

Context-Sensitive Truth-Theoretic Accounts of Semantic Competence

STEVEN GROSS

Abstract: According to cognitivist truth-theoretic accounts of semantic competence, aspects of our linguistic behavior can be explained by ascribing to speakers cognition of truth-theories. It's generally assumed on this approach that, however much context-sensitivity speakers' languages contain, the cognized truth-theories themselves can be adequately characterized context-*insensitively*—that is, without using in the meta-language expressions whose semantic value can vary across occasions of utterance. In this paper, I explore some of the motivations for and problems and consequences of dropping this assumption.

1. Introduction

According to cognitive truth-theoretic accounts of human semantic competence, aspects of our linguistic behavior can be explained by ascribing to speakers cognition of truth-theories—that is, cognition of a finite number of axioms assigning semantic values to syntactically simple lexical items, a finite number of combinatorial axioms assigning semantic values to syntactically complex expressions on the basis of their syntax and their constituents' semantic values, and a finite number of production schemata the application of which enables the derivation of 't-sentences' stating truth-conditions for sentences of the language with which the speaker is competent.¹ Such theories are compositional: they assign complex expressions semantic value as a function of their structure and the semantic values of their constituents. Thus their attraction in part is that ascribing them can explain semantic creativity (our ability to grasp readily the meanings of novel sentences composed of familiar constituents) and semantic systematicity (the various patterns our competence exhibits, such as the fact that speakers capable of grasping 'Fred runs' and 'Mary hops' are also capable of grasping 'Mary runs' and 'Fred hops').

One major task facing this approach is to accommodate the *prima facie* phenomenon of context-sensitivity in all its variety—the various ways the truth-conditions associated with a sentence can seem to vary across occasions of use, as

Thanks for helpful feedback to audience members at ESPP Congress 2003 where an excerpt from this paper was presented. Thanks also to an anonymous referee for extremely useful comments.

Address for correspondence: 433 Logan Hall, Department on Philosophy, University of Pennsylvania, Philadelphia, PA 19104, USA.

E-mail: gross2@phil.upenn.edu

¹ This approach to semantic competence stems from Davidson, 1967. Larson and Segal, 1995, is a textbook presentation of cognitivist neo-Davidsonianism.

the truth-conditions associated with 'I am tall', for example, can seem to vary with the speaker, the time of utterance, and the contextually relevant standard or comparison class for tallness.²

Truth-theorists tend to accept the thesis of *Meta-Insensitivity*, according to which one can adequately characterize the cognized truth-theories context-insensitively—that is, without using expressions in the meta-language whose semantic value can vary across occasions of utterance.³ In particular, on this view, one can accommodate *object-level* context-sensitivity without adverting to context-sensitivity in the meta-language. A standard hypothesis, for instance, is that speakers cognize that certain features of context help determine an expression's semantic value, so one can accommodate context-sensitivity by positing appropriate contextual parameters in the truth-theory. The normal form for a statement of truth-conditions, for example, might then be, not that of Tarskian t-sentences ('[S] is true iff P'), but rather that of *conditionalized* t-sentences:

For any u, x_1, \dots, x_n , if u is an utterance of [S] such that $A(x_1, \dots, x_n)$, then u is true iff $B(x_1, \dots, x_n)$.⁴

For instance:

For any u, s, t , and k , if u is an utterance of [I am tall] by s at t such that the contextually relevant comparison for tallness is k , then u is true iff s is taller than k at t .⁵

If one can thus represent *all* sources of context-sensitivity by variables, then the cognized constraints on semantic value that remain can be characterized context-insensitively.

It's been questioned, however, whether such parameterization strategies can succeed in accommodating all context-sensitivity. In the next section, I outline some of

² 'Prima facie' because, as we'll note, one might try accommodating some cases by denying that they exhibit context-sensitivity.

From the first, Davidson, 1967, p. 33, recognized context-sensitivity as a 'very large . . . fly in the ointment'. The literature on the topic is now enormous. I survey some of the variety of prima facie context-sensitive constructions in Gross, 1998/2001, Chapter 1.

³ See, e.g., Davidson, 1976, p. 175, and Higginbotham, 1988, pp. 29 and 31.

⁴ '[S]' stands for the semantically relevant syntactic representation of 'S'—that is, its LF representation.

An alternative to conditionalizing t-sentences is to relativize the ascription of truth to a sequence. For example, while Higginbotham, 1988, from whom I borrow talk of the 'normal form,' advocates conditionalized t-sentences, Larson and Segal, 1995, advocate relativized t-sentences, from which conditionalized t-sentences can be derived with the help of extra-semantic axioms and production schemata. This difference is not relevant to our concerns.

⁵ The sample t-sentence is meant for illustrative purposes only—there are other treatments of 'I am tall's context-sensitivity. More generally, in presenting sample t-sentences, I abstract away from details irrelevant to our topic.

the reasons one might offer. But my main topic is not these criticisms themselves, but rather a strategy truth-theorists might pursue if such criticisms should prove correct. What results if one drops *Meta-Insensitivity*? In section 3, I sketch how this would allow us to sidestep the outlined criticisms while preserving compositionality. One obvious problem, however, if one allows context-sensitivity in the characterization of cognized truth-theories, is accounting for our competence in comprehending utterances of certain context-sensitive sentences entered in contexts other than ours, as when I understand what the words 'I am tall' express when *you* utter them. After raising this problem in section 4, I discuss two ways one might address it in sections 5 and 6. Section 7 then examines what some of the consequences might be if one abandons *Meta-Insensitivity*. Throughout, the emphasis is on, not *advocating* truth-theoretic accounts of semantic competence that utilize context-sensitivity, but rather making them visible enough to become objects of assessment.

2. Three Criticisms

Accommodating context-sensitivity in a truth-theoretic approach to semantic competence is no small task. Arguably, its *prima facie* complexity provides in itself sufficient motivation at least to consider dropping *Meta-Insensitivity*, insofar as doing so makes available further resources for handling the phenomena. But not all will find the consequences attractive. So, I motivate my discussion by sketching three kinds of reason one might advance against the possibility of accurately accommodating context-sensitivity in a context-insensitively characterized meta-language: first, *complexity* criticisms, according to which the dependence of at least some sentences' truth-conditions on context is just too complex and open-ended to be captured without residue in parameterized t-sentences; second, *indexicality* criticisms, according to which context-insensitively characterized truth-theories fail to capture how certain sentences represent how things are in an essentially indexical manner; and, third, *empirical* criticisms, according to which the empirical evidence simply doesn't bear out the claim that speakers cognize such parameterized truth-theories. These brief remarks won't establish the success of any of these criticisms, but only indicate the issues to a minimally sufficient extent that an examination of context-sensitively characterized truth-theories might seem worthwhile.

2.1 Complexity Criticisms

One might argue that it's simply not possible, for some context-sensitive sentences, to state in context-insensitive language accurate conditionalized truth-conditions. It would then follow that there's no truth-theory that has such a conditionalized t-sentence as a *theorem*—in particular, none that we cognize. Why think this? Some have suggested that the way the truth-conditions of at least some sentences vary with context is just too *complex* to be captured in a finite number of parameters or is *open-ended* in a way that resists capture in a fixed list of contextual

features.⁶ The argument often proceeds in part via a display of cases meant cumulatively to suggest the unlimited variety of potentially relevant contextual features: the weight of a person's clothes seems sometimes to matter, but sometimes not, to whether she satisfies 'weighs over 150 pounds' (compare a scenario in which we're discussing boxers' weight-classes to one in which we're deciding whether to admit someone onto an already over-loaded lifeboat); that a red leaf was painted green can sometimes matter, but sometimes not, to whether it satisfies 'is green' (compare a scenario in which we're trying to ascertain its species to one in which we're trying to match it to a textile); that the speaker had breakfast yesterday morning can sometimes matter, but sometimes not, to whether an her utterance of 'I've had breakfast' is true (compare a scenario in which we're wondering whether she's hungry to one in which we're wondering whether she's even gotten out of bed before noon); etc.⁷ The display must do more than show that there are *many* ways that context can matter to truth-conditions; for no one will disagree that language and linguistic competence are complex. The work is done by the 'etc.': there must be reason for thinking that it could go on indefinitely.⁸

Showing how one can in fact accommodate the examples produced and those of their ilk is of course the surest way to rebut this kind of criticism.⁹ Pending completion of that task, however, one can at least dilute the collective probative force of the examples by noting the following.

First, not everything 'relevant' to a context-sensitive sentence's truth-conditions on some occasion of use need be captured by an accurate conditionalized *t*-sentence. There may be, for example, many things relevant to something's being what is demonstrated in a use of 'that' on some occasion, but arguably the truth-conditions for the sentence need only contain a variable for what is demonstrated: it needn't bother with what *determines* what's demonstrated. But then it can seem fairly easy *generally* to find *some* parameter to mark a sentence's context-sensitivity. Indeed, when all else fails, mightn't it be an option just to build in reference to the context itself? For instance, even if many things might be relevant to whether someone satisfies 'weighs over 150 pounds,' it might

⁶ Talk of 'open-endedness' might suggest considerations of how a language might develop over time. The proponent of the complexity criticism, however, should resist this way of cashing out the phrase. An account of semantic competence arguably aims to capture something synchronic, not diachronous change.

⁷ Such cases are relied on, for example, by Travis, 1985 (and elsewhere). For discussion from the perspective of relevance theory, see Carston, 2002, Chapter I.

⁸ Arguably, for this criticism to have force, it must be possible indefinitely to continue the display with respect to just *one* lexical item. Suppose one accepts the Context Thesis of Szabó 2001, according to which the context-sensitivity of a complex expression can always be traced to that of a constituent. There are a finite number of lexical items in a natural language and thus a finite number of possible constituents. So, if there's a non-finite number of ways context can affect truth-conditions, this would already have to be the case with respect to just one lexical item. Continuing the display with *other* lexical items would be relevant rather to establishing the *pervasiveness* of the phenomenon.

⁹ See, e.g., Szabó, 2001, in response to Travis.

seem you can nonetheless accurately state its conditionalized truth-conditions along something like the following lines:

For all *u* and *c*, if *u* is an utterance of [Pat weighs over 150 pounds] in context *c*, then *u* is true iff Pat weighs, in the way intended in *c*, over 150 pounds.¹⁰

Second, even if there were a *putatively* context-sensitive sentence for which it seemed impossible to produce an accurate conditionalized t-sentence, the possibility exists of challenging whether it's *in fact* context-sensitive. One might redescribe the case as involving some other sort of content-underdetermination—for example, as involving *ambiguity*, requiring syntactic resolution of the sentence to which truth-conditions are to be assigned;¹¹ or as reflecting the varying *pragmatic* effects of uttering a truth-conditionally invariant sentence. Perhaps predications of 'green,' for example, are ambiguous *inter alia* in sometimes requiring etiologically natural coloring and sometimes not. Or, perhaps they are not ambiguous; instead, utterances of, for example, 'The leaf is green,' despite appearances to the contrary, *do* possess context-invariant truth-conditions: [The leaf is green] is true iff the leaf is green. The appearances to the contrary would then be explained by an account of how something false was uttered in order to communicate something else.¹²

These moves are eminently challengeable. But doing so requires some conception of what constitutes an acceptable encoding of contextual determinants

¹⁰ In Gross, 1998/2001, pp. 66–7, I object that such proposals are covertly meta-linguistic and so fail to provide a statement of truth-conditions knowledge of which would suffice for understanding. (Cf. the argument discussed below in section 7.6.) Perhaps, however, one might respond that, from *our* needing to analyze a locution meta-linguistically in order to understand or introduce it, it doesn't follow that the locution itself, or what it expresses, is meta-linguistic.

I should note also that adverting to context itself as a parameter requires some conception of what individuates contexts. (If one does *not* advert to context itself as a parameter, then the parameters to which one *does* advert—assuming their individuation conditions are clear—themselves *provide* an answer to the question 'what is a context?'.) Cf. Gross, 1998/2001, Chapter 2.

¹¹ I'm assuming here that lexical ambiguity is a syntactic phenomenon: one needs to determine which word has been uttered, not which meaning of one word was intended.

¹² Perhaps one might elaborate upon the truth-conditions by offering an analysis of being green (not to be confused with offering a lexical axiom for 'green')—for instance, maintaining that something is green iff it looks green to an appropriate and appropriately situated perceiver; so that, in the scenario where we're trying to ascertain the species, an utterance of 'The leaf is not green' would be false.

I remark further on these various moves, with reference to cases of 'part' adjectival context-sensitivity in Gross, 1998/2001, Chapter 3. (An adjective is 'part' context-sensitive if its satisfaction depends upon the contextually relevant part of the subject of predication. For example, 'It's green,' said of a certain house, could be true or false depending upon whether (the relevant parts of) its interior or exterior is what's contextually relevant.)

and some conception of how the various sources of content-underdetermination—syntactic, semantic, and pragmatic—ought to be distinguished. Insofar as these are motivated by empirical considerations arising from the project of accounting for semantic competence, the criticism becomes one of the *third* kind: that ascribing cognition of a context-insensitively characterized truth-theory is not empirically supported.¹³

2.2 Indexicality Criticisms

One might argue that certain sentences are such that, although one can capture their truth-conditions context-insensitively, ascribing cognition of a truth-theory that delivers the resulting t-sentences fails to account for the role their indexicality plays in our cognitive economy.¹⁴

Consider, for example, the sentences ‘Steven Gross has won the lottery’ and ‘I have won the lottery’. Standard context-insensitive conditionalized truth-conditions for them might look roughly like this (ignoring irrelevant details):

For all *u* and *s*, if *u* is an utterance of [Steven Gross has won the lottery], then *u* is true iff Steven Gross has won the lottery.

For all *u* and *s*, if *u* is an utterance of [I have won the lottery] and *s* is the speaker, then *u* is true iff *s* has won the lottery.

But now suppose I am an amnesiac reading a newspaper item about the lottery. I utter ‘Steven Gross has won the lottery’. The nurse informs me that *I* am Steven Gross, which leads me to express my newly gained knowledge by uttering ‘I have won the lottery’.¹⁵ On the context-insensitive conditionalized view, my grasp of the relevant contextual features enables me to instantiate the conditionalized t-sentences to yield an assignment of de-conditionalized

¹³ Another kind of complexity criticism stems, not from the putative context-sensitivity of sentences, but rather from the putative sensitivity of some sub-sentential expressions to *sentential* context. Searle, 1983, Chapter 5, for example, claims that examples such as ‘Tom opened the door’ and ‘Sally opened her eyes’ pose problems because, although the literal meaning of ‘opens’ doesn’t vary across these sentences, the term’s contribution to truth-conditions does. (Searle appeals as well to cases in which it’s *indeterminate* whether a term applies, but this is a distinct matter—even if such indeterminacy can have its source in the complex context-sensitivity of the relevant terms.) Again, some of the moves mentioned above may be deployed in an attempt to deny the claim. For example, one might invoke ‘sense-generalization,’ identifying the semantic value of the term with a determinable, allowing that on different occasions claims concerning more specific determinants can be pragmatically communicated. (On sense-generalization, see Atlas, 1989; on related cases of polysemy, see Pustejovsky, 1995.)

¹⁴ This criticism is, for example, part of Ludlow’s case (1999, Chapter 6) against untensed metalanguages. (Classic discussions of the relevant feature of indexicality include Castañeda, 1966, and Perry, 1979.)

¹⁵ Exclaiming joyously is a kind of uttering in the sense intended.

truth-conditions for these particular utterances. But then, since Steven Gross *is* the speaker, it would seem that the utterances have identical truth-conditions: any difference between the two utterances disappears.¹⁶ Let u1 refer to my utterance of ‘Steven Gross has won the lottery’ and u2 refer to my utterance of ‘I have won the lottery.’ The utterances’ de-conditionalized truth-conditions would seem to be:

U1 is true iff Steven Gross has won the lottery.

U2 is true iff Steven Gross has won the lottery.

Thus it would seem that the truth-theory fails to capture how these utterances differ in virtue of the uttered sentences’ differing semantic properties.

It may be reasonably objected that the t-sentence just suggested for u2 might not accurately display how *I* would represent who the speaker is. Now that I know that I am Steven Gross, and given how salient this fact is, the first-person pronoun might better capture how I represent the value that the contextual parameter *who the speaker is* has on this occasion. An expression of the t-sentence that displayed how *I* represent the utterance’s de-conditionalized truth-conditions would then be:

U2 is true iff I have won the lottery.

Allowing context-sensitivity into the statement of de-conditionalized truth-conditions is consistent with excluding it from the statement of the speaker’s standing semantic knowledge. So, in raising this objection, a defender of *Meta-Insensitivity* doesn’t on that account give away the game. But it may be responded that, even though it may be more accurate to deploy the first-person pronoun in representing who I hold the speaker to be, it’s not the *truth-theory* that requires this. And so it would remain the case that there’s a failure to capture how the utterances differ in virtue of the semantic properties of the uttered expressions.

One could still object to this kind of criticism on various grounds. For example, one might challenge the claim that the utterances’ difference reflects a *semantic* difference, or the claim that this difference should be captured in particular in the cognized truth-theory, the ascription of which might be claimed only to be *part* of an explanation of semantic competence. Again, insofar as empirical considerations drawn from the project of accounting for semantic competence are adduced in reply, the criticism trails into the third category.

2.3 Empirical Criticisms

Finally, one might argue that empirical evidence simply doesn’t favor attribution of cognition of a context-insensitively characterized truth-theory. Even if,

¹⁶ Because the utterances contain different words, there’s a different *route* to the derived truth-conditions. But this doesn’t show up in what I thus come to know.

using only context-insensitive terms in the meta-language, one *can* state accurate truth-conditions for context-sensitive sentences, and even if there's no requirement (empirical or otherwise) that, for certain sentences, one must reflect in the truth-theory the highlighted aspect of their indexical nature, it still doesn't follow that our semantic competence *in fact* involves cognizing a context-insensitively characterized truth-theory. If empirical evidence is lacking for positing the requisite contextual variables or denying particular claims of context-sensitivity, and if evidence supporting some alternative exists, it's irrelevant whether criticisms of the first two kinds can be met.

Thus, for example, relevance theorists have argued that positing contextual variables to handle problematic cases is *ad hoc*, as is redescribing them as involving syntactic or pragmatic underdetermination—while their alternative can handle such cases uniformly by appealing to independently motivated mechanisms.¹⁷ Proponents of truth-theories, for their part, have attempted to provide independent motivation for hypothesized contextual variables. Stanley and Szabó, for example, argue that binding requirements support the presence of certain phonologically unrealized contextual variables in LF. The situation here remains unsettled.¹⁸

Turning to another example of the third kind of criticism: Peter Ludlow has adduced psycho-linguistic evidence against standard truth-theoretic treatments of tense and temporal indexicals. Such treatments accommodate this kind of context-sensitivity by positing contextual variables for times in a context-insensitively characterized truth-theory—yielding, for example, t-sentences such as 'For all, u, t, and r, if u is an utterance of [Emma had coughed] at t and r is the contextually given reference time, then u is true iff there's an event done by Emma that's a coughing that occurs before r and r occurs before t.' Ludlow argues, however, that such treatments don't square with some suggestive developmental and acquired impairment data, which suggest that

¹⁷ According to relevance theorists, automatic processes that tend to maximize positive cognitive effect and minimize processing effort generate a content from an LF. Because positive cognitive effect and processing effort vary with context, so does the generated content. In generating this content from an LF, the automatic processes don't make use of any truth-theoretically organized semantic information. See Sperber and Wilson, 1986/95, and Carston, 2002. For the charge of ad hocness, see, e.g., Sperber and Wilson, 2002.

¹⁸ See Stanley and Szabó, 2000—and, for some responses, see Bach, 2000, Carston, 2000, and Brenehy, 2002. This is part of a much broader debate: one wouldn't expect such empirical matters to turn simply on one sort of consideration. For example, Stainton, 1994, argues for a relevance theoretic approach to non-sentential assertions, and Carston, 1988, argues that various sorts of pragmatic intrusion pose problems for truth-theorists but not for relevance theorists. Truth-theorists reply that such criticisms rely on an impoverished understanding of the relevant syntax and semantics—see Ludlow, forthcoming, and Stanley and King, forthcoming, respectively. Further, truth-theorists (and others) argue that relevance theory provides insufficient constraints to explain why certain utterances *can't* be used to communicate various claims—see, e.g., Stanley, 2002.

competence with complex tenses doesn't require mastery of the words for such concepts as BEFORE.¹⁹

This example is more interesting for my purposes because, unlike relevance theorists, Ludlow doesn't advance this criticism as part of an attempt to support an *alternative* to truth-theoretic accounts of semantic competence. Rather, he argues for an alternative truth-theoretic account: in particular, one that allows tensed—and thus context-sensitive—language in the meta-language. My interest, recall, is examining how truth-theorists might respond should any of these criticisms prove correct. Abandoning truth-theoretic accounts of semantic competence altogether is but one option. Another is to pin the blame on *Meta-Insensitivity*. Insofar as there is *some* reason to think that cognition of a truth-theory is involved in semantic competence (e.g., because of the explanation it would provide of creativity or systematicity), there is some reason to explore this alternative.

The three kinds of worry indicated here might not all pan out (indeed, *none* might). It's thus worth bearing in mind that, in supplying *different* reasons for considering the use of context-sensitivity in the meta-language, they also suggest different places where it might be needed. The best cases for complexity criticisms (e.g., complexly context-sensitive adjectives) are not best cases for indexicality criticisms, and the best cases for indexicality criticisms (e.g., the first-person pronoun) are not best cases for complexity criticisms. There doesn't seem to be any 'essential' indexicality in color predications—while providing a *t*-sentence that at least accurately states the truth-conditions for a sentence containing the first-person pronoun seems quite straightforward (assuming no other constituent of the sentence gives rise to problems). Just which truth-theoretic axioms might call for context-sensitive characterization, and accordingly (as we'll see) just what problems this might introduce, depends on where continued discussion of these worries leads.

3. Side-Stepping the Criticisms While Preserving Compositionality

How might dropping *Meta-Insensitivity* help one respond to the three kinds of criticisms canvassed? Allowing the use of context-sensitive expressions in characterizing cognized truth-theories opens up the possibility of capturing object-level context-sensitivity by *coordinating* it with context-sensitivity in the meta-language. The strategy would avoid complexity criticisms, since there's no attempt to parameterize *all* semantically relevant contextual features and so no occasion to worry whether the task can be successfully completed. One might then say: something satisfies 'green' iff it is

¹⁹ Ludlow, 1999, pp. 138–41. This is not Ludlow's central point against standard treatments of tense—for instance, the indexicality criticism (and its relation to some further points) plays a more prominent part in his discussion. But he notes (p. 138) that, even if one were to ignore his other points (or, presumably, if one could rebut them), this suggestive psycholinguistic evidence would support his proposal over standard treatments. Of course, the psycholinguistic evidence too might be subject to rebuttal: perhaps it's possible that the semantic module sub-personally deploys the concept BEFORE before the child herself has mastered and lexicalized it. But my point here is to display a *kind* of criticism, not to endorse the particular example supplied.

green.²⁰ Likewise, once context-sensitive expressions may be deployed in the meta-language, there's no bar to representing what the speaker cognizes about the target sentence in a way that captures the relevant aspects of its indexicality. Bracketing an obvious issue to be joined presently, the t-sentence for 'I am tall' could read: 'I am tall' is true iff I am tall.²¹ Finally, empirical challenges to posited parameters are irrelevant to accounts that don't posit the challenged parameters. Of course, there's now the task of empirically supporting the ascription of context-sensitively characterized truth-theories. But, pending evidence to the contrary, this can be seen as a project, not a problem. Since the proposal isn't much considered, this remains largely uncharted territory.²²

²⁰ Let me stress that the suggestion is *not* that one refrain from positing *any* contextual parameters. Dropping *Meta-Insensitivity* merely makes available further resources. It's consistent *not only* with a divide and conquer strategy (perhaps there's good reason to introduce parameters in certain sorts of cases and good reason to handle others using meta-linguistic context-sensitivity—and, moreover, perhaps good reason to treat other cases in some third fashion), *but also* with combining parameterization and the use of meta-linguistic context-sensitivity in the treatment of a *particular* context-sensitive term or construction. It's thus possible, for example, to hypothesize contextual variables in t-sentences involving color predications and yet *also* allow the color term in the meta-language to be context-sensitive. Good reasons for positing such parameters—and perhaps further for positing phonologically unrealized variables in LF—can thus be absorbed into this approach.

In the case of contextually determined quantifier domain restrictions, this might provide a response to one objection I raise in Gross, 2004, to Putnam's critique of ontology. Putnam argues that ontological questions raised by philosophers are nonsensical, any impression to the contrary arising from a failure to appreciate fully the context-sensitivity of quantifier phrases and related terms. *Inter alia* I question whether his position can be squared with current approaches to accommodating context-sensitivity in accounts of semantic competence. The present suggestion may allow him to do so.

²¹ At least this is how the t-sentence would be characterized from the perspective of the cognizer herself—similarly for other examples I will use. Among the complexities introduced by use of context-sensitivity in the meta-language is that a change in ascriber can alter how the truth-theory must be characterized. Presumably, ascribed from *my* perspective, the axiom *you* cognize would be characterized as follows: 'I am tall' is true iff *you* are tall. (The need for such shifts parallels, but is distinct from, the possible shifts discussed below between axioms used in production and those used in comprehension.) Sample t-sentences provided below should all be understood as being presented from the cognizer's perspective.

Note that it's a further thesis held by some proponents of cognitivist truth-theoretic accounts of semantic competence that it's nomologically the case that for speakers to cognize truth-theories is for them to represent them in a Language of Thought (LOT). An adequate characterization of cognized truth-theories would then be one that captured in natural language more-or-less how speakers represent them in LOT. Given this, characterizing truth-theories from the cognizer's perspective is particularly natural.

(I return below in 7.5 to the consequences of dropping *Meta-Insensitivity* for theorists ascribing truth-theories.)

²² The only sustained exception of which I'm aware is Ludlow's defense (1999) of a tensed meta-language. Brief discussion of using context-sensitivity to characterize semantic competence with respect to 'part' adjectival context-sensitivity can be found in Gross, 1998/2001, pp. 70–5. The possibility is mentioned in passing in Lepore and Loewer, 1989, regarding attitude ascriptions. It no doubt crops up elsewhere as well. (For example, as a referee points out, in Rumfitt, 1993, pp. 440–4. Rumfitt proposes an account of indirect speech reports that, if I understand it aright, builds into the truth-theory itself an *analysis* of 'same-saying' derived from a remark of Davidson. This then requires using in the *truth-theory* the kind of context-sensitivity that in 2.2 arose in *deconditionalizing* a t-theorem.)

But we can here at least note that dropping *Meta-Insensitivity* doesn't threaten compositionality: it can remain the case that a complex expression's semantic value is a function of its semantically relevant syntactic structure and the semantic values of its (immediate) constituents, even if context-sensitive expressions are used to characterize the conditions under which it has a certain semantic value. To see this, we must first distinguish the *sentences* used to characterize a truth-theory, and what is thereby characterized—the truth-theory itself. This is crucial once context-sensitivity is admitted into the characterization, since for that very reason the truth-theory *sentences* can characterize *different* truth-theories in different contexts. The truth-theory sentences comprise the sentences used to characterize the lexical and combinatorial axioms as well as the sentences that can be generated from them by applications of the production schemata. It's most plausible to assume that any context-sensitivity found in the truth-theory sentences can be traced to context-sensitivity in the sentences for the lexical axioms and *perhaps* in the sentences for the combinatorial axioms—as opposed to its being introduced by the production schemata. To drop *Meta-Insensitivity*, then, is to allow that the axiom sentences and the sentences generated from them can express different theorems—different valuation-conditions—on different occasions of use. But in *each* context what these meta-level sentences express can be valuation-conditions conforming to compositionality. Each truth-theory expressed will then assign expressions valuation conditions as a function of their syntactic structure and the valuation conditions of their constituents. Given how the world is, their semantic values will thus likewise compose.

Thus, one important reason some favor truth-theoretic approaches to semantic competence in the first place—viz., that they can compositionally explain such features of our semantic competence as its creativity and systematicity—remains in place. The ascription of context-sensitive truth-theories thus retains at least this much empirical support. And so there's at least this much motivation for examining their problems and prospects.

4. A Problem

The most obvious problem facing context-sensitively characterized truth-theories is accounting for our ability to comprehend certain utterances of others—for example, my ability to comprehend *your* utterance of 'I am tall'. At least, this is a problem on the most natural first proposal for exploiting context-sensitivity in the meta-language (and on the supposition that 'I' is among the object-level terms to be so handled).

In the previous section, I said that dropping *Meta-Insensitivity* opens up the possibility of capturing object-level context-sensitivity by *coordinating* it with context-sensitivity in the meta-language. A natural idea is to coordinate object-level and meta-level context-sensitivity by using in the truth-theory expressions whose context-sensitivity *parallels* that of object-level context-sensitive expressions—that

is, whose semantic values contextually vary *in the same way*. Abstracting from details of implementation, we might suppose that a truth-theory should yield theorems expressed by such t-sentences as ‘[I am tall] is true iff I am tall’.²³ This would seem to work well enough in an account of language *production*, since in this case—perhaps with some tweaking—we can guarantee that the target-sentence and the t-sentence are interpreted relative to the same context of utterance.²⁴

In addition to handling language production, this would seem to work as well in at least a fair range of cases of comprehending someone *else’s* use of context-sensitive expressions. I can understand what your words ‘That is tall’ express if I (or my semantic module) can use context-sensitive expressions in parallel fashion to yield the t-sentence ‘[That is tall] is true iff that is tall’. Of course, I might sometimes *fail* to use them in a parallel fashion—for example, I might understand the relevant contextual determinants of content otherwise than you do (mistakenly taking you to be demonstrating *that* thing instead of *this* one, for instance) and thus assign inaccurate truth-conditions to the uttered sentence. But the possibility of miscommunication is a fact of life that any account of semantic competence must accommodate.

The problem, however, is that an agent sometimes *can* understand someone else’s words even when the agent’s use of the term would *not* assign the same semantic value as the speaker’s: that’s what happens when I understand *your* utterance of ‘I am tall’. The truth of that sentence out of *your* mouth doesn’t depend on *my* height. But assigning it truth-conditions using parallel context-sensitivity in the meta-language would seem to have just that result. The challenge, then, is to show how cognition of context-sensitive truth-theories can play a role in explaining *this* aspect of our semantic competence.

The problematic sentences, note, are those containing context-sensitive constituents whose contextual determinants at least in part ‘automatically’ fix their

²³ The meta-language and object-language needn’t be the same, so it’s not being suggested that to use a parallel context-sensitive expression in the meta-language is necessarily to use the very expression that appears in the target-sentence.

²⁴ Cf. Davidson, 1967, pp. 33–4. Whether tweaking is needed depends on how ‘[I am tall]’ is understood—in particular, whether it refers to a type or token. If it refers to the very token (to be) produced, then the guarantee is in place. (And if it refers to some other token (or, some other token not appropriately related to the produced token), then, although the guarantee is no longer in place, we no longer have a case of *production*.) But if it refers to the type, then, since context-sensitive sentence-types, considered independently of any particular occasion of use, aren’t themselves either true or false, the form of the t-sentence must be modified so as to tie the type to the very context relative to which the t-sentence is interpreted. For example, one might modify the proposed t-sentence to read ‘[I am tall]’, understood as if uttered in this very context, is true iff I am tall’. (More radically, one might contemplate abandoning the constraint that the t-sentence must *itself* be true; so long as cognizing it yields the right results, its truth might be neither here nor there. Then one might countenance a t-sentence that in different contexts assigned a sentence-type—falsely, let’s suppose—different truth-conditions. I question the necessity of *true* t-sentences for the purposes of a truth-theoretic account of semantic competence in Gross, forthcoming.)

semantic value independently of speakers' intention—for example, 'here', 'now', tense morphemes, etc.²⁵ Insofar as contextual determinants are a function of speakers' *intentions*, it's possible for a hearer to use the relevant context-sensitive expression in the same way by appropriately configuring her own intentions. This is because our intentions are largely under our control. This is not an option, however, when the contextual determinants are fixed automatically. But it doesn't follow that a speaker and hearer can't *ever* use these 'automatic' terms in the same way: they can when the automatic contextual determinants relative to each happen to agree—for example, when we simultaneously refer to the present moment using 'now'.²⁶ In sum, the problem thus arises, not for the comprehension of *every* context-sensitive expression, but only for automatic ones; and not for *every* use of an automatic expression, but only for some: those for which the speaker's and hearer's use of the term would not agree, but there's no barrier to comprehension.

It's worth emphasizing some obvious consequences of this. First, the problem does *not* arise for the use of meta-linguistic context-sensitivity in handling the complex, open-ended context-sensitivity claimed for examples like 'green' above. There's no in principle bar to a comprehender's using these terms just as the speaker does, and so no in principle bar to the comprehender's so using the terms in assigning truth-conditions. Second, the problem simply doesn't arise *at all* if truth-theoretically characterizing our competence with *automatic* context-sensitive terms does *not* require positing meta-linguistic context-sensitivity. As noted above, it's possible for someone to advocate the use of context-sensitivity in the characterization of cognized truth-theories, but not for *these* terms.

But if one does contemplate using context-sensitivity to characterize the truth-theoretic aspect of our semantic competence with automatic terms, one must address this problem—what for ease of reference I'll call, despite the attempt above to limit its extent, the *comprehension* problem. I comment now on two responses briefly suggested by Ludlow: first, sometimes coordinating the use of automatic terms with complementary, rather than parallel, context-sensitivity; and, second, appealing to simulation theory.²⁷

²⁵ I say 'in part' because, for example, the semantic value of 'here' and 'now' are not automatic with respect to their *extent*. (My formulation could still use some qualification. For even the automatic aspects needn't be completely independent of speaker's intentions: I may be here—and thus my use of 'here' may have the semantic value it does—in part because of my intention to be here, which might even be an intention I have in virtue of my intention to use the word 'here' in that way. One must distinguish, among the speaker's intentions, those relevant to the determination of semantic value *in the relevant way*.)

²⁶ This never occurs with 'I', but that's something peculiar to it. (Cf. Millikan's argument (1990) that, since there's no variation in reference, there's no essentially indexical first-person representation in thought, but rather a referentially-invariant SELF concept. Millikan also discusses the harder cases of the concepts HERE and NOW. For some critical discussion, see Ezcurdia, 2001.)

²⁷ Ludlow, 1999, pp. 63–4, presents the problem as one of comprehension *generally* for sentences containing certain terms—not just as a problem for comprehending certain utterances of them.

5. Complementary Context-Sensitivity

Above I said that a natural first thought is to employ in the meta-language context-sensitivity *parallel* to that in the object-language. But stating a sentence's truth-conditions using context-sensitive language that varies in context *in the same way*, we've seen, can lead to the wrong results when the contexts relative to which the target sentence and the t-sentence, respectively, are to be interpreted differ. A first response abandons the idea that appropriately *coordinating* object-level and meta-level context-sensitivity always requires that the context-sensitivity be *parallel*. What is needed in such cases, rather, is an expression whose context-sensitivity is appropriately *complementary*. For example, I can assign accurate truth-conditions to your utterance of 'I am tall' if I employ the t-sentence '[I am tall] is true iff you are tall'. Such complementary context-sensitivity in the t-sentence, however, yields the wrong results in language *production*. So, some way is needed to distinguish when *parallel* context-sensitivity should be used and when *complementary* context-sensitivity should be used. One idea is to build this into the relevant lexical axioms by making their semantic valuations conditional on whether speech is being produced or comprehended. The lexical axiom for 'I' might then look something like:

If I am uttering it, then [I] refers to me; but if I am comprehending your speech, then [I] refers to you.²⁸

It's not obvious, however, how best to develop this strategy. I'll mention three issues.

First, the division between production and comprehension doesn't always neatly align with when *parallel* context-sensitivity is needed in the meta-language and when *complementary* context-sensitivity is needed. *Modulo* issues to come, our bifurcated lexical axiom for the first-person pronoun works well enough, since one person can never use 'I' to refer to what another would in using it.²⁹ But, as we've noted, *sometimes* an automatic context-sensitive expression would have the *same* semantic value whether used by the speaker or the hearer. In these cases it would seem that comprehension also requires parallel context-sensitivity in the t-sentence. Thus the axiom:

If I am uttering it, then [here] refers to here; but if I am comprehending your speech, then [here] refers to there,³⁰

²⁸ Ludlow, 1999, credits something like this idea to Nolan, 1970.

²⁹ There are potential counter-examples. If I and my addressee know that A has mistaken B for me, I might describe A's thoughts about and actions towards B using first-person pronouns: 'Look, now John is waving at me'. Again, if C is in some sense standing in for me (perhaps he's a stage-hand standing where I would stand, as the director and I confer on blocking), I might use the first-person pronoun—at least seemingly—to refer to C: 'But now I'm standing out of the spotlight!' But suppose such cases (of shifted and deferred reference?) are otherwise handled.

³⁰ This is the axiom Ludlow, 1999, in fact supplies as an example in his brief mention of this response to the comprehension problem.

would seem inappropriate for cases in which I'm where you are—at least insofar as 'there' must refer to some place other than where we are. One might conclude that the bifurcation of the relevant lexical axioms ought not to turn on whether we have a case of production or comprehension: we need some other way of indicating when the parallel clause should be used and when the complementary clause. But it's difficult to see how one could specify such conditions without adverting to whether the object-level and meta-level expressions agree in semantic value; and an account of comprehension shouldn't presuppose this ability to comprehend.

A response to this problem is to continue bifurcating the relevant axioms according to whether the utterance is being produced or comprehended, while employing in the comprehension clause's consequent context-sensitivity that is both complementary and *inclusive*. By the latter, I mean context-sensitive expressions whose properties don't *preclude* their agreeing in value with the parallel context-sensitive term used in the production clause's consequent. Indeed, perhaps 'there', contrary to what was suggested above, can itself be such a term. But if one balks at allowing someone to use 'there' to refer to here (at least if she *believes* that 'here' and 'there' would then be on this occasion co-referential), then instead one might coin a term without this restriction ('there*') or employ in the t-sentence the complementary, inclusive phrase 'where you are'. If this approach is to scale up, something similar must be available for the bifurcated t-sentences of other automatic terms as well.³¹

A second issue is the possibility of a mismatch between whose speech is being comprehended and to whom the comprehender does (or would) refer in

³¹ Ludlow, 1999, though he mentions bifurcating axioms as a way of addressing the comprehension problem, doesn't then develop its application to his own use of a tensed meta-language. His axioms for simple tense morphemes are as follows:

Val(x, PAST) iff x was true.
 Val(x, FUT) iff x will be true.
 Val(x, PRES) iff x is true.

But suppose we wish to accommodate comprehension of past utterances. Perhaps bifurcated his axioms might read:

If I'm speaking, then Val(x, PAST) iff x was true; if I'm comprehending your utterance, then Val(x, PAST) iff x was true before you spoke.
 If I'm speaking, then Val(x, FUT) iff x will be true; if I'm comprehending your utterance, then Val(x, FUT) iff x was or will be true after you spoke.
 If I'm speaking, then Val(x, PRES) iff x is true; if I'm comprehending your utterance, then Val(x, PRES) iff x was true when you spoke.

This suggestion mimics Ludlow's use of when-clauses to accommodate complex tenses (and temporal anaphora). It assumes, however, that comprehension always treats utterances as having been entered in the past—that is, comprehension, on this suggestion, would never treat utterances of others as having been entered in the same (temporally extended) 'now' in which comprehension occurs.

using 'you'. Suppose, for example, you and I are eaves-dropping on some third party's musings while engaged in conversation with one another. 'I'm ready to order', the third party says. It's not the case that her utterance is true iff *you* are ready to order.

It might be responded that my use of 'you' can shift rapidly, as my attention shifts from you—my interlocuter—to the third party whose speech I'm trying to understand. Indeed, one might allow that I can simultaneously use 'you' in two different ways: I refer to you in asking 'Can you also hear her?' while (let's suppose) my semantic module tacitly deploys a second-person representation referring to the speaker in comprehending her speech.³²

Alternatively, it might be responded that the mismatch poses no problem for the proposed bifurcated axiom, because the axiom covers when *I* speak and when *you* speak. If my use of 'you' in the meta-language wouldn't refer to you, but to the speaker, then it's irrelevant that *had* you been the speaker, 'I' would have referred to you. The obvious problem with this reply, however, is that it amounts to an admission that the axiom fails to cover all cases. So, perhaps we should augment the reply with a proposal that the axiom contain a further clause: if I'm comprehending *her* (or *his*) speech, then 'I' refers to her (or him).³³

The third issue is whether any of these proposals adequately address the concerns raised by the indexicality criticism sketched above.³⁴ The criticism was that sentences such as 'Steven Gross has won the lottery' and 'I've won the lottery' express beliefs that differ in their upshot for cognition and action, which should be reflected in an account of semantic competence—but context-insensitively characterized truth-theories fail to capture such differences. Now, the proposed bifurcated axioms might capture the difference as far as *production* goes. But this is less clear with respect to comprehension. It is true that all of the variations so far mentioned deploy context-sensitive language in the comprehension clause as well. But it doesn't follow from the mere fact that context-sensitive expressions are used that the axiom encodes the relevant functional impact of the object-level indexicals. For example, though it's the case that, when you speak, 'I' refers to you, it's not clear that my cognizing this enables me to represent the special first-personal way you express yourself.

It might be responded that it's only in one's *own* case—that is, in production—that capturing this aspect of indexicality is crucial. But it's hard to see how one

³² It would seem there's no bar in principle from *explicitly* using 'you' simultaneously in two ways. To call in my debts efficiently, I might telephone two people on separate phones and then say 'You owe me \$100' into both mouthpieces at once. I would thus simultaneously occupy two conversational contexts.

³³ This suggests another possibility. Perhaps we should just dispense with the 'you' clause altogether, letting the 's/he' clause cover the cases for which the 'you' clause was intended. 'Mismatch' worries arise again, but responses parallel to those offered for the 'you' clause could be made.

³⁴ Of course, this is not an issue if one's not motivated to drop *Meta-Insensitivity* by indexicality worries.

might motivate this. If it's important that the truth-theory reflect the role of indexicality in one's own cognitive economy, then it would seem important that it reflect its role in that of others too. What others say provides access to their cognitive economy, and what you thus take the state of their cognitive economy to be is a part of (and thus affects) your own. If correctly grasping the differences in others' beliefs doesn't require that the relevant functional differences be captured in the truth-theory's *comprehension* clauses, then why must they be captured in the *production* clauses? An asymmetry threatens to undermine the indexicality critic's insistence that the functional difference correspond to something *semantic*, more specifically something captured in a cognized truth-theory.

Alternatively, one might respond by modifying the proposed axioms. For example, one might add reflexives to capture the relevant indexicality: if I'm comprehending your (her) speech, then 'I' refers to you yourself (she herself). It might be objected that this fails fully to capture the phenomenon. Suppose, for example, that you wish to warn me of some imminent danger where we are co-located. You say, 'There are landmines here'. It seems that the full import of what you say is not captured by my cognizing that what you say is true iff there are landmines where you yourself are. Your point is that there are landmines *here*. But perhaps a proponent of the present response can consistently maintain that those truth-conditions capture the indexical aspect of what you said (how it's essentially tied to your perspective), even though it requires extra-truth-theoretic information (to the effect that you yourself are *here*) for me fully to appreciate your warning. That my coming fully to appreciate your warning requires further inference needn't tell against my assigning appropriately indexically-expressed truth-conditions to your utterance.

Perhaps, then, there are ways of addressing each of these three issues. But it's arguably not enough to come up with ways to generate workable t-sentences. The suggested maneuvers can seem *ad hoc* (especially those advanced in reply to the first two issues); so it would be nice to have some independent empirical evidence as well suggesting that speakers do indeed cognize such truth-theories. I don't have any to supply, but neither do I have reason to believe that none could be found. Such independent evidence would be needed as well if one wanted to argue for this approach in favor of the alternative to which we now turn.

6. Simulation

A second way one might respond to the comprehension problem is to appeal to 'simulation' theory. As Ludlow puts it:

... one strand of thinking would be that when we interpret others we try to 'simulate' their thoughts. That is, perhaps when we interpret another's

utterance we try to 'project ourselves into' that person's egocentric space. If that is how we interpret others, then of course we want the axioms to reflect the egocentric perspective of the speaker. Anything less would get in the way of our ability to simulate, and hence understand, our interlocutor.³⁵

Unlike the first strategy, this suggestion would retain the idea that the meta-language exploits only *parallel* context-sensitivity. But it offers an alternative understanding of how the cognized truth-theory is exploited in comprehension. Just what this alternative is—how a truth-theory might be used in simulating another speaker—requires clarification, as psychologists and philosophers have used the term 'simulation' in a variety of ways.

Talk of simulation sometimes refers to something we are capable of doing consciously and deliberately. Applied to the interpretation of others, this would involve a replication in reality, imagination, or otherwise in thought of some aspects of another's situation and accompanying cognition—as when I come to experience what you experience by exposing myself to the same kind of physical forces to which you are exposed; or I come to a conclusion about what you can see by imagining myself situated and oriented as you are; or I come to a conclusion about what you believe by consciously considering what I would believe in your circumstances. But, though comprehension of another's speech may indeed sometimes be accompanied by and perhaps involve such simulation, it's clearly neither required nor done in many cases, including many of the problematic cases.³⁶ So, this would seem an unpromising direction in which to develop the simulation response.

A second way theorists talk of simulation is as part of an account of the sub-personal information processing—for the most part inaccessible to consciousness—that explains our various cognitive capacities. So construed, the simulation response would refer to the particular way cognized truth-theories—themselves largely inaccessible to consciousness—are utilized in the information processing that enables semantic comprehension.

What makes a bit of sub-personal processing count as a simulation in this sense? This question could be answered in various ways, but a usefully broad answer is that a process is a simulation when its components are run 'off-line'—that is, in an appropriately functionally modified way.³⁷ This can be illustrated with what's perhaps the most familiar example of a simulation theory: it's hypothesized that the prediction of others' intentional behavior

³⁵ Ludlow, 1999, p. 63. On simulation theory generally, see the articles collected in Carruthers and Smith, 1996, and Davies and Stone, 1995.

³⁶ Even when conscious simulation does accompany comprehension, it's a further question whether the simulation contributes to the comprehension of sentences' truth-conditions as opposed to a fuller understanding of what is thereby comprehended.

³⁷ See Nichols *et al.*, 1996, pp. 41–5. More generally, I shall draw upon their particularly clear framing of debates concerning simulation (developed as well in other papers by Stich and his co-authors).

involves running the sub-personal mental components responsible for one's own intentional behavior, but with inputs modified to reflect the other's current mental state and with outputs leading, not to the initiation of action, but rather to one's forming a belief about what the other will do. Thus, on this hypothesis, one could be said to predict what another will do by sub-personally simulating her to see what one would do oneself if one were relevantly like her and similarly situated. On the broad characterization, however, simulation is not *restricted* to hypotheses concerning our capacity to predict and explain the intentional states and actions of others. It covers as well accounts of other sorts of capacities—for example, the hypothesis (Currie, 1995) that mental imagery is produced by running off-line components implicated in visual experience.

The suggestion would then be that semantic comprehension involves using, within a sub-personal simulation, truth-theories implicated in language production. There are various way one might develop this idea further. I'll sketch two. The first applies to semantic comprehension the simulation approach to action explanation generally. Given the centrality of action prediction and explanation in discussions of simulation, this is perhaps the most natural first thought one has in invoking it. The broad characterization, however, invites us to consider invoking simulation without appeal in particular to a simulation account of action explanation. The second idea accordingly points to a more specifically semantic simulation strategy.

6.1 Simulation and Action Explanation Generally

The first idea assumes that one already hypothesizes a simulation theory for the understanding of others' actions generally, and then sees semantic comprehension as an element of a particular instance of this: the explanation of another's *linguistic* behavior. To see how this would go, we need first to say something about how simulation might underlie generally the capacity to explain others' actions. Above, we mentioned the hypothesis that simulation might underlie the capacity to *predict* others' actions. But it's not as obvious how simulation might enable action *explanation*. Prediction begins with some conception of the other's mental state and situation and uses off-line simulation to yield a likely action. But things are reversed with explanation: it begins with the action and proposes the mental state that led to it. An off-line simulation using one's own sub-personal practical reasoning might thus seem doubly useless: we lack the relevant inputs (they're what we're after, so if we had them we wouldn't need the simulation after all), and the output would be what we already know. How then might simulation be involved in the explanation of others' actions? One suggestion is that we have a component that somehow generates *hypotheses* concerning the relevant mental states of others, which serve as inputs to an off-line simulation. The function of the simulation is then to *test* hypotheses by seeing if the action to which they would lead matches the action observed. If it doesn't, the hypothesis generator is sent back to work; if it

does, the hypothetical mental state is embedded in an ascription and sent to the belief box.

To apply this to semantic comprehension, we need some sense of how language production might proceed—in particular, what role truth-theories might play according to cognitivist truth-theoretic approaches. A very crude account of the relevant aspects of language production might run as follows.³⁸ An intention to communicate that *P* is formed, which, perhaps given the speaker's (possibly tacit) beliefs about how communication might be effected, leads to the formation of an intention to utter a sentence with the appropriate truth-conditions. In order to identify such a sentence, the cognized truth-theory is then accessed, and an LF is found—call it *S*—for which the truth-theory yields the *t*-sentence: *S* is true iff *P*. Assuming the communication is to be effected via speech, the intention to utter PF(*S*)—the phonological form of *S*—is then formed, which leads the speaker's articulatory systems to produce the appropriate sounds.

Simulation, then, might enable semantic comprehension as follows. An hypothesis concerning the relevant speaker's intention is generated. This intention—rather than our own—is fed into the systems that would normally lead to our own linguistic behavior. The cognized truth-theory is accessed, as above, to yield an LF. But instead of this leading to the production of certain sounds, the result is sent to a system that checks the PF against the actual sound produced. If it is insufficiently similar, the hypothesis generator is sent back to work; if it is sufficiently similar, the hypothetical intention is embedded in an ascription to the speaker and sent to the belief box. So, the hearer winds up believing that the speaker intended to communicate that *P*.

Of course, the whole point of invoking simulation is to help with the comprehension problem. Does it on this approach? Only if we posit that somewhere along the way certain context-sensitive expressions are replaced with appropriately complementary context-sensitive terms. For suppose the speaker utters the sentence 'I am tall' with the intention to communicate that she (herself) is tall, and suppose the hypothesis is indeed generated that the speaker intends to communicate that she (herself) is tall. Feeding *this* intention into the process that leads to speech production yields the PF <She is tall>, which fails to match the sounds in fact produced. If I am the hearer, then a match is achieved only if the input is the intention to communicate that *I* am tall. But we don't want the intention to communicate that *I* am tall to be embedded in an ascription and placed in the belief box. For this simulation to yield the right result we need some way to input an intention to communicate something about *my* height that will nonetheless enable us to conclude that the speaker intends to communicate something about *her* height.

Here are two ways this might be done. One possibility is that the generated hypothesis—that the speaker intends to communicate that *she* is tall—is transformed

³⁸ Among the many ways this is very crude is in its ignoring the calculation of pragmatic effects.

into the intention to communicate that *I* am tall prior to being fed into the language production systems. Then, after a match is achieved, the originally hypothesized intention can be placed in the belief box. The second possibility supposes that the hypothesized intention generated in the first place is already formed as if *I* were the subject under consideration. This wouldn't avoid the need for transformation, however. For transformation would still be required before the intention could be placed in the belief box. Either way, comprehension would require the ability to effect transformations between my own and the speaker's perspective at least with respect to the concepts whose expression gives rise to the problematic cases. This ability, however, would not be captured within the cognized truth-theory *itself*. The simulation proposal differs from the previous strategy in requiring only *parallel* context-sensitivity to characterize the cognized truth-theory, but that's because the work done by complementary context-sensitivity is done elsewhere. If one thought that the simulation hypothesis would somehow free the comprehender from needing at some point to use some other term to comprehend the producer's use of an automatic expression (at least in cases in which context would assign their uses different values), this was an illusion. It only frees the cognized truth-theory of the task.

So, we have one proposal for accounting for semantic comprehension as an element in a more general sub-personal simulation account of action explanation. No doubt there are other ways—in part because there are no doubt other ways of developing the general simulation account of action explanation. Indeed, it would be good if there were other ways. In particular, the intention-hypothesis-generator on the current proposal can seem an embarrassingly dark black-box, especially in this application to semantic comprehension. Surely what typically enables us to generate an hypothesis concerning a speaker's communicative intentions is that she produced those words with that meaning (in this context).³⁹ How can an hypothesis be generated *prior* to adverting to what we know about the words' semantic properties?⁴⁰ Absent some alternative development, this provides some reason to favor the second simulation idea to which we now turn.

6.2 Specifically Semantic Simulation

Even if the intention-hypothesis-generator black-box can be undarkened or avoided, there remains reason to seek an alternative development of the simulation strategy. The first suggestion presupposed a simulation-theoretic

³⁹ Sometimes we can reasonably hypothesize what someone intends to communicate without reference to the meaning of the sounds the person produces. But these are exceptional cases. Cf. Sperber and Wilson, 2002a, p. 11.

⁴⁰ Perhaps a dummy hypothesis is generated in order to start a process of mutual adjustment? How would this go? Abductive inference is a difficult topic generally, however, not just in this case.

account of our capacity to explain others' actions generally. But the simulation-theoretic account of action explanation is controversial. Some maintain that we explain others' actions on the basis of a cognized 'theory of mind': rather than sub-personally simulating the practical reasoning that led to action, we reason from general beliefs about why intentional agents act as they do.⁴¹ It's thus important to see that responding to our comprehension problem by invoking simulation need not commit one to a position concerning action explanation generally. When sub-personal simulation is characterized broadly as running a component 'off-line' (and in particular is not restricted to our capacity to explain and predict action), space opens up for a wide range of simulation hypotheses, with the possibility that some will and some won't hold up empirically.⁴²

So, how else might one develop the simulation strategy—in particular, without assuming (though not necessarily precluding) a simulation account of action explanation generally? The second suggestion hypothesizes that the information contained in a cognized truth-theory is utilized in two distinct ways by the processes leading to language production and by those involved in language comprehension, respectively. Comprehension thus involves simulation because the semantic information used in production is used in comprehension in a functionally modified way. The rough idea—crudely simplifying again—is this. In production, an appropriate representation is generated of what the speaker intends to communicate by linguistic means. For the representation to be appropriate, it must not only express the relevant intention; it must also be properly formatted to enable computational processing, drawing upon the information contained in the truth-theory, to extract an LF. The idea is that the processing takes this representation and, drawing upon the truth-theory, searches for an LF whose truth-conditions it expresses. The processing then continues by generating the appropriate phonological form to input to the articulatory systems.

In comprehension, on the other hand, the heard sounds are assigned an LF-representation that serves as input: processing now draws upon the information contained in the cognized truth-theory to generate an assignment of truth-conditions—that is, it finds a truth-theorem for this LF, and then detaches its right-hand-side for whatever further processing awaits it. The processing that's used in comprehension is distinct from that used in production: the former starts

⁴¹ One prominent contention is that 'theory theory' can, and simulation approaches cannot, explain the propensity for certain false action predictions. According to theory theorists, the false predictions result from false beliefs about how people act: action prediction is thus 'cognitively penetrable.' But, if a simulation approach were correct, these false beliefs would not impact the process by which we arrive at action predictions. An appeal to these false beliefs thus can't explain the false predictions—nor (it's claimed) do simulation theorists have an empirically plausible alternative explanation to offer. See Stich and Nichols, 1992, section 5, for the objection (developed further in other papers), and Heal, 1996, for a response. Stich and Nichols, 1997, concedes some ground.

⁴² Cf. Stich and Nichols, 1997, pp. 307–9.

with the left-hand-side of a truth-theorem and looks for its right-hand-side; the latter starts with a right-hand-side and looks for its left. In this way, this second simulation strategy differs from the first. The first simulation strategy supposes there's only one way the information contained in cognized truth-theories is utilized. According to it, in both production and comprehension we move from left-hand-side to right-hand-side.⁴³

The second strategy is like the first, however, in supposing that post-truth-theoretic processing is distinct in production and comprehension. Comprehension leads not (directly) to speech, but rather to further cognition, including to the formation of beliefs concerning what the speaker said or communicated. If the current hypothesis is to help solve the comprehension problem, it's thus required that this further processing include, once again, the 'transformation' of automatic context-sensitive terms: my coming to believe that he said that he (himself) is tall, on the basis of his uttering 'I am tall', can't consist in my placing in my belief-box a representation to the effect that *I* am tall. Any simulation strategy must account for how one 'pops out' of the simulation—or otherwise distinguishes it from the real thing. This is especially crucial in an application involving representations containing automatic context-sensitive terms.

It might be objected that the proposal isn't a simulation strategy at all, even on the broad understanding of simulation in play. Simulation, on this broad understanding, involves running a component off-line—that is, in an appropriately functionally modified way. But it would seem that the function of cognized truth-theories is as much to contribute to comprehension as to contribute to production.⁴⁴ Each activity could thus be claimed to involve a functional modification of the other. Some sort of asymmetry must be introduced in order to deem one a *simulation*—in particular, we need reason to think of the cognized truth-theories' role in *production* as in some sense more fundamental than its role in comprehension. Speaking of it as the *normal* function adds nothing; nor is it clear that it should matter what the *typical* or *usual* function is, if what's meant is just a numerical tally (it's not clear production would win in any case);⁴⁵ nor, finally, is it clear that either evolutionary or developmental primacy should matter. Still, it can seem there's a clear enough sense in which, according to this proposal, in comprehension I (or the semantics module) sub-personally make as if I'm speaking,

⁴³ Truth-theoretic pedagogical practice, incidentally, tends to downplay the task of production. In learning one's way around t-theories, one's characteristically asked to generate t-sentences for (and thus truth-conditions for) sentences or their LF's, but not to generate t-sentences for (and thus to find a corresponding object-level LF for) some sentence in the meta-language. (See, for example, the exercises and sample derivations in Larson and Segal, 1995.) The *first* simulation strategy goes so far as to hypothesize that our sub-personal use of t-theories *always* proceeds in the direction opposite of our student exercises.

⁴⁴ Similarly, the proponent of a simulation account of action explanation presumably must allow that contributing to the explanation of others' actions is part of the function of the components involved in practical deliberation.

⁴⁵ Nichols *et al.*, 1996, speak of a component's normal or usual function.

while in production I don't make as if I'm comprehending: if nothing else, in production, but not necessarily in comprehension, the first-person pronoun is assigned by the truth-theory the reference it does in fact have on that occasion.⁴⁶ Whether or not this is so, however, the objection is beside the point: even if one couldn't validate the 'simulation' label in this case (however useful thinking in terms of simulation might have been heuristically), what matters is that we have a proposal for responding to the comprehension problem. It might be fruitful to think through how best to characterize simulation strategies generally, but it's no objection to this particular proposal if in the end it's not so classified.⁴⁷

There are then various strategies available for addressing the comprehension problem for certain uses of automatic context-sensitive terms. They all require much development if they are to be serious proposals. But even as schematic suggestions they exhibit significant differences that might empirically distinguish them. The 'complementary context-sensitivity' strategy and the 'specifically semantic simulation' strategy agree that production and comprehension differ, as far as cognized truth-theories are concerned, in terms of what's input and what's output and thus in the direction of biconditional t-theorems relevant for processing purposes. But the 'specifically semantic simulation' strategy rejects the 'complementary context-sensitivity' strategy's hypothesis that production and comprehension differ further in that, for a *specific* class of terms, different *clauses* of cognized truth-theoretic axioms are relevant for their production and comprehension respectively. According to the 'specifically semantic simulation' strategy, there is no further difference in production and comprehension so far as cognized truth-theories are concerned. However, it must posit—and develop an account of—a non-truth-theoretic difference: the presence of terms generating the comprehension problem must trigger an independent process of 'transformation'. Some such

⁴⁶ I say 'not necessarily' to allow (mis)comprehending one's own past speech.

The particular reason given above to consider cognized truth-theories' role in production to be primary doesn't obviously carry over to other examples of simulation theories in the broad sense that Nichols *et al.*, 1996, offer. For example, the visual system's role in mental imagery would have to be secondary to its more fundamental function for some other reason. Closer to home is Halle and Stevens' account (1962) of phoneme recognition, which Nichols *et al.*, 1996, pp. 44–5, suggest might be 'the earliest detailed "off-line simulation" account in cognitive science':

On their account, our capacity for phoneme recognition exploits our phoneme production system. Crudely put, the idea is that when given phonemic input (i.e., when spoken to), the phoneme production system is taken off line (detached from its normal function) to generate hypotheses for matching the phonemic input. Recognition occurs when a match is found between the phonemic input and one of the hypothetical outputs of the phoneme production system.

It would need to be clarified in what sense phoneme *production* is the system's normal (or primary, or fundamental) function.

⁴⁷ It's worth noting that Stich and Nichols, 1997, p. 299, now argue that unqualified talk of simulation is theoretically useless and that the term is perhaps best retired.

account is required as well by the ‘simulation of action explanation generally’ strategy, which also rejects the positing of bifurcated axioms. But this strategy differs from other two in denying that production and comprehension differ in the directionality of the biconditional each makes use of. Where production and comprehension differ on this strategy is rather in the larger set of sub-personal processes invoked in each. In particular, it must endorse the hypothesis that a process of simulation operative in action explanation generally accesses truth-theoretic information to explain in particular others’ linguistic behavior.

7. Consequences, Non-Consequences, and Further Issues

I turn now to some further consequences of dropping *Meta-Insensitivity*—and some non-consequences as well. Examining the former helps clarify how dropping *Meta-Insensitivity* requires a reconception of the place of truth-theories in an account of linguistic competence.⁴⁸ I discuss the *non*-consequences, not only to add further clarification, but also to respond to objections I’ve frequently encountered.

7.1 Semantic Competence and the Cognition of a Single Truth-Theory

First, dropping *Meta-Insensitivity* is inconsistent with holding that semantic competence consists in cognizing a *single* truth-theory. For if we employ context-sensitive expressions in the meta-language, then what they express—and so what one cognizes—can *itself* vary across conversational contexts. The whole idea of the strategy, after all, is to exploit meta-level context-sensitivity that’s coordinated with object-level context-sensitivity. But then the context-sensitive t-sentences, for example, express *different* theorems in different contexts.

There are other reasons for denying that semantic competence consists in cognizing a single truth-theory, however; so it might be questioned whether this consequence requires any adjustment from what a truth-theorist should already accept. One might hold, for instance, that cognition of a truth-theory explains but one *aspect* of semantic competence. According to some, semantic competence in addition involves cognizing various bits of further information contained in the lexical entries, but not in the truth-theoretic lexical axioms, of various terms. Some of the information Higginbotham (1989) terms an ‘elucidation of meaning’—for example, that knives are for cutting—need not appear in the lexical axiom for the elucidated term (‘knife’), but might rather appear elsewhere in some sort of lexical entry. But this concedes just that there may be more to semantic competence than what’s captured truth-theoretically: it doesn’t preclude that the truth-theoretic (and thus the combinatorial) aspect of semantic competence is captured by a single truth-theory. In dropping *Meta-Insensitivity*, one must abandon this further claim as well.

⁴⁸ Some of these consequences, of course, may seem unattractive to some truth-theorists. They would thus provide further motivation for rebuffing the kinds of criticism with which we began.

Closer to this consequence, however, is Davidson's (1986) response to malaprops. Our ability to interpret them suggests to Davidson that only knowledge of a *series* of truth-theories, properly deployed, would suffice for knowledge of a human natural language.⁴⁹ Roughly, the suggestion is that semantic competence involves the ability to adjust one's cognized truth-theory to the idiosyncratic needs of current conversation—for example, by allowing 'derangement' to be assigned the semantic value one normally assigns to 'arrangement' in order to interpret someone who confuses the terms. Thus, again, it might seem that there's already independent reason for denying that semantic competence consists in the cognition of a single truth-theory.

Accepting Davidson's point can seem a small concession, however. Proponents of truth-theoretic conceptions of semantic competence should of course allow that cognized truth-theories change over time (for example, when one learns a new word). Adjustments for idiolectal differences—even one-off malaprops—thus might not phase someone who holds that semantic competence involves cognition of a single *more-or-less* stable truth-theory. Meta-linguistic context-sensitivity, on the other hand, in addition to offering a further challenge distinct from that posed by malaprops, is less easily shrugged off as a fringe phenomenon. Moreover, if Travis and others are right about the complex context-sensitivity of most terms, and if deploying coordinated context-sensitivity in characterizing cognized truth-theories is the proper truth-theoretic accommodation of them, then meta-linguistic context-sensitivity introduces a potentially much more pervasive source of change in cognized truth-theory.

7.2 Invariant Semantic Knowledge vs. Knowledge Particular to a Context

One attraction of *Meta-Insensitivity* and parameterization strategies is the natural division they assume and enshrine between, on the one hand, the more-or-less context-invariant standing knowledge of language that speakers carry from one conversational context to the next and, on the other, the ever-changing features of those contexts that speakers must track in addition in order to determine the truth-conditions associated with context-sensitive sentences on a particular occasion of use. If, however, one must use context-sensitivity in the meta-language in order to adequately characterize cognized truth-theories, then one cannot distinguish at least *this* way what speakers bring to bear *on* a context and what they know *of* a context. A second consequence of dropping *Meta-Insensitivity*—basically a corollary of the first—is thus that what speakers cognize in cognizing truth-theories is not the *invariant* contribution that expressions contribute to meaning in virtue of their semantic properties.

This is not to say that there is nothing invariant in a context-sensitive characterization of object-level context-sensitivity. After all, the simple examples we have mentioned have all deployed unchanging truth-theoretic *sentences* to characterize a

⁴⁹ Davidson never goes so far as to suggest that speakers in fact cognize truth-theories or that such cognition in fact explains or constitutes their semantic competence. His claim is merely that knowledge of a truth-theory (or, a series of them) *would suffice* for such competence. Having noted this, I'll now treat Davidson's discussion as if it takes place within a *neo-Davidsonian* cognitivist setting.

speaker's changing cognitions. The situation is analogous to one in which someone believes, day after day, what is on each day expressed by the sentence 'Today is the first day of the rest of my life'. The sentence used to characterize what she believes persists, but the belief thereby characterized changes. Similarly, although there may be stability in the language used to *express* what the speaker cognizes ('[I am big] is true iff I am big'), there is not stability in *what* the speaker cognizes.

It may be that in addition the speaker unchangingly cognizes *that* the truth-theoretic sentences are true (or, express truths). But what is thereby cognized is not itself the basis of the truth-theoretic aspect of semantic competence. For just as one can know that a sentence expresses a truth without knowing what that sentence expresses (perhaps you simply have it on good authority that the sentence is true), always cognizing that a t-sentence expresses a truth is not yet always to cognize the truth-theorems it expresses.⁵⁰ Adapting Perry, 1980, we may say that the speaker *accepts* the truth-theoretic sentences. It might be thought that *this* provides a single, more-or-less stable *attitude* speakers have, so that in this sense the account of semantic competence can after all be said to yield cognition of something invariant. But this depends on how talk of acceptance is itself understood. If acceptance of a sentence just is belief that the sentence (always) expresses a truth, then what was just said above applies. If accepting a sentence just is believing in each context what it expresses, then we have a multiplicity of attitudes after all. It's unclear whether talk of acceptance could be cashed out some other way—and also unclear whether it could be deemed something *sui generis*.

Closely related to this consequence is another. Higginbotham (1988, p. 31) complains that context-sensitive t-sentences 'will not reveal *how* truth-conditions vary systematically from context to context' (emphasis added—cf. Davidson, 1967, p. 34, and Burge, 1974, p. 212). It might thus seem a consequence of dropping *Meta-Insensitivity* that one must deny that revealing how truth-conditions vary systematically from context to context is something a cognized truth-theory must do—at least so far as providing an account of semantic (or, more broadly, linguistic) competence is concerned. This isn't quite right, since one way to violate *Meta-Insensitivity* would be to take a truth-theory that *does* yield t-sentences that context-insensitively characterize how truth-conditions vary systematically across contexts and then *add* context-sensitive axioms to it.⁵¹ The particular *motivations* for dropping *Meta-Insensitivity* with which we began, however, question precisely whether we cognize (or, indeed, whether there is) such a truth-theory. Thus dropping *Meta-Insensitivity* for *these* reasons would indeed have this consequence.

7.3 The Boundary between Semantics and Pragmatics

Another remark of Higginbotham helps make even more explicit the upshot dropping *Meta-Insensitivity* has for distinguishing semantics and pragmatics. He holds (1988, p. 29) that, 'The rules of semantics, since they relate form to

⁵⁰ Nor—switching scope—is cognizing that a t-sentence *always* expresses a truth yet to cognize the truth-theorems it expresses.

⁵¹ This is what Rumfitt, 1993, does.

meanings, are independent of context; . . . they come alive only in use. Disentanglement is therefore necessary . . .'. But if an adequate characterization of cognized truth-theories requires a context-sensitive meta-language, then there is not a context-independent truth-theoretic assignment of semantic value. It would thus follow either (a) that Higginbotham is wrong that the rules of semantics (at least so far as an account of semantic competence is concerned) are independent of context, or (b) that in so ascribing truth-theories one goes beyond an account of semantic competence, crossing into the realm of pragmatics. Either way, dropping *Meta-Insensitivity* has upshot for if and how one draws the semantics-pragmatics boundary. If (a), then semantics (so far as an account of semantic competence is concerned) does not relate forms to meanings (or valuation conditions) in a context-independent way. Insofar as drawing a semantics/pragmatics border remained important, it would have to be drawn some other way. If (b), then this paper's central suggestion—that, even in the face of criticisms such as those sketched above, one might preserve a place for cognized truth-theories by dropping *Meta-Insensitivity*—should not be formulated as a defense of truth-theoretic accounts of *semantic* competence. For, according to (b), once one drops *Meta-Insensitivity*, one is no longer offering an account specifically of *semantic* competence. Using the term 'semantics' in this way, however, does not preclude the suggestion that by dropping *Meta-Insensitivity* one can retain a place in an account of *linguistic* competence for cognition of truth-theories (and the explanatory power such ascription is claimed to have).

7.4 Unexplained Abilities

As noted, if one drops *Meta-Insensitivity*, then a speaker's semantic competence doesn't consist in—or, involve—cognizing *one* truth-theory. Rather, it involves the cognition of various truth-theories, according to the speaker's conversational context: a speaker's semantic competence thus involves the ability to deploy the appropriate truth-theory in a given context.⁵² It can seem to follow that the hypothesis that speakers cognize truth-theories doesn't yield a net reduction in unexplained cognitive abilities. If semantic competence involved the cognition of just *one* more-or-less stable truth-theory, then the proponent of a truth-theoretic account of semantic competence could claim to have cashed out talk of semantic competence—talk of a certain ability or know-how—in terms of a body of information possessed, in terms of cognition-*that*.⁵³

⁵² In light of the previous point, I should note that some might prefer I say, not that *semantic* competence would involve this ability, but rather that this aspect of linguistic competence would—likewise for elsewhere in this paper.

⁵³ One's possession of this information might itself rest on further abilities—for example, the ability to use the concepts expressed in the truth-theory. But these are abilities requiring explanation anyway independently of an account of semantic competence. (I am assuming here (against Dummett, 1975) that an account of semantic competence can be 'modest'—that is, that its job is just to explain our grasp of how words express concepts, not also to explain our grasp of the concepts our words express.)

If, however, semantic competence doesn't consist in the cognition of a *single* truth-theory, but rather in the ability to deploy the appropriate truth-theory in a given context, then we have replaced one kind of ability talk with another, which now must itself be explained. Not that this would constitute an *objection* to dropping *Meta-Insensitivity*: rather, it marks a debt that would have to be satisfied.

In fact, the net unexplained cognitive abilities remains unchanged—though there is a shift in how the questions are posed. Cognition of a context-*insensitively* characterized truth-theory might explain (the relevant aspects of) *semantic* competence, but there would remain the question of how language-users track the value of contextual variables such that they can de-conditionalize conditionalized t-sentences in order to arrive at truth-conditions for particular utterances. How does a language-user track such contextual features as who is being demonstrated, what is the relevant part of the subject of predication, who is being addressed, etc.? When *Meta-Insensitivity* is dropped, such questions (supposing the relevant object-level terms receive a context-sensitive truth-theoretic treatment) are transposed into questions concerning how the language-user deploys the *appropriate* truth-theory in a given context: how does 'that' in the meta-language shift its reference so as to refer to appropriate object on a given occasion, how does the semantic value of 'green' shift to align with the particular usage in play, etc.? Allowing context-sensitivity in the meta-language thus does indeed introduce new explanatory questions—but they replace parallel extra-semantic questions facing the proponent of a parameterization strategy.⁵⁴

⁵⁴ These questions (what fixes the semantic values the context-sensitive terms of the *meta*-language have on a given occasion—in particular, how are they fixed so as to yield the *appropriate* truth-theory?) take a more specific form if one accepts the further thesis that language-users cognize truth-theories by representing them in a Language of Thought (LOT)—which in turn, if one also rejects *Meta-Insensitivity*, requires allowing that LOT representations can be context-sensitive. One must, however, be careful to distinguish questions of the form 'in virtue of what does such-and-such a mental representation in LOT have the character it does (i.e., in virtue of what, does its content alter with context in the way that it does)?' from questions of the form 'in virtue of what, does it have the content that it does on a particular occasion?'. Answers to the former provide part of the answers to the latter. (Note that both kinds of question concern foundational semantics, not descriptive semantics. That is, they ask in virtue of what representations have certain semantic properties, as opposed to asking what semantic properties they have. Cf., e.g., Stalnaker, 1997.) Automatic context-sensitive LOT representations—such as NOW, for instance (bracketing the non-automatic aspect of its extent)—would seem to have their character in virtue of their conceptual role and their content in virtue of how the world is independent of the agent's relevant intentions. The story would obviously be more complex for non-automatic terms for which intentions and mechanisms of deference can come into play. It's worth noting, however, that there's room to consider whether some of the mechanisms of *ad hoc* concept construction posited by relevance theorists to handle some of the problematic cases to which we've alluded can be co-opted for an account of some of the sub-personal mechanisms involved in fixing the content of a context-sensitive LOT representation on a particular occasion. Cf. Carston, 2002, Chapter 5.

7.5 The Context of the Theorist

Denying *Meta-Insensitivity*, we've seen, potentially raises problems when speaker and comprehender are in different contexts—that is, when the same term as used by the one would have a different semantic value when used by the other. Similar issues arise for the theorist *ascribing* cognition of context-sensitively characterized truth-theories.⁵⁵ If only context-*insensitive* terms were used in the meta-language, the respective contexts of the theorist and the ascriber would be irrelevant. They become relevant, however, once *Meta-Insensitivity* is dropped. The theorist, for example, doesn't want to utter: S cognizes that [I am tall] is true iff I am tall. With respect to automatic context-sensitive terms, such as 'I', the theorist ascribing cognition must have available context-sensitive expressions appropriately complementary to those that would be used from the speaker's perspective, as in: S cognizes that [I am tall] is true iff she herself is tall. With non-automatic context-sensitive terms such as 'tall' (or at least [is tall]), an accurate ascription would seem to require that the theorist's intentions be such that she use the meta-linguistic term in the *same way* the speaker would (tall for a seven year old, tall for a basketball player, or what have you). But since speakers may use 'tall' differently on different occasions, this means that there's no *one* context from which a theorist thus adverting to context-sensitivity in the meta-language can ascribe the various truth-theories users of this term cognize. The theorist can, however, without worrying about aligning contexts, indicate for any given speaker what truth-theoretic sentences would—as used from the speaker's perspective—express what the speaker cognizes. This is what I have been doing throughout this paper: sample context-sensitive truth-theoretic sentences have been displayed, not used in an ascription.

7.6 Knowing a Semantic Theory without Knowing the Language

Once *Meta-Insensitivity* is dropped, there is the threat (or so it has been claimed to me) that truth-theorists can be hoist by the petard of an argument they often deploy against alternative approaches to semantics—viz., that one could know these alternative semantic theories and yet still not know the language they are theories of. Truth-theorists have used this argument against, for example, structural semantics, model-theoretic semantics, and Horwich's minimalist semantics.⁵⁶ To take a simple translational example, knowing a theory that yields such information as that 'blau' means in German what 'azul' means in Spanish needn't enable one to know German (for example, if one doesn't know Spanish). It is claimed that more

⁵⁵ See fn. 21 above.

⁵⁶ See, e.g., Lepore, 1983, and Higginbotham, 1988 and 1999. (See Horwich, 1999, for his defense of minimalism in semantics. His use of the term 'minimalism' is unrelated either to Chomsky's use or to the various uses of the term in debates concerning the semantics/pragmatics border.)

sophisticated approaches fall prey to essentially the same point.⁵⁷ The threat is that, once *Meta-Insensitivity* is dropped, the argument can be turned back upon cognitivist truth-theoretic accounts of semantic competence. For, as we've seen, to deploy context-sensitive expressions in characterizing the cognized truth-theories is to deny that semantic competence consists in cognizing (much less knowing) a *single* truth-theory. Thus, every truth-theory (that which is expressed in some context by some truth-theoretic sentences) is such that knowing just it leaves one well short of knowing a natural language.

Truth-theorists, however, can avoid being thus hoisted by their own petards when they drop *Meta-Insensitivity*. For the objection they raise to alternative semantic theories is that knowledge of them yields no semantic competence *at all*. The result of dropping *Meta-Insensitivity* is rather that knowledge of a truth-theory (at least of the sort we cognize) can at best yield *partial* semantic competence—competence with respect to a certain context. For to drop *Meta-Insensitivity* is to concede that, using one set of truth-theoretic sentences, one can't from the standpoint of one context characterize, for any context-sensitive expression and any context, the conditions for the expression having a certain semantic value in that context—or at least this is not how *our* semantic modules tacitly do it. To overcome this partiality, what one must have is the ability to cognize the appropriate truth-theory in a given context, which is just what cognizing context-sensitively characterized truth-theories is supposed to give us.

7.7 The Replacement Thesis

Consider the *Replacement* thesis: for any possible utterance of a context-sensitive sentence, there's a context-*insensitive* sentence such that, had *it* been uttered, it would have said the same thing. What the thesis itself says depends on the relevant notion of sameness of content. But suppose some clarification has been accepted. It might then be thought that a consequence of dropping *Meta-Insensitivity* is that one must also deny *Replacement*. For, if *Replacement* were true, it would seem one could in principle characterize a cognized truth-theory *without* using context-sensitive terms. Some consider *Replacement* obviously true or at least find its denial unattractive, and have therefore raised this consequence as an objection to dropping *Meta-Insensitivity*. I won't here comment on the viability of *Replacement*, but I do want to argue that it's unobvious whether it's indeed a consequence of dropping *Meta-Insensitivity*. Whether it is might depend on one's *reasons* for dropping the latter.

The objector's claim needn't be that *Replacement* guarantees the success of the parameterization strategy (or in some other way guarantees that semantic

⁵⁷ What force this argument should have depends on what a semantic theory ought to do. Some semanticists hold that the job of semantic theory is simply to state what information sentences encode (relative to a context). A semantic theory, on this view, needn't also play a role in an account of semantic *competence*—so, it's just irrelevant that knowing the semantic theory needn't suffice for knowing the language. Cf. Soames, 1989.

competence involves the cognition of *one* context-insensitively characterized truth-theory). Consider, for instance, the context-sensitive t-sentence '[You're here] is true iff you're here'. *Replacement* doesn't guarantee a context-insensitive statement of truth-conditions that obtains for *any* use of 'You're here'. What it guarantees is that each use of the t-sentence has a context-insensitive replacement (just as each of use of 'You're here' would). But different uses of the t-sentence—addressed, for example, to different people—would have *different* replacements. The *Replacement* thesis thus might allow us only to move from context-sensitively characterized truth-theories to a *collection* of context-insensitively characterized truth-theories—one for each context of utterance.

Still, *Replacement* at least guarantees that one could characterize insensitively each of the truth-theories speakers cognize from one context to the next. But then, claims the objector, *Meta-Insensitivity* is upheld after all. There remains a substantial departure from standard treatments of object-level context-sensitivity: a single parameterized truth-theory will have been replaced by a *collection* of truth-theories. But, nonetheless, meta-linguistic context-sensitivity is *not required* to characterize them, if the objector is right. Thus it would seem that the denial of *Meta-Insensitivity* requires the denial of *Replacement*.

This conclusion, however, can be blocked. My formulation of *Meta-Insensitivity* refers to how one might *adequately* characterize cognized truth-theories. This is because truth-theoretic accounts of semantic *competence* do more than just supply truth-theories: they ascribe *cognition* of such theories, which arguably constrains not only the relevant class of truth-theories, but also what counts as an adequate characterization of one. In particular, there is arguably a constraint on accurate ascription that context-insensitive replacements would fail to satisfy. Loosely put, so as to remain neutral among the many competing accounts of attitude ascription, an accurate ascription should capture, not only the content of an ascriber's attitude, but also relevant aspects of how she herself represents that content. There is thus room to argue that, even if context-insensitive replacements preserve content expressed context-sensitively, an ascription in context-sensitive terms might capture aspects of the ascriber's attitude that would be omitted by context-insensitive replacements.⁵⁸

Replacement may yet be false⁵⁹—and, if *Meta-Insensitivity* is also false, it may be that some of the considerations used against the latter can be deployed against the

⁵⁸ It might now seem however that, once *Meta-Insensitivity* is dropped, *ascriptions* of truth-theories will provide counter-examples to *Replacement*, since they must use context-sensitive terms in their that-clauses. It's difficult to think through this line of thought further without a clearer understanding of what standards for sameness of content are in play with *Replacement* and without agreement as to the semantics of attitude ascriptions. But let me note that on at least some views of ascriptions—in particular, those stemming from Davidson's (1968) paratactic analysis—it's possible to maintain that terms occurring in the that-clause of an attitude ascription are not *used* in the relevant sense. Cf. Gross, 2002a.

⁵⁹ *Replacement* is indeed denied by, e.g., Carston, 2002, Chapter I, who cites others as well. I remark on the matter a bit further myself in Gross, 1998/2001, pp. 17–9, and 2002b.

former as well. Perhaps complexly context-sensitive predicates might stymie attempts *both* to parameterize their satisfaction conditions *and* to supply context-insensitive replacements for each possible use of them. And surely those moved to deny *Meta-Insensitivity* by indexicality criticisms will use them as well against *Replacement*, however the latter's notion of 'same content' is cashed out. But these links will be neither here nor there to those moved to deny *Meta-Insensitivity* rather on empirical grounds concerning the nature of the relevant cognitive processing.

Meta-Insensitivity and *Replacement* are two distinct claims: the line from denying *Meta-Insensitivity* to denying *Replacement*, if there is one, might depend on one's *reasons* for denying *Meta-Insensitivity*, and even then might not be direct or obvious. The denial of *Replacement* is thus at best a *potential* consequence that would have to be made out. If it *is* a consequence, whether it is dire is a further question.

8. Conclusion

I have suggested a move proponents of truth-theoretic accounts of semantic competence might make in the face of certain criticisms that have been raised against them—viz., that they abandon the claim that cognized truth-theories be characterized using only context-*insensitive* expressions in the meta-language. And I have discussed some of its problems and consequences. I have also noted along the way, however, that the move is only as interesting as the empirical grounds (direct or indirect) that might be mustered for an account that has it as a part. Besides a reference to the explanatory power accorded compositionality, there has been no attempt to supply such empirical grounds here. But it is legitimate, especially in an unsettled and still somewhat underdeveloped area, to consider in advance what the possibilities might be—and thus for what one might need to *seek* evidence. And, if I have succeeded rather only in displaying why the move is unattractive, that's progress of a sort too.

*Department of Philosophy
University of Pennsylvania*

References

- Atlas, J. 1989: *Philosophy without Ambiguity: A Logico-Linguistic Essay*. Oxford University Press.
- Bach, K. 2000: Quantification, qualification, and context: a reply to Stanley and Szabó. *Mind & Language*, 15, 262–83.
- Brenehy, R. 2002: The current state of (radical) pragmatics in the cognitive sciences. *Mind & Language*, 17, 169–87.

- Burge, T. 1974: Demonstrative constructions, reference and truth. *Journal of Philosophy*, 71, 205–23.
- Carston, R. 1988: Implicature, explicature, and truth-theoretic semantics. In R. Kempson (ed.), *Mental Representations*. Cambridge University Press.
- Carston, R. 2000: Explicature and semantics. *UCL Working Papers in Linguistics*, 12, 1–46.
- Carston, R. 2002: *Thoughts and Utterances*. Oxford: Blackwell.
- Castañeda, H. 1966: 'He': a study in the logic of self-consciousness. *Ratio*, 8, 130–57.
- Currie, G. 1995: Visual imagery as the simulation of vision. *Mind & Language*, 10, 25–44.
- Davidson, D. 1967: Truth and meaning. *Synthese*, 17, 304–23. Cited as reprinted in (Davidson, 1984).
- Davidson, D. 1968: On saying that. *Synthese*, 19, 130–46.
- Davidson, D. 1976: Reply to Foster. In G. Evans and J. McDowell (eds), *Truth and Meaning*. Oxford University Press. Cited as reprinted in (Davidson, 1984).
- Davidson, D. 1984: *Inquiries in Truth and Interpretation*. Oxford University Press.
- Davidson, D. 1986: A nice derangement of epitaphs. In E. Lepore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Oxford: Blackwell.
- Davies, M. and Stone, T. (eds) 1996: *Mental Simulation: Evaluations and Applications*. Oxford: Blackwell.
- Dummett, M. 1975: What is a theory of meaning? (I). In S. Guttenplan (ed.), *Mind and Language*. Oxford University Press.
- Ezcurdia, M. 2001: Thinking about myself. In A. Brook and R. DeVidi (eds), *Self-Reference and Self-Awareness*. Amsterdam: John Benjamins.
- Gross, S. 1998/2001: *Essays on Linguistic Context-Sensitivity and its Philosophical Significance*. Ph.D. Thesis, Department of Philosophy, Harvard University. Cited as published in *Studies in Philosophy: Outstanding Dissertations*. London: Routledge.
- Gross, S. 2002a: Vagueness, indirect speech, and the world. *Protosociology*, 17, 149–64.
- Gross, S. 2002b: Is context-sensitivity eliminable? Some remarks. *The Dialogue*, 2, 21–38.
- Gross, S. 2004: Putnam, context, and ontology. *Canadian Journal of Philosophy*, 34, 507–54.
- Gross, S. Forthcoming: Can empirical theories of semantic competence really help limn the structure of reality? *Noûs*.
- Halle, M. and Stevens, K. 1962: Speech recognition: a model and a program for research. In J. Fodor and J. Katz (eds), *The Structure of Language*. Englewood Cliffs, NJ: Prentice-Hall.
- Heal, J. 1996: Simulation and cognitive penetrability. *Mind & Language*, 11, 44–67.
- Higginbotham, J. 1988: Contexts, models, and meanings: a note on the data of semantics. In R. Kempson (ed.), *Mental Representations*. Cambridge University Press.
- Higginbotham, J. 1989: Elucidations of meaning. *Linguistics and Philosophy*, 12, 465–518.
- Higginbotham, J. 1999: A perspective on truth and meaning. In L. Hahn (ed.), *The Philosophy of Donald Davidson*. Chicago, IL: Open Court.
- Horwich, P. 1999: *Meaning*. Oxford University Press.

- Larson, R. and Segal, G. 1995: *Knowledge of Meaning*. Cambridge, MA: MIT.
- Lepore, E. 1983: What model-theoretic semantics cannot do. *Synthese*, 54, 167–87.
- Lepore, E. and Loewer, B. 1989: You can say *that* again. *Midwest Studies in Philosophy*, 14, 338–356.
- Ludlow, P. 1999: *Semantics, Tense, and Time*. Cambridge, MA: MIT.
- Ludlow, P. Forthcoming: A note on alleged cases of non-sentential speech. In R. Elugardo and R. Stainton (eds), *Non-Sentential Speech*.
- Millikan, R. 1990: The Myth of the Essential Indexical. *Noûs*, 24, 723–34. Revised and expanded in R. Millikan, *White Queen Psychology and Other Essays*. Cambridge, MA: MIT, 1993.
- Nichols, S., Stich, S., Leslie, A. and Klein, D. 1996: Varieties of off-line simulation. In P. Carruthers and P. Smith (eds.), *Theories of Theories of Mind*. Cambridge University Press.
- Nolan, R. 1970: *Foundations for an Adequate Criterion of Paraphrase*. The Hague: Mouton.
- Perry, J. 1979: The problem of the essential indexical. *Noûs*, 13, 3–21.
- Perry, J. 1980: Belief and acceptance. *Midwest Studies in Philosophy*, V, 533–542.
- Pustejovsky, J. 1995: *The Generative Lexicon*. Cambridge, MA: MIT.
- Rumfitt, I. 1993: Context and content: the paratactic theory revisited and revised. *Mind*, 102, 429–454.
- Searle, J. 1983: *Intentionality*. Cambridge University Press.
- Soames, S. 1989: Semantics and semantic competence. *Philosophical Perspectives*, 3, 575–596.
- Sperber, D. and Wilson, D. 1986/95: *Relevance: Communication and Cognition*. Oxford: Blackwell. 2nd edition with postface, 1995.
- Sperber, D. and Wilson, D. 2002: Truthfulness and relevance. *Mind*, 111, 583–632.
- Sperber, D. and Wilson, D. 2002a: Pragmatics, modularity, and mind-reading. *Mind & Language*, 17, 3–23.
- Stainton, R. 1994: Using non-sentences: an application of relevance theory. *Pragmatics and Cognition*, 2, 269–84.
- Stalnaker, R. 1997: Reference and necessity. In B. Hale and C. Wright (eds), *A Companion to the Philosophy of Language*. Oxford: Blackwell.
- Stanley, J. 2002: Making it articulated. *Mind & Language*, 17, 149–68.
- Stanley, J. and Szabó, Z. 2000: On quantifier domain restriction. *Mind & Language*, 15, 219–61.
- Stanley, J. and King, J. Forthcoming: Semantics, pragmatics, and the role of semantic content. Forthcoming in Z. Szabó (ed.), *Semantics vs. Pragmatics*.
- Stich, S. and Nichols, S. 1992: Folk psychology: simulation or tacit theory? *Mind & Language*, 7, 35–71.
- Stich, S. and Nichols, S. 1997: Cognitive penetrability, rationality and restricted simulation. *Mind & Language*, 12, 297–326.
- Szabó, Z. 2001: Adjectives in context. In I. Kenesei and R. Hamish (eds), *Perspectives on Semantics, Pragmatics, and Discourse: a Festschrift for Ferenc Kiefer*. Amsterdam: John Benjamins.
- Travis, C. 1985: On what is strictly speaking true. *Canadian Journal of Philosophy*, 15, 187–230.