) to happen. Similarly, *it's raining* means that it's raining in some contextually specified place. Semantic minimalism denies these claims. Instead, the semantic minimalist claims that *Mary is ready* simply means that Mary is ready, and *It's raining* simply means that it's raining. There is no specification of a semantic complement of *ready* or of the location of the raining.

Unsurprisingly, these aspects of semantic minimalism have been met with numerous objections, which CL canvass and respond to (cf. Chapters. 11–12 of CL, 2005). According to one (cluster) of these objections, which I will call the Problem of Irrelevance, semantic minimalism results in a theory that is utterly irrelevant to psychology. Human linguistic practices are essentially driven by much more complex representations of meaning than what semantic minimalism offers; we *never* intend to communicate just the Semantic Minimalist's meaning of an expression. In short, there is no real cognitive role for a minimalist theory to play. Thus, a semantically minimalist theory is otiose, and so should be rejected as a style of semantic theorizing. In response to the Problem of Irrelevance, CL argue that there is a cognitive role for Semantic Minimalist meanings. They hold that we in fact do sometimes rely on such minimal meanings in our communicative practices.

In this paper, I wish to revisit the Problem of Irrelevance. I think that it relies on a common but false assumption about the nature of linguistic theorizing. The falsity of this assumption can be brought out with some considerations about the general nature of scientific theorizing. When we treat semantic as scientific theories about human language, we can see that there is no need to make the kind of head-on response to the Problem that CL attempt (which is just as well, since I'll argue that their attempt is unconvincing). Instead we can assert something much stronger: like any other scientific theory, a semantically minimalist theory may be on the right track even if speakers *never* intend to communicate the semantically minimal meaning of an expression. Indeed, we'll see that semantic minimalism may be on the right track even if speakers *cannot* (intend to) communicate such minimal meanings. If these claims are correct, then the Problem of Irrelevance has no force whatsoever.

This paper is structured as follows. In §2, I present the Problem of Irrelevance, and I critically examine CL's reply to it. In §3, I explore some properties of scientific models, and argue that semantic theories have these properties too. In particular, I argue that, in precise parallel to other scientific models, the particular outputs of a good semantic theory may be quite different from the communicative facts that form the raw data of a semantic theory. I then show how these facts invalidate the Problem of Irrelevance. I conclude in §4.

Before beginning, a couple caveats are in order. First, in this paper, I remain officially neutral about semantic minimalism. Although I will defend CL's view against

an objection, this defense should not be construed as an endorsement of it. Second, there are a number of issues concerning how to understand the position(s) of CL's critics, such as whether they are really offering different positions, and whether they aim to be working on (the foundations of) a scientific psychology (e.g., Stainton, 2006). The present paper has no stake in any of these issues. My interest concerns CL's methodology and a criticism of their theory. Thus, it doesn't really matter where the criticism comes from; e.g., even scientists with ultimately similar views to CL's can err in criticizing CL's view.

## 2 The Problem of Irrelevance

A number of authors have presented the Problem of Irrelevance for minimalist theories, and CL devote an entire chapter to discussing it (cf. CL, 2005, chap. 12 and citations therein). For instance, Recanati writes:

That minimal notion of what is said is an abstraction with no psychological reality, because of the holistic nature of speaker's meaning. From a psychological point of view, we cannot separate those aspects of speaker's meaning which fill gaps in the representation associated with the sentence as a result of purely semantic interpretation, and those aspects of speaker's meaning which are optional and enrich or otherwise modify the representation in question. They are indissociable, mutually dependent aspects of a single process of pragmatic interpretation. (Recanati, 2001, 88)

Recanati's pessimism about minimalist semantic theories is driven in large part by his view that a semantic theory must capture the entire "content of the statement as the participants in the conversation themselves would gloss it" (Recanati, 2001, 79–80). Recanati expresses this in his

"Availability Principle", according to which "what is said" must be analyzed in conformity to the intuitions shared by those who fully understand the utterance – typically the speaker and the hearer, in a normal conversational setting. This in turn supports the claim that the optional elements... (e.g., the reference to a particular time in "I've had breakfast") are indeed constitutive of what is said, despite their optional character. For if we subtract those elements, the resulting proposition no longer corresponds to the intuitive truth conditions of the utterance. (Recanati, 2001, 80)

The Availability Principle does not logically imply that minimalism is wrong, but it does tightly constrain minimalist theories. Minimalist theories, almost by definition, do not capture all of what Recanati calls the "intuitive truth conditions" of utterances of sentences. Thus, given the Availability Principle, a minimalist theory would need to defend itself by showing that it does in fact exactly capture the intuitive truth conditions of our utterances, at least sometimes. A bit more carefully, we can characterize the Availability Principle as follows:

(AP) A semantic theory is acceptable only if it correctly characterizes the intuitive truth conditions often enough within some psychologically interesting range of cases.

For present purposes, I will assume that (AP) is an appropriate formulation of Recanati's Availability Principle; any divergences between the two will not matter in this paper. One might strengthen (AP) further by specifying the particular range of cases in which a semantic theory must get things right, and by specifying how often the theory must get things right. I won't worry about such strengthenings, though; since what I have to say will apply equally to all such versions of (AP).

CL do not challenge the Availability Principle. Instead, they seem to accept it, and argue that there is a psychologically interesting range of cases where minimalist theories capture the intuitive truth conditions of utterances. This range of cases involve contexts where participants in the conversation (speakers, audience, others) must appeal to the "shared fallback content" of an utterance. In such cases, the minimalist's semantic content is what all participants can be assumed and expected to know even when some parties have mistaken or incomplete communication-relevant information. Moreover, this content is what can be conveyed by indirect quotation or reproduction to other parties in different contexts from the context of utterance (CL, 2005, 184–185). For example, suppose Jolene utters *Mary is tall*, and the audience disagrees about whether Jolene means that Mary is tall for a girl, or for a Samoan, or for a basketball player, or for a college student, or for something else. In such a state of ignorance, CL argue, the audience can still give some sort of characterization of the meaning of Jolene's utterance. For instance, they can report that Jolene said that Mary is tall, and that she didn't say that Mary is short, or that Peter will be tall. This agreement about the meaning of Jolene's utterance is the shared fallback content of Jolene's utterance, and, CL claim, it is precisely what minimalist theories characterize as the semantic content of the utterance.

Despite the cleverness of CL's argument, it's hard to see why an opponent of semantic minimalism would find it compelling. Opponents such as Recanati might concede that CL had identified a psychologically interesting range of communicative contexts, but deny that the minimalist's theory expresses the intuitive truth conditions in these cases. In particular, it would be natural for them to hold that in the cases CL imagine where the context is unknown or incomplete, our communicative expectations revert to default interpretations that nonetheless contain more content than the minimalist theory offers. For instance, in a context like the one just imagined, it would be reasonable to argue that we would expect and assume the various parties to interpret Jolene's utterance of *Mary is tall* as meaning that Mary is tall for an X, where the particular value of X is assumed and expected to normally be given by the context in some appropriate way. In the present case, the context failed to supply a value for X, so the shared fallback content includes something like a representation of a variable that has not been assigned a value. Thus, an opponent of minimalism could conclude, even in the cases CL imagine, there is still much more semantic structure present in the content of utterances than the minimalist predicts. Thus, the minimalist still needs to satisfy (AP) if her theory is to be acceptable.

Where does this leave us? At best, we seem to have reached an impasse of intuitions. CL claim to have found a way for minimalist theories to satisfy (AP), and their opponents have reasonable grounds for denying this. Fortunately, I think some

headway can be made on this issue. Unsurprisingly, I think that part of the problem lies with (AP) itself. In the next section, I will argue that we have good reasons for doubting both the truth of (AP) and the soundness of the Problem of Irrelevance.

### 3 Linguistic Theories as Scientific Theories

Before getting started, let's look at the general structure of the Problem of Irrelevance. The argument appears to take the following form:

- (3i) In all relevant ranges of cases, the intuitive truth conditions of our utterances contain much more content than what is characterized by minimalist theories.
- (3ii) If (3i) is right, then from a psychological point of view, we cannot separate the minimalist aspects of meaning from those aspects supplied by a more enriched view of meaning (often enough, in any relevant range of cases).
- (3iii) Hence, minimalist aspects of meaning cannot be separated from those aspects supplied by a more enriched view of meaning (often enough, in any relevant range of cases).
- (3iv) But if we can't separate minimalist from non-minimalist elements of meaning (often enough, in any relevant range of cases), then minimalist theories are unacceptable.
- (3v) Hence, minimalist views are unacceptable.

Premise (3i) is an empirical claim; CL attempt to refute it by appealing to cases where we employ shared fallback content. However, we've seen that their attempt at best leads to an unresolved clash of intuitions between them and their opponents. That leaves only premises (3ii) and (3iv). Premise (3ii) comes from the quote of Recanati above (2001, 88), and premise (3iv) comes from (AP). (To see this, notice that if we can never separate out the minimalist aspects of meaning, then there must always be some non-minimalist aspects present, so the minimalist aspects of meaning never characterize the intuitive truth conditions in the utterance. Hence, by (AP), minimalist theories are unacceptable.)

In this section, I argue that neither (AP), (3ii), nor (3iv) is credible, largely for the same reasons. In order to do this, I will assume that linguistic theories, and semantic theories in particular, are scientific theories (or "models", as I occasionally call them). I consider this assumption, and my use of it in this paper, to be utterly trivial. Nonetheless, I want to register this assumption up front, since one still occasionally encounters philosophers who distinguish "doing philosophy" from "doing science". I'll say a bit more about this assumption later on, but for now let's turn to the central discussion.

In order to see what is wrong with (AP), (3ii), and (3iv), it will be useful to step back from linguistic theorizing and examine some aspects of the methodology of the (other) sciences. Ultimately, I maintain that when we look at the details of how a typical quantitative empirical theory is related to the raw data, we can precisely identify counterparts in the linguistic theory. The upshot of all this will be that just as other scientific theories of particular phenomena need not – and very often do

not – directly account for the raw data, so too linguistic theories will often fail to fully explain the raw linguistic data that drives our theories. And this divergence between the linguistic theory and the data is a perfectly normal aspect of perfectly normal science. In short, no non-linguistic theories would ever be constrained to observe appropriate counterparts of (AP), (3ii), or (3iv). These features of scientific theories that free them from such constraints are found in linguistic theories too. Hence, (AP), (3ii) and (3iv) are unacceptable for linguistics for just the same reasons that their counterparts are unacceptable in the sciences.

To set this discussion in context, it's worth observing that Chomsky has long defended parallels between linguistics and other sciences (e.g., Chomsky, 2000). My discussion is certainly of a piece with Chomsky's, but it differs in one crucial respect. Rather than attending to purely verbal characterizations of historical examples, I focus on some quantitative aspects of current scientific methodology. Thus, we will be able to address certain issues with considerably more precision and detail than Chomsky offers. It is perhaps also worth mentioning that the present approach also makes possible a critical examination and principled rejection of some claims of Chomsky's that in terms of its methodology, linguistics is just like the other sciences. I discuss this issue in detail elsewhere (Johnson, 2007b).

To get things started, let's take a simple example. Suppose we are studying the relationship between different quantities of a given additive  $X$  used in some manufacturing process and the amount of some type of atmospheric pollution  $Y$  generated by the process. The industry standard is to use  $n$  units of  $X$  per ton of product, but for a period of time, certain companies used more or less than  $n$  units. The relation between the varying amounts of  $X$  used and  $Y$  emitted are given as black diamonds in the plot below. (Zero on the  $x$ -axis represents the use of  $n$  units of  $X$ ; other values represent the respective deviations from this standard; ignore the two curves and white diamonds for the moment.)

Given this data, there are many – infinitely many, in principle – possible relations that could hold between  $X$  and  $Y$ . One extreme option would be to insist that every aspect of the data is crucial to understanding how  $X$  and  $Y$  are related. In such a case, a researcher might look for a function that captured the data precisely, as in the very complex one depicted with a solid line. The resulting theory will then perfectly predict the behavior of  $Y$  on the basis of the behavior of  $X$ . The raw data, in the form of a set of pairs of measurements  $\{ \langle X_i, Y_i \rangle : i \in I \}$ , is fully accounted for.

Despite its success at capturing the data, the first approach is almost never adopted. A vastly more common strategy is to assume that the real relation between  $X$  and  $Y$  is much simpler, and that  $Y$  is influenced by other factors that are unrelated to  $X$ . One might, e.g., hypothesize that relationship is given by the simple function:

$$(4) \text{ Predicted value of } Y_i = f_1(X_i) = \beta_0 + \beta_1 X_i + \beta_2 X_i^2$$

for some fixed numbers  $\beta_0, \beta_1, \beta_2$ . Once these numbers are determined from the data, we get the much simpler curve given by the dashed line. In the present example, the values of  $\beta_0, \beta_1$ , and  $\beta_2$  were determined by seeking those values for which  $\sum_{i \in I} [(Y_i - f_1(X_i))^2]$  is as small as possible. If we simply wanted a polynomial

function that captured the data perfectly, with probability 1,  $f_2$ , a polynomial of degree 29, will do so, for the given raw data set of size 30:

$$(5) \text{ Predicted value of } Y_i = Y_i = \beta_0 + \beta_1 X_i + \beta_2 X_i^2 + \dots + \beta_{29} X_i^{29}$$

Although  $f_1$  doesn't predict the behavior of the original data as well as its rival  $f_2$ , many other theoretical considerations speak in its favor. For example, suppose we got hold of another sample of data, given by the white diamonds above. Then we might ask how well the two functions captured this new data. One way to do this would be to compare the sizes of the discrepancies between what  $f_1$  and  $f_2$  predict about the value of  $Y$  for given values of  $X$  in the new data set. E.g., we might examine the ratio:

$$(6) \frac{\sum_{i \in I'} [(Y_i - f_1(X_i))^2]}{\sum_{i \in I'} [(Y_i - f_2(X_i))^2]}$$

Here  $I'$  indexes the second set of measurements, and  $f_1$  and  $f_2$  are assumed to have had the particular numerical values of their parameters  $\{-\beta_0, \beta_1, \beta_2\}$  in the case of  $f_1$ , and  $\{\beta_k : 0 \leq k \leq 29\}$  in the case of  $f_2$  – fixed by the first data set. In the particular case presented in Fig. 1, (6) yields a value around  $6 \times 10^{-31}$ , indicating that there is vastly less discrepancy between the new data and what  $f_1$  predicts than there is between this data and what  $f_2$  predicts. (In a typical case, this number would be even smaller, since there are  $n - 3$  and  $n - 30$  degrees of freedom associated with  $f_1$  and  $f_2$  respectively; (6) does not factor in this discrepancy.) From a God's-eye view, this is unsurprising, because  $f_1$  is the form that actually generated the data (although the data themselves determined the values of the  $\beta_i$ s).<sup>1</sup> The extra structure in the curve given by  $f_2$  errs by capturing much of the variance in the data that is unrelated to the true relation between  $X$  and  $Y$ . (There's much more to be said about the general issues of model construction and model evaluation; cf. e.g. Forster and Sober, 1994; Burnham and Anderson, 2002; Bamber and van Santen, 1985 for further relevant discussion.) In short, a bizarre model like  $f_2$  that captures all the

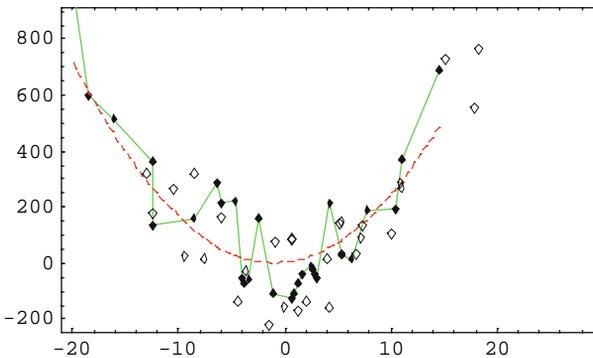


Fig. 1 Deviations from the normal amount of additive X versus amount of pollutant Y

(original) data is vastly inferior to the far more standard model like  $f_1$  that doesn't. In particular, the simpler model does a massively better job at predicting the general trends of new data as it arrives.

How then is  $f_1$  related to the actual raw data? Typically, this connection is formed by adding a "residual" or "error" term to our equation:

$$(7) Y_i = f_1(X_i) + \varepsilon_i = \beta_0 + \beta_1 X_i + \beta_2 X_i^2 + \varepsilon_i,$$

The term  $\varepsilon_i$ , whose value varies as  $i$  varies, expresses whatever deviation is present between the model and the raw data (i.e.,  $\varepsilon_i = Y_i - f_1(X_i)$ ). In practice, scientific models of complex phenomena *never* perfectly fit the data, and there is *always* a residual element ( $\varepsilon_i$ ) present. This is so even when the system under study is completely deterministic, with no random influences present in the data whatsoever. E.g., the true model might be something like

$$(8) Y_i = f_1(X_i) + f_3(Z_{li}, \dots, Z_{ki})$$

In such a case,  $Y$  is always an exact function of  $X$  and  $Z_1, \dots, Z_k$ . However, the influence of the  $Z_j$ s may be very small, very complicated, unknown, poorly understood, etc. Thus, for any number of reasons, it may be natural to model the phenomena with  $f_1$ , all the while realizing that the existence of residuals in the raw data show that there is more to the full story than  $f_1$  presents.

With this brief overview of scientific theorizing in hand, let's return to linguistics. In the relevant respects, linguistic theorizing is similar to the statistical case. Like  $f_1$ , semantic theories are models of a complex phenomenon (the interpretation of language). The raw data of a sample of this phenomenon aren't represented as numerical measurements, as in the statistical example. Instead, they are given as assessments about certain types of idealized<sup>2</sup> linguistic behavior: what sorts of things would typical speakers communicate by uttering a given sentence, and under what conditions? In other words, the raw data of semantic theorizing are the intuitive truth conditions of our utterances, as we do or would make them in various contexts. Proceeding in a manner similar to the statistical researcher, the minimalist theorist begins by hypothesizing that there is some relatively simple structure – i.e. simple in comparison to the complexity of the raw data – that accounts for much of the collective behavior of the raw data. In order to obtain this relatively simple underlying general structure, some aspects of the raw data (i.e., the intuitive truth conditions) will have to be ignored, just as we ignore some aspects of variance in the statistical case. Here the analogy between the minimalist and statistical strategies is quite tight. In the statistical case, the raw data was decomposed into two parts: the pollution level  $Y$  is the additive combination of the effects of  $X$  (given by the population model  $\beta_0 + \beta_1 X + \beta_2 X^2$ ) and some other effects (given by  $\varepsilon$ ) that are not part of the present study. The general structure of this statistical model is thus:

$$(9) \text{Raw Data} = \text{(i) Effects of processes under study} \text{ (ii) Interacting in some way with (iii) Residual Effects}$$

But, and this is a crucial point, (9) is also the general structure of the minimalist's semantic theory. The minimalist theory supplies some aspects of meaning that are

hypothesized to capture much of the general behavior of the *totality* of the data set. By assumption, the outputs of this theory needn't capture all of the raw data (i.e., the intuitive truth conditions of the utterances in the data set). In fact, the semantic theory needn't *ever* capture all of the intuitive truth conditions. Such an outcome is absolutely standard science. For instance, our pollution researcher would not assume that there will be some raw datum  $Y_i$  such that  $Y_i = f_1(X_i)$ , with no contribution from the residual effects. Indeed, it is quite typical to expect that  $\varepsilon_i$  will never equal 0, particularly when the phenomenon under study is extremely complex. (Indeed, when the phenomena are quite complex, a model may be considered significant even if it captures as little as 16% of the variation in the raw data (e.g. Putnam, 2000, 487).) In the case of semantics, the intuitive truth conditions (on analogy with the  $Y_i$ ) are perhaps always partly determined by the minimalist theory of meaning (on analogy with the  $f_1(X_i)$ ), and partly determined by other aspects of communication that interact with the minimalist theory in some way (on analogy with the  $\varepsilon_i$ ). These other aspects of communication are just the familiar processes that are hypothesized to interact with the hypothesized minimalist linguistic capacities: background beliefs, indexical-fixing elements, demonstrations, "performance" capacities of speaker/hearers, etc.

As an aside, it's worth noting (again) just how "Chomskian" my interpretation of semantic minimalism is. In many places, Chomsky cautions that linguistic theories are not obliged to capture all the facts about various grammaticality judgments, or all of various details present in collections of data. By seeking out more general patterns, we may be able to learn about a speakers' linguistic "competence", which can be masked by additional "performance" factors that are also realized in the empirical data. For instance, Chomsky writes:

Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of the language in actual performance. (Chomsky, 1965, 3)

By idealizing away from various extraneous factors, we can "smooth" out the raw data and thereby focus on some of the more significant forces and structures underlying human linguistic abilities.

To return to our main theme, a minimalist semantic theory is a *theory* about the nature of the raw data. Like any other scientific theory, one of its essential rights and obligations is to characterize those parts of the raw data it considers to be truly part of the phenomenon under study, and what other parts are due to extraneous processes; cf. (i) and (iii) in (9). From the standpoint of ordinary scientific theorizing, this point is too obvious to merit mention. But in the present case, for some reason, the point appears to some to be completely optional. Others have noticed this odd practice, too. For instance, Paul Pietroski writes that

one shouldn't try to define the scope of semantic theorizing in advance of investigation. . . . One cannot determine a priori which facts a semantic theory should explain. In any domain, what a theory *should* explain depends on what gets discovered. And we can discover

semantic facts, just as we can discover biological facts. In my view, insensitivity to this point still pervades philosophy of language, despite Chomsky . . . and the remarkably successful research program he initiated. (Pietroski, 2005, 26–27)

(Pietroski provides ten citations spanning six decades where Chomsky has urged this point; many more citations could be provided.) The fact that a semantic theory gets to characterize its own scope also means that, like any other scientific theory, a semantic theory should be judged by the ordinary, complicated but familiar, criteria of successful theories: simplicity, elegance, predictive fecundity, integration with other successful theories which collectively account for the raw data (or, more typically, hopefully someday will account for the raw data), etc. Methodologically speaking, demanding that a semantic theory sometimes exactly characterize the intuitive truth conditions of utterances appears to be just like demanding that statistical models should (at least for some interesting range of values) be like the complex  $f_2$ , instead of the like the more common  $f_1$ . Such a demand would be bizarre and deeply incorrect in the statistical case; I submit it is no better motivated in the case of semantic theorizing.

The points just made are absolutely crucial to the understanding and evaluation of a semantic theory. Moreover, they form the core of why (AP), (3ii), and (3iv) should be rejected. The problem with (AP) and (3iv) is that they place an unwarranted constraint on theory construction. In any other study of complex phenomena, the demand that theories perfectly capture the raw data across some interesting range of cases would grind theorizing to a halt. The only work that could be done would be purely descriptive and utterly unpredictable and unexplanatory characterizations of the available data, along the lines of  $f_2$ . (3ii) should be rejected because it is one of the rights (and obligations) of a theory to provide a theoretically useful characterization of the phenomena it addresses. Typically this is done by constructing a good theory, and then declaring that the relevant phenomena are whatever the theory addresses. In the case of scientific models, this is the essence of why they are constructed and confirmed using statistical methods. There is no reason to deny such rationale to semantic theories, contrary to the announcement in (3ii). Thus, the Problem of Irrelevance rests on two assumptions which we have little reason to believe. So we have little reason to think that the Problem of Irrelevance poses a real problem at all.

There is another way to view the flaws with (AP), (3ii) and (3iv) that helps to bring out why the Problem of Irrelevance might appear sound. The Problem of Irrelevance fundamentally turns on an equivocal interpretation of “separability”. As Recanati and many others have made clear, the raw data of linguistic theory – the intuitive truth conditions of our utterances – are probably nearly always substantially influenced by pragmatic factors. In this sense, Recanati is probably right that pure semantic content is “inseparable” from pragmatic factors: in actual language use, you rarely if ever find the former alone, without the latter. Indeed, the contribution of pragmatics is probably thorough enough to make this last claim hold very generally. This notion of inseparability makes (3ii) plausible (recall that (3ii) says that if intuitive truth conditions typically contain more content than what the minimalist

theory supplies, then we can't separate minimalist from non-minimalist aspects of meaning). However, such an interpretation of inseparability makes it unclear why (AP) and (3iv) should hold (recall that (3iv) says that minimalist theories are acceptable only if we can (sometimes) separate minimalist and non-minimalist aspects of meaning; (AP) says that a semantic theory is acceptable only if it (sometimes) correctly characterizes the intuitive truth conditions). After all, we've seen that it's no criticism of a theory that it treats the raw data as being a product of multiple sources. If this is what separability amounts to, then the claim that minimalist and non-minimalist aspects of meaning are inseparable simply begs the question against minimalist theories. Minimalists will certainly hold that their theories are reasonable, even though these theories purport to explain the two types of aspects of meaning using different mechanisms.

How then might we make (AP) and (3iv) plausible? To do this, we need a notion of inseparability that applies to linguistic *theories*. That is, (AP) and (3iv) are plausible if inseparability here means that no reasonable total theory of language will treat minimalist and non-minimalist aspects of meaning as effects of (relevantly) distinct processes. That is, in order for (AP) and (3iv) to be plausible, the relevant notion of inseparability must require that all aspects of the intuitive truth conditions be explained by the same mechanisms. Then (AP) and (3iv) are virtually tautologies, hence trivially true. But now there is no support for (3ii). Why should the fact that the intuitive truth conditions of our utterances *do* contain both minimalist and non-minimalist aspects of content be sufficient to license the restriction that any theory of semantic content *must* capture all of these aspects? Such a view clearly begs the question against minimalist semantic theories, which propose to capture only part of the intuitive truth conditions of our utterances. Moreover, when (3ii) is interpreted this way, it is independently implausible for reasons we have already discussed; (3ii) would require semantic theories to account for all of the raw data, in dramatic contrast to the way all other theorizing proceeds.

In short, (3ii) is a kind of converse of (AP) and (3iv). (3ii) becomes plausible when we interpret inseparability as implying that our intuitive truth conditions are almost always the product of minimalist and non-minimalist aspects of meaning; however, this makes (AP) and (3iv) question-begging and implausible. On the other hand, (AP) and (3iv) are plausible when we interpret inseparability as implying that no reasonable theory will separate minimalist and non-minimalist aspects of meaning; however, this makes (3ii) question-begging and implausible. The problem with either interpretation of inseparability is the same: there's no apparent way to motivate the inference from the empirical fact that our utterances typically contain minimalist and non-minimalist aspects of meaning to the normative claim that all these aspects must be captured by a semantic theory. Thus, we can locate the structural flaw in the Problem of Irrelevance in an equivocation regarding inseparability.

In the final analysis, a lot of the present discussion hangs on the current epistemological situation with respect to linguistics. I've stressed that part of a scientific theory involves carving up the data in one way or another. This aspect of theorizing is especially important in linguistics because of how little we know about how language works. If we had a better idea of how it works, this additional knowledge

would likely constrain what structures, processes, etc. could be plausibly employed by a theory. As it is, though, our lack of knowledge about language leaves us with very few such constraints on theories (at least at the level of detail that is currently relevant here – i.e., a level where we may ask whether minimal propositions play some psychological role concerning language).

The difficulty with constraining which entities, etc. a theory may reasonably posit is further compounded by the fact that there is tremendously strong evidence that many entities may not be introspectively accessible by anyone. To see this, notice that many aspects of cognitive processes are not consciously perceivable by persons undergoing them. The many decades of literature on both human vision and judgment and decision-making have presented countless studies showing that many aspects of these types of cognitive processes are not consciously accessible to us. Similarly, people lack conscious awareness of many hypothesized aspects of language and its processing (Johnson, 2007a; cf. Townsend and Bever, 2001 for an overview). E.g., people aren't aware of such items as traces, PRO, movement, Merge, etc. that are posited by *some* linguistic theories simply because of the theoretical work they do. But opposing theorists cannot simply announce that the former theories are false, because they employ some entity which they *feel* is not truly part of language or its processing. The entity in question may well be part of the end product or part of the processing, or both, even if people are not generally consciously aware of it. (Indeed, even in the case of semantics, there is evidence that people have substantial difficulties becoming aware of certain aspects of the meanings of relatively ordinary sentences (Johnson, 2007a).)

In the present case, then, it's not enough for Recanati to simply declare that CL's minimal propositions won't find any place in a scientific theory of language. He also needs to supply enough *uncontroversial* details of a theory of language to show that no such minimal propositions will play any role in a completed correct theory. Lacking such details, CL's theory is not objectionable on these grounds. (I take it that this is just a general point about the nature of scientific theories in general. Much of the philosophy of science concerns how scientific theories will often posit unobserved entities, where the justification for treating these entities as real comes only later, as the theory is confirmed by the usual holistic criteria – simplicity, accuracy, integration with larger theories, predictive fruitfulness, etc. These are the criteria by which a minimalist theory should be judged, not by apriori speculation about what entities will appear in a completed theory of language.)

#### 4 Rescuing the Attack on Minimalism?

In the previous section, I levied a charge against the Problem of Irrelevance. To those familiar with Chomsky's work, this charge may seem familiar. To use Chomsky's words, Recanati and the other proponents of the Problem of Irrelevance are guilty of "methodological dualism" (e.g., Chomsky, 2000, chap. 4). Methodological dualism is characterized by the insistence on a set of standards for some area of research (semantic theories in this case) that would be utterly inappropriate and trivially

unacceptable in any other area of inquiry. Chomsky has long argued that the study of language and mind is held to just such dualistic standards. Furthermore, he has argued, such standards would cripple any other developed discipline such as chemistry or physics, were they applied there.

Now of course, there is nothing *a priori* wrong with this sort of methodological dualism. Indeed, maybe language and mind are “special” in some important sense. Maybe language and mind are so radically unlike every other complex phenomena that humans have theorized about that such dualistic standards are appropriate. Perhaps assumptions like (AP) are appropriate constraints on semantic theories, even though their corresponding forms would be a travesty anywhere else. Recanati appears to believe something like this, because he provides an argument in favor of (AP). If this argument is compelling, perhaps it could justify the unique standards he imposes on semantic theories. Thus, it will behoove us to briefly examine this argument.

The essence of Recanati’s argument is given in the passage below:

Suppose I am right and most sentences, perhaps all, are semantically indeterminate. What follows? That there is no such thing as ‘what the sentence says’ (in the standard sense in which that phrase is generally used). . . . If that is right, then we cannot sever the link between what is said and the speaker’s publicly recognizable intentions. We cannot consider that something has been said, if the speech participants themselves, though they understand the utterance, are not aware that that has been said. This means that we must accept the Availability Principle and its maximalist consequences. (Recanati, 2001, 87–88)

This argument is puzzling. Given the context of this argument, Recanati’s claim that most or all sentences are “semantically indeterminate” amounts to the claim that minimalist semantic theories don’t capture the intuitive truth conditions of most or all sentences. It is unclear, though, why such a claim should be taken to imply that “there is no such thing as ‘what the sentence says’ (in the standard sense in which that phrase is generally used)”. I take it that “what the sentence says” here refers to the content that a (minimalist or other standard) semantic theory ascribes to a sentence. If that is correct, then claiming that there “is no such thing” is question-begging. We’ve seen that it’s part of the job of a theory of some complex phenomenon (semantic or otherwise) to carve out the exact sub-portion of the phenomenon that it directly deals with, leaving the remaining parts for further theorizing. Recanati’s claim that there is no such thing as “what is said” in this context is like saying there is no such thing as the true population model  $f_1$  in the statistical case. I conclude that this inference is invalid, for largely the same reasons that (AP), (3ii), and (3iv) are unjustified. Hence, I do not believe that Recanati has successfully resurrected his case against minimalism.

How might one defend a claim like (AP)? (I focus on (AP), although similar considerations also hold for similar claims like (3ii) and (3iv).) The most obvious strategy, of course, would be to come up with a rival semantic theory that has this form, and show that it is to be preferred overall, using the standard criteria for the selection and confirmation of scientific theories. This strategy reverses Recanati’s: instead of defending a type of theory as preferable *because* it respects (AP), one

would need to defend (AP) because a type of theory that is independently determined to be preferable respects it.

Barring such a straightforward approach, there appear to be two general types of defenses of (AP). Both types attempt to show that semantics is unlike all other empirical disciplines insofar as residual effects contribute nothing to the raw data in some (or enough) instances. On the one hand, one might try to provide a *metaphysical* defense of such a principle. For instance, one might argue that no matter how integrated the phenomenon of language is within a larger interacting system (e.g., human cognition, social interactions, etc.), the phenomenon must occasionally deliver “pure” manifestations of itself. Unfortunately, this strategy looks hopeless, at least insofar as it is an empirical fact about language that the raw data we use to study it (i.e., the intuitive truth conditions of utterances) are the expression of multiple different systems. (Moreover, it also appears that we have strong independent reasons for individuating and distinguishing these systems, so that the previous claim is not arbitrary.) If, as the empirical research very strongly suggests, this is the case, then it’s an open empirical question whether the raw data is ever a pure expression of language only, or whether it always also contains other, extra-semantic, aspects (e.g. Gleitman and Liberman, 1995) Thus, it looks like any such metaphysical argument may run afoul of the empirical facts.

On this note, it is also an open, and quite interesting, question whether one could construct a context, perhaps in the form of a carefully designed experimental setting, in which subjects would reliably offer all and only what the minimalist characterizes as the meaning of our sentences. Indeed, one way to study a minimalist theory would be to look for a single context in which subjects reliably offered the minimalist’s truth conditions as the intuitive truth conditions for a broad range of expressions. Finding “interesting” contexts of this sort (e.g., *not* contexts like “Give me the minimalist truth conditions of the following utterance. . .”) could prove quite interesting. Finding such contexts could also falsify (3i) in the Problem of Irrelevance.

The second type of strategy for defending (AP) is *epistemological*. In this vein, one might argue that any semantic theory must be sophisticated enough to occasionally completely explain exactly how all the component effects work to produce the raw data. As I noted in the previous paragraph, it’s an interesting question whether there are relevant contexts in which the minimalist’s truth conditions are identical to the intuitive truth conditions. But let’s leave that issue aside. In general, we cannot demand that a theory of some complex phenomenon have this level of sophistication, for the simple reason that the theory would have to be so complex that we wouldn’t know how to go about constructing it. With a highly complex phenomenon like language, it’s vastly more reasonable to identify some (hopefully) “natural part” of the phenomenon, and work with that, letting the resulting theory determine the nature of the natural part as inquiry progresses. To demand that our theories account for all of the raw data would be to ask for a theory that we don’t have any idea how to construct (and couldn’t use if we ever did build it). In the present case, it would have to be at least a near-total theory of human psychology. And to place this kind of requirement on a theory would be to grind research to a halt. Hence such a requirement should be rejected.

Incidentally, it's worth observing that Recanati himself appears to reject any such requirement on semantic theories. After all, the raw data of linguistics are *individual persons'* assessments of the intuitive truth conditions of utterances in contexts. Recanati's discussion suggests that he is following standard practice by abstracting away from the "residual" disagreements between persons (or between one person at different times). One wonders why such a maneuver would be legitimate for Recanati here, if it is not also legitimate in the case at hand. Of course, abstracting away from differences in judgments is necessary if one hopes to carve out a project that can be addressed and worked on. Precisely similar considerations are what motivate many semanticists, minimalists included, to further abstract away from the (idealized) intuitive truth conditions of utterances in the hopes of uncovering some general structure.

## 5 Conclusion

We've seen that the Problem of Irrelevance fails primarily because it ignores a basic fact about research into complex phenomena. This fact is that each theory gets to determine what part of the phenomena it addresses, and typically this is only a very proper subpart of the total phenomena. The requirement that a theory accommodate all of the intuitive truth conditions often enough in some relevant range of cases is a restriction on semantic theories that has no precedent in any of the developed sciences. Indeed, it is far more typical to assume that a given theory will not account for the data. Theories typically don't account for all the raw data for both the metaphysical reason that the data are typically the expression of multiple interacting processes, and for the epistemological reason that any such theory would be too hard for us to even begin constructing, and probably too hard for us to confirm or even to understand, if we were able to construct it.

On a final note, I have occasionally encountered the view that my assumption that semantic theories are scientific theories is somehow optional, and that one could avoid all these problems by stipulating that Recanati's position is a "philosophical" view. I confess I simply don't understand this position (not least because I do not understand what a distinctively "scientific" or "philosophical" theory would be). Nonetheless, there are a few things that can be said in response. First, I mean very little by calling a semantic theory "scientific". Semantic theories deserve this appellation, I suggest, primarily because their construction and confirmation centrally involve employing some of our best known methods for obtaining knowledge about a particular empirical phenomenon (the interpretation of language). From this perspective, it is unclear how one could reasonably defend the importance of a "non-scientific" theory of language. Moreover, my present use of the idea that semantic theories are a type of scientific theory is, I believe, especially uncontentious. So if one wants to resist my conclusions by taking this route, one needs to show why the particular features of semantic theorizing that I've appealed to are not part of some other form of theorizing *that can be independently motivated as an interesting and worthwhile project*. I see no way of doing this.

**Acknowledgment** I am grateful to Rob Stainton for much useful feedback.

## Notes

1. The data were generated using the formula  $Y_i = 3 + 4X_i + 2X_i^2 + \varepsilon_i$ , where  $\varepsilon$  was normally distributed with a mean of zero and a standard deviation of 100. The values for  $X$  were randomly drawn from a normal distribution with a mean of zero and a standard deviation of 10.
2. The notion of idealization in linguistics and the other sciences has been discussed at great length in many places (e.g., Liu, 2004; Chomsky, 1986, and citations therein). Since the primary data of interest in the present paper concerns “intuitive truth conditions”, the idealizations at play here are substantially less (although by no means absent!) than in other areas of linguistics.

## References

- Bamber, D., & van Santen, J. P. H. (1985). How many parameters can a model have and still be testable? *Journal of Mathematical Psychology*, 29, 443–473.
- Burnham, K., & Anderson, D. R. (2002). *Model selection and multimodel inference* (2nd ed.). New York: Springer.
- Cappelen, H., & Lepore, E. (2005). *Insensitive semantics*. Oxford: Blackwell's.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (1986). *Knowledge of language*, Westport, CT: Praeger.
- Chomsky, N. (2000). *New directions in the study of language and mind*. Cambridge: CUP.
- Forster, M., & Sober, E. (1994). How to tell when simpler, more unified, or less ad hoc theories will provide more accurate predictions. *British Journal for the Philosophy of Science*, 45, 1–35.
- Gleitman, L., & Liberman, M. (Eds.). (1995). *An introduction to cognitive science, volume I: Language*. Cambridge, MA: MIT Press.
- Johnson, K. (2007a). Tacit and accessible understanding of language. *Synthese*, 156, 253–279.
- Johnson, K. (2007b). The legacy of methodological dualism. *Mind and Language*, 22, 366–401.
- Liu, C. (2004). Laws and models in a theory of idealization. *Synthese*, 138, 363–385.
- Pietroski, P. (2005). *Events and semantic architecture*. Oxford: OUP.
- Putnam, R. (2000). *Bowling alone*. New York: Touchstone.
- Recanati, F. (2001). What is said. *Synthese*, 128, 75–91.
- Stainton, R. J. (2006). Terminological reflections of an enlightened contextualist. *Philosophy and Phenomenological Research*, 73, 460–468.
- Townsend, D., & Bever, T. (2001). *Sentence comprehension*. Cambridge, MA: MIT Press.