

1

Semantics and Metasemantics in the Context of Generative Grammar

Seth Yalcin
University of California, Berkeley

1 Introduction

Metasemantics, I will take it, inquires into the nature of certain properties investigated by natural language semantics.¹ It seeks a certain sort of ‘more fundamental’ characterization of these properties. It asks whether and how these semantic properties might admit of some illuminating reduction to, or unification with, nonsemantic properties.

It would be good to put this more precisely. But beyond a superficial level, this kind of question can hardly be framed in abstraction from fairly specific assumptions about what sorts of things natural language semantics is in the business of explaining, and about what the best theory of those explananda looks like. One needs such assumptions to have some reasonable handle on what the ‘semantic properties’ are whose ground metasemantics asks about. Any clear statement of the problems of metasemantics is thus bound to be about as controversial as the picture of first-order semantic theory it assumes. And needless to say, there is not exactly consensus about just what picture of first-order semantic theory is the right one. Varieties of semantic theory abound; and even when theorists agree about the proper shape of a semantic theory, they may disagree about what such a theory is fundamentally directed at explaining.

The science of linguistic meaning is at an early stage. Metasemantic speculation is thus hazardous, and risks being idle. Be that as it may, in this chapter, I make a preliminary attempt at framing some questions of metasemantics. I aim in particular for a characterization of the problems of metasemantics that meshes with natural language semantics in the modern, broadly model-theoretic tradition as informed

¹ I needn’t quarrel with other ways of using ‘metasemantics’.

by generative grammar. I make essentially no effort to defend that tradition directly; I simply take a certain theoretical perspective from within that tradition, and ask how some core metasemantic questions might be stated from within it. Worse, I aim only to frame questions, not hypotheses. I do not assume that metasemantics constitutes a subject matter with anything like clear boundaries (anymore than ‘meta-chemistry’ or ‘meta-geoscience’); I assume just that whatever else it is, it seeks some more fundamental description of the facts semantic theory traffics in, in some way to be clarified.

Here are the main contours of the viewpoint that emerges. The questions at the core of metasemantic inquiry, I suggest, have less to do with issues about content than is often assumed. Rather, these questions are better framed in terms of the notion of semantic value, ‘semantic value’ being a term of art of semantic theory. They are not questions of supervenience or cause, but of the ground of certain properties, properties articulated—in the current state of understanding, anyway—largely using the notion of semantic value. Specifically, they concern the ground of an aspect of a certain mental state postulated in linguistic theorizing, one we could informally describe as the state of knowing a language. They are questions whose answers are heavily constrained by the assumption of compositionality. Finally, they are questions barely on the horizon of systematic inquiry, and ones probably best pursued at present simply by pursuing semantic theory.

2 Not about content

2.1 *Semantic value and content distinguished*

I begin with the distinction between semantic value and content. Metasemantic questions foremost concern the compositional semantic values of expressions. They are not, except perhaps indirectly, questions about content.

This may sound confusing. After all, with some exceptions, it has been usual in contemporary philosophy to work with a notion of content (propositional content, informational content, representational content, proposition) according to which contents perform an array of explanatory work, and that work just stipulatively includes playing the role of the compositional semantic values of declarative sentences in context. (What other explanatory work contents/propositions are understood to perform varies among theorists, but usually the work includes playing a certain psychologically important role—‘the objects of the attitudes’ role—on which more shortly.²) This tendency, however, is problematic. As emphasized in various ways by Dummett (1973, 1993), Lewis (1980), Stanley (1997a, 1997b), Yalcin (2007, 2012a), Ninan (2010b, 2012b),

² The list commonly also includes the demand that contents or propositions be ‘the fundamental bearers of truth values’ and ‘the objects of assertion’ (or the ‘assertoric content’ of assertions). As far as I can see, the former may be added to the content role compatible with what I say in the present chapter. I discuss the status of the object of assertion role briefly below.

and Rabern (2012a, 2012b), among others, there is little motivation for theorizing under the assumption that the compositional semantic value of a sentence relative to context is the sort of thing that also plays certain of the other key roles paradigmatically associated with the notion of content. We should take care to distinguish the explanatory roles corresponding to ‘content’ and ‘semantic value’, respectively; and we should not assume the roles have the same kind of realizers.

I say we should distinguish the explanatory roles corresponding to these two notions. I am taking it we do best to understand these notions functionally. By ‘content’, I have in mind *the realizers of the content role*; by ‘semantic value’, I have in mind *the realizers of the semantic value role*. To clarify what we are talking about, then, we could ask: What exactly are these roles—what are we, or should we be, trying to explain with these notions? What are semantic values and contents, respectively, supposed to *do* in the theories employing these notions? These are questions of terminology, but they are not merely terminological, since of course, good terminology should be relevantly isomorphic to the structure of the domain being theorized about. Detailed answers to these questions generally make for substantive debate.

We don’t need (and anyway, would not have the space to provide) detailed answers here. Instead I will sketch some vague and partial proto-answers, ones that could be fleshed out in myriad ways. What I want to do is to bring into view some familiar and recognizable directions for cashing the ‘semantic value’ and ‘content’ roles out, enough to see that the two notions are (on theoretically sensible construals) located in separable, if related, explanatory projects. The point is to say enough about them to see why we ought to break them off from each other. Naturally, even my vague partial characterizations of these notions will be contestable. But I don’t think I go far out on a limb. What little I have to say about these notions is not particularly original, and is in the spirit of much, though certainly not all, prevailing usage.

Our main preoccupation in this chapter is with the notion of semantic value, but consider first content. I will use ‘content’ (or ‘informational content’, ‘propositional content’, ‘representational content’, ‘proposition’, etc.) so that contents are, at least:

- (C1) those things, the having of which by mental states largely grounds or constitutes the intentionality of those states;
- (C2) those things, the having of which by mental states largely grounds or constitutes the causal efficacy of those states vis-à-vis action (or a certain class of actions), at a level of abstraction approximating folk psychological description and explanation.

So understood, the notion of content has its home in a theory which attempts to explain the representational properties of mental states, the production of behavior at a certain high level of abstraction, and the character and explanatory power of folk psychological explanations of rational action. Granting it makes sense to seek such

a theory in the first place—contestable, obviously—a not uncommon starting point postulates a single sort of realizer for the above work.³

This characterization of the content role is incomplete and rough, and a full account should extend and refine it. But it gives some initial handle on what one could be trying to explain with the notion—on what, vaguely, is supposed to discipline ‘content’ talk in theoretical inquiries employing this notion. I do not claim this reflects any consensus about how philosophers in fact use ‘content’—there is no consensus—only that it corresponds to one common and important thread.

My point in isolating this aspect of the content role is to suggest that if anything actually does this work, the thing that does this work should not also be expected to realize the semantic value role (or more specifically, *the semantic value of a sentence in context*-role). What is the semantic value role? In attempting to articulate it, we might start indirectly, with a sense of the sort of data that is commonly leveraged to constrain theories in natural language semantics, at least at present. For example:

- (D1) **PRODUCTIVITY FACTS.** Speakers of a given language can understand and produce complex expressions in that language that they have never before encountered.
- (D2) **ENTAILMENT FACTS.** Some sentences in a language entail others; some sentences are inconsistent with others. Competent speakers manifest knowledge of such facts.
- (D3) **COMMUNICATION FACTS.** Speakers of a common language can transfer an abundant range of information systematically using that language.
- (D4) **ACCEPTABILITY FACTS.** Some sentences or discourses in a language are judged to be unacceptable, or uninterpretable, or marked, by speakers of that language, while others are not.
- (D5) **TRUTH/APPROPRIATENESS FACTS.** Some sentences in a language are judged to be true or appropriate by speakers relative to actual or stipulated scenarios, while others are judged false or inappropriate relative to such scenarios.⁴

³ Compare the approaches to content taken in, e.g., Stalnaker (1984), Dretske (1988), Lewis (1994), Braddon-Mitchell and Jackson (2007: ch. 7). I would read these theorists as working with a notion of content directed heavily, if not exclusively, at these putative explananda.

Of course, a theory of content might begin with the presumption of a single kind of realizer, but evolve to recognize a family of varieties of contents, with each suited for more precisely delineated explanatory purposes (as, e.g., a compositional semantic theory might evolve to incorporate a secondary level of semantic values in order to accommodate focus or presupposition projection).

⁴ Often it is said that the explananda of semantics include the truth-conditions of sentences (Lewis 1970; and Heim and Kratzer 1998, among countless others). If one says this, however, one should take care not to beg the question against non-truth-conditional semantic systems. We do better to describe the kind of data such theorists have in mind as entailment and consistency data, or as truth judgment data. The hypothesis that declarative sentences have truth-conditions is, *inter alia*, a means of providing a specific explanation of this kind of data.

(Of course, the data points rarely fall squarely under just one of these headings; and this list is in no way meant to be exhaustive of the kinds of data that are or might be used to constrain theories in natural language semantics.) Semantic theory does not attempt to explain all of this data ‘by itself’; rather it is supposed to play a certain part in the full explanation, in tandem with other theories and assumptions.

Merely listing some data does not suffice to settle the topic of inquiry, of course. It is one thing to describe the sort of data a theory gets constrained by as a matter of actual practice at some point in its history, and another thing to say what sort of questions the theory is ultimately directed at answering. In this chapter, I will take it that natural language semantics is a certain chapter in a general theory about what it is to speak and understand a language, broadly in the spirit of Chomsky (1965, 1986). The notion of semantic value, as deployed in the context of such a theory, serves fundamentally to characterize aspects of the mental state of knowing a language, in a sense I will try to elaborate. Let me sketch how this notion is supposed to make contact with data of the sort just sketched, in the context of this conception of what semantics explains.

With an eye towards explaining the productivity facts, we make two assumptions. First, understanding a language (fragment) is understood (modeled) relationally: it is assumed to be, in part, a matter of knowing the semantic values of the expressions of the language. The state of knowing that the semantic value of e is m in language L is an aspect of the broader state of knowing, or understanding, or ‘cognizing’ L , a state which also underwrites knowledge of the syntax of L . This state is one we postulate in the context of linguistic theorizing. It is not the ‘knows’ of the epistemologist. It is not a propositional attitude, and it is not of a kind with folk-psychological intentional talk. (Compare Chomsky 1968a: 169. More on this below.) The second assumption is that the semantic values of complex expressions are generally determined compositionally. With these, we have the possibility of an abstract explanation of the productivity facts (D1), as knowledge of a finite stock of basic semantic values, together with knowledge of a generative syntax and a finite body of composition rules, could suffice to ground knowledge of the semantic values of the countless complex expressions of the relevant fragment.⁵

Second, semantic values are assumed to be the sorts of things consequence and consistency relations are articulated in terms of: when $\Gamma \models \phi$ holds, this is (at least partly) *because of* the semantic values of (the sentences in) Γ and of ϕ , respectively. Hypotheses about semantic values can thereby serve to predict, and ground, entailment and consistency facts, hence knowledge of such facts (D2).

Third, we assume that competent speakers can leverage their mutual knowledge of the semantic values of the expressions of their language to coordinate upon and

⁵ I don’t mean here to be presuming in advance a significant gap between syntax and semantics. On some conceptions, the compositional semantic properties of expressions just are, for the most part, certain of their syntactic properties (e.g., Chomsky 1986). This viewpoint is, I think, compatible with the conception of semantic theory I have in mind, though not entailed by it.

communicate information in context—at least to whatever rough extent speakers in fact accomplish this with language, a delicate question, and one difficult to broach in the absence of highly substantive further assumptions (about which more below). Thereby semantic values can serve to partly ground facts of information transfer (D3); and hence theorizing about semantic values may be constrained, to some extent, by the facts of linguistic communication.

Fourth and fifth, we assume that relevant judgments of acceptability (D4) and truth (D5) are substantially (not invariably) explained (causally) by appeal to the state of knowing the language—especially to the aspect of the state which consists in knowing the semantic values of the expressions of the language. Semantic values are hypothesized to ground certain markedness features of sentences—for instance, markedness properties owing to composition failure, or mismatched semantic features, or to entailment properties (# ‘Since most of us are here, no one is here’, etc.). They are hypothesized to ground certain judgments of ambiguity, both lexical and structural. And they are taken to ground, at least partly, certain evaluations of sentences as true (false) or (in)appropriate by subjects relative to actual or hypothetical scenarios.⁶

I take it this gives a basic initial handle on what the realizers of the semantic value role are supposed to do. And I suggest this yields a construal of the notion of semantic value plausible for much of what we actually find in semantic theory. Naturally we expect the explanatory role of this notion to evolve in the course of semantic theory, and the relative significance of various explananda on hypotheses about semantic values to shift as inquiry proceeds.

If we approach the explanatory roles associated with the notions of content and semantic value, respectively, roughly along the lines just sketched, it clearly does not follow that the realizers of the sentential semantic value role should be identified with whatever realizes the content role. Why then has this identification, or something like it, been a common background assumption in the philosophy of language? Why might such an identification seem attractive?

Much of the temptation apparently stems from the demand that semantic values play a significant role in serving to ground facts of the (D3)-variety—facts of communication. After all, one conceivable way that semantic values could in principle serve to help explain cases of systematic information transfer is if (i) linguistic communication were largely a matter of coordinating on items of content, and (ii) the compositional semantic values of sentences (evaluated relative to context) were just *identified* with the relevant item of informational content transferred. Assumptions along these lines have been common in the philosophy of language. If they were

⁶ Although speaker judgments form a nontrivial portion of the evidence constraining semantic theory, semantic theory is (as many have stressed) not a theory of speaker intuition. Again, it is part of a theory of human linguistic competence—a theory of what it is to speak and understand a language. For more on the role of intuition in linguistic theorizing, see Schütze (1996, 2010); Kepsler and Reis (2005); Sprouse and Almeida (2012); Schütze and Sprouse (2014), and references cited therein.

correct, then it would be clear enough how knowledge of semantic values could be leveraged to communicate informational content.

A second motivation for the identification flows from the demand that semantic values play a significant role in serving to ground facts of the (D5)-variety—in particular, judgments of truth and falsity. If one assumes that the realizers of the content role should be truth-evaluable entities, and that speaker judgments about truth and falsity tend to be tracking the truth or falsity of contents somehow determined by what is said, then one simple way the semantic values of sentences in context could help to explain the relevant class of truth value judgements is if those semantic values just were items of content.

Nevertheless, there are good reasons, both conceptual and empirical, to reject the idea that the semantic values of sentences in context and contents are the same kinds of things.

2.2 *Against identification: conceptual considerations*

Beginning on the conceptual side, a notable objection is due in essence to Lewis (1980). It is possible that, owing to the operators the language in question contains, the semantic value of a sentence relative to context must be some complicated intension, variable with respect to an array of parameters—say, parameters for world, time, location, standard of taste, orientation, standard of precision, state of information, etc. The details here will be a contingent matter concerning the particular architecture of the language in question. It has to do with what expressions (if any) are best semantically modeled as intensional operators. It may also be, owing to the presence of quantifiers or lambda abstractors in the language behaving semantically along standard lines, that the semantic value of a sentence (relative to context) must be something variable with respect to assignment functions (Rabern 2013). Such considerations bear directly on the question what the semantic value of a sentence is. But—here is the main point—they don't directly bear on the question what realizes the content role. Views about the nature of content are not constrained by such considerations. In advance of inquiry, there is no reason to think that the things that best realize the content role—the things that, *inter alia*, play the role in the explanation of intentionality and action described above—need to be variable in just the respects required to give a compositional semantics adequate for the idiosyncrasies of some particular natural language. Developing an adequate compositional semantics for intensional operators and variable binders—a subproject of explaining the sort of facts described in (D1)–(D5) above, *inter alia*—is just a different problem from the problem of developing a theory of content. We should expect answers to questions about what realizes the semantic value role to be peculiar to particular language systems. But in advance of theorizing, there is little reason to expect the realizers of the content role to exhibit this kind of sensitivity.

Some theorists will want to respond as follows: declarative sentences in context ought at least correspond in some relatively direct way to items of content—in

particular, to the content characteristically communicated by utterances of that sentence (in some sense of ‘characteristically communicated’). After all, as recently acknowledged, we expect semantic values to play some part in explaining facts of linguistic communication and information transfer. In response, let me make three points.

First, it may be questioned whether some notion of content will in fact appear, or loom centrally, in the best theoretical account of linguistic communication. Granting that the semantic properties of expressions play some significant role in successful ordinary linguistic communication, it does not follow that this role is largely that of picking out propositional contents. Such a picture has its applications, but only under severe idealizations. Our commonsense ways of describing linguistic communication seem to involve some notion of content (or so philosophical orthodoxy has it); but the best theoretical account of the phenomena will presumably move beyond ordinary talk and commonsense preconceptions, departure from ordinary talk being the natural expectation in scientific inquiry of any sort. The way that the best theory of linguistic communication moves beyond ordinary talk may well end up failing to vindicate much of that talk, and may fail to call for any of the technical notions of content developed in the philosophical literature. There seems to be little reason, anyway, for confidence that this is not the case in the current early state of inquiry.⁷

Second, even allowing that linguistic communication does centrally involve some degree of coordination on items of content at some interesting level of abstraction, it remains open that coordination on items of content is a highly approximate, more-or-less affair, with perfect coordination on content not being especially important, and rarely or never happening. If this is the reality of the situation, then we should not expect sentential semantic values in context to determine unique contents.

Third, even if we imagine there to be a relatively direct mapping from sentential semantic values evaluated relative to context to unique contents, there is no reason we must assume that the relevant mapping is the identity relation. It would suffice if sentential semantic values plus context can (often enough) systematically determine the relevant contents, in some way mutually known by competent speakers. That is, it would suffice if there were commonly known pragmatic (or ‘postsemantic’) bridge principles indicating how an item of informational content is to be recovered, given an input context and sentential semantic value.⁸ This would suffice to allow semantic values to perform their share of the work in explaining facts of information transfer.

I take this last point to basically echo Lewis (1980); instructive further discussion occurs in Ninan (2010b, 2012b) and Rabern (2012a, 2012b). Ninan and Rabern are

⁷ Compare Chomsky (1986, 1992, 1995, 2000: 154) and Pietroski (2003). Note that item (D3) above speaks neutrally of ‘information’, in some sense that remains to be clarified by theory. Whether the relevant notion of information is to be glossed in terms of content is left open.

⁸ As an example, Kaplan’s definition of truth at a context could be construed as a principle of this sort. For relevant discussion see Yalcin (2007, 2011) and Ninan (2010b, 2012b).

mainly concerned to separate the semantic value of a sentence in context from its ‘assertoric content’ (or the ‘object of assertion’). (Compare Dummett 1973, 1993.) Let me explain the difference in emphasis between their discussions and the present one. I agree that insofar as there is a useful notion of specifically *assertoric* content, this notion must be disentangled from the notion of semantic value; but I think that contrast should be understood as a special case of the more general contrast between semantic value and content in the sense I emphasize here. (Perhaps they would not disagree.) I avoid framing in terms of ‘assertoric content’ for the following reasons. First, for all I have yet said to flesh the content role out, it is not obvious what it would even mean to say that sentences in context bear content. The role I have carved out so far refers only to mental states bearing content. One might of course go on and elaborate a meaning for ‘assertoric content’ in various ways, as many have; but however this gets cashed out, there will usually be an assumption to the effect that the things that realize the assertoric content role are of the type of the things that play the content role as I have (incompletely) characterized it. Thus it seems to me to cut to the chase to focus on the distinction between semantic value and content *simpliciter*, understanding the latter as a notion anchored primarily in an account of the intentionality of the mental and of folk psychological prediction. Second, in connection with epistemic modal and conditional talk, I have argued that the usual notion of the object of assertion breaks down in various cases, hence is of limited application (Yalcin 2007, 2012a). Third and related, it is clear that in theory, one could predict facts of informative communication without the assumption that every sentence determines a discrete item of informational content. Certain dynamic semantic systems are well known to achieve this result, for example (for relevant examples and discussion, see Groenendijk and Stokhof 1991, Yalcin 2012b, Rothschild and Yalcin 2012). So while I agree that one way to ground information transfer facts is to hypothesize that the semantic values of sentences in context determine assertoric contents, that hypothesis seems neither conceptually necessary nor empirically plausible in full generality.⁹

Analogous points can be made in connection with the demand that sentential semantic values help to explain truth and falsity judgments. There are ever so many ways sentential semantic values might ground such judgments without their being literally identified with items of content.¹⁰

These considerations do not lean on any empirical assumptions about what operators or binders English, or whatever, contains. I am merely emphasizing the difference in what constrains theories of semantic value and of content, respectively, given that

⁹ Lewis (1980) emphasizes the point that sentential semantic values (in context) may be *too rich* to be the realizers of the content role. The possibility of a dynamic semantics of the sort just described highlights a way that semantic values may be *too poor* to effect a general association between sentences in context and contents.

¹⁰ For discussion of some options in a particular kind of static semantic setting, see Yalcin (2011). For discussion of some options in a dynamic semantic setting, see Stokke (2014). But these barely scratch the surface.

we articulate their explanatory roles as above, and noting what the possibilities are in advance of inquiry into the semantic values of sentences. Looking at their quite distinct explanatory roles, it would be surprising if the realizers of the sentential-semantic-value-evaluated-relative-to-context role turned out to be the same kind of thing as the realizers of the content role.

2.3 *Against identification: empirical considerations*

Might it nevertheless just happen to be true that the realizers of these two roles turn out to be the same kind of object—say, in connection with English? That is a logical possibility. But as in effect already suggested, there is empirical reason to doubt that this possibility is realized in connection with actual natural languages, for there are a number of cases where otherwise quite well-motivated sentential semantic values seem not to be natural choices for the realizers of the content role. Let me briefly mention a few basic cases, referring the reader elsewhere for extended discussion.

First example: on several standard treatments, pronouns are modeled as variables, and sentences containing ‘referential’ (unbound) pronouns are taken to be open sentences. Given the standard semantics for quantifiers and lambda binders, the semantic value of an open sentence relative to context will be an object which is variable in truth with respect to assignment functions, in order that it compositionally predict the correct results for relevant embeddings of that clause (as evaluated at that context). Hence the compositional semantic value of a sentence in context will generally be, at least, a function from assignments to truth values. While a possible view about the realizers of the content role is that they are objects which are variable in truth value with respect to assignment functions,¹¹ it is fair to say that this is far from anything like a standard view about the nature of content. (It is notable, moreover, that virtually no one imagines that the standard variable-based semantics for pronouns entails, or recommends, such a view of content.) See Rabern (2013), where the point is pressed with great clarity.

Second example: it is known that the state of information relevant to the evaluation of an epistemic modal is shiftable systematically by attitude environments and by conditional environments (Morgan 1970; Stephenson 2007; and Yalcin 2007). Such facts recommend a semantic theory which relativizes extensions to some additional element beyond a context, a world, and an assignment function—for example, an accessibility relation (Van Rooij 2005; and Kratzer 2012), or a state of information (Yalcin 2007; Kolodny and MacFarlane 2010; MacFarlane 2011; and Klinedinst and Rothschild 2012). Thus even relative to a context and an assignment function, epistemic modal sentences appear to be variable with respect to something beyond a possible world. While a possible view about the realizers of the content role is that they are objects which are variable in truth value with respect to accessibility relations or states

¹¹ See Stalnaker (2008); Cumming (2008); and Ninan (2010a, 2012a) for some models that might encourage something like assignment-relativity for content.

of information in addition to worlds, this is far from anything like a standard view about the nature of content, and not a view that has seen any sustained development in the literature.

Third example: on the dynamic semantic theory of Heim (1982), the semantic value of (1):

- (1) A man walked in.

is an update operation on the ‘file’ of the conversation (a certain kind of context change potential). Informally speaking, the indefinite serves to ‘open a new file’ labeled with a new variable index—thereby it introduces a new ‘discourse referent’—and the sentence serves to eliminate those ways of mapping that index to an object which fail to map it to a man who walked in. Dynamic semantic approaches broadly in this spirit have been widely defended in the literature for a variety of constructions. But, to my knowledge, no one has ever suggested that such context change potentials would be apt realizers for the content role, as that role was (incompletely) described above. This is unsurprising, since superficially, it does not appear that they would be apt for this role. First, the structure needed to dynamically update and track discourse referents in conversation, in the specific way required by, for example, a proper treatment of (in)definiteness in natural language, is not obviously structure required by the realizers of the content role. Second, context change potentials generally underdetermine truth-conditions, whereas the realizers of the content role are widely imagined to determine truth-conditions.

Examples could easily be multiplied. Given what contents and semantic values, respectively, are postulated to explain, this should not be surprising.

Note that the point here is independent of the question of structure. Even if one thinks that the semantic values of sentences in context are structured entities (not functions of a certain sort, as I have assumed), and also thinks that the content of a sentence in context is a structured object, still, it will be implausible to maintain that these structured objects are of a kind, for just the reasons mentioned.¹²

I conclude we do better not to theorize under the hypothesis that the semantic values of sentences in context are literally of a kind with the realizers of the content role. Talk of ‘semantic content’ thus tends to blur a conceptually important distinction. So too does talk of a ‘character-content’ distinction.

2.4 *Relation to the debate about pragmatic enrichment*

In a separate (and much larger) literature, a number of theorists have defended the thesis that the semantic value of a sentence in context might, in some sense, underdetermine the propositional content asserted. Some remarks are in order concerning the relation between this literature and the foregoing.

¹² Rabern (2012b: ch. 2) contains further discussion. Collins (2007b) raises some other worries for taking the semantic values of sentences to be structured propositions.

Récanati has in many places developed the view that “[i]n general, even if we know who is speaking, when, to whom, and so forth, the conventional meaning of the words falls short of supplying enough information to exploit this knowledge of the context so as to secure understanding of what is said” (Récanati 1989: 298; see also Récanati 1993, 2010). Bach has in many places argued that, often, “the conventional meaning of the sentence determines not a full proposition but merely a *propositional radical*; a complete proposition would be expressed, a truth condition determined, only if the sentence were elaborated somehow” (Bach 1994: 127). Related themes are developed in Sperber and Wilson (1986), Bezuidenhout (2002), Carston (2002), Wilson and Sperber (2012), among many others. Citing these and related works, Soames has shifted to the position that “. . . the semantic content of S in a context *constrains* what S is used to assert, without always *determining* what is asserted, even when S is used with its normal literal meaning” (Soames 2008: 280, original emphasis; see also Soames 2010). I take it that most if not all of these authors would want to agree that, in general, the compositional semantic value of a sentence in context need not be identical to an item of informational content (a propositional content, a proposition). However, the focus in this literature has not been on the conceptual distinction emphasized here. Rather, it has to a great extent focused on the status of the following thesis:

Context-sensitivity traces to syntax. The contribution of context to the determination of the assertoric content of a sentence uttered in context is restricted to that of fixing the values of indexicals, free variables (e.g., referential pronouns), and other semantically context-sensitive syntactic material present in the logical form of sentence.

Bach, Récanati, and the others cited above are concerned especially to argue against this thesis. (Stanley 2000 is a *locus classicus* of this debate.) By contrast, we have not been particularly concerned with this thesis. Even if it were completely correct—which appears unlikely, though it depends, *inter alia*, on what ‘assertoric content’ means—it would still be necessary to distinguish the semantic value of a sentence relative to context from its content in that context, in the manner recently pressed. The explanatory roles of the notions would remain clearly distinguishable, and worth distinguishing. Second and more concretely, Rabern’s (2013) important point about the semantics of variables and variable binders, noted above, would still apply. Even if context is understood to ‘fix the values of free variables’, it must be recognized that the way in which this would work would not be akin to the way in which context is understood to fix the extensions of indexicals in a Kaplan (1977/1989)-style framework. Variables have semantic values that vary with the choice of assignment function, not with the choice of context. Thus insofar as context might ‘fix the values of free variables’, this would have to owe to some substantive further principle applying to whole unembedded sentences, a principle best understood as outside of (and indeed, defined in terms of) the compositional semantics proper. This would be an

extra-semantic bridge principle of the sort mentioned above, one spanning the gap between semantic value and content.

Thus the distinction I am pressing is plausibly well-motivated even on a highly ‘minimal’ conception of the scope for pragmatic enrichment. Plausibly it is needed already for the most basic sorts of variable-binding.

2.5 *Two varieties of meta-question*

With the distinction between compositional semantic value and content now in view, we may distinguish ‘descriptive’ and ‘meta-’ questions for each domain. Thus, concerning semantic values, we can ask, for any given expression *e*:

- (2) What is the semantic value of *e*?
- (3) In virtue of what does *e* have the semantic value that it has?

Call the first (following Stalnaker 1997) a question of descriptive semantics. We could also think of it as a modeling question (viz., what serves as the best theoretical model of the semantic value of *e*?). I take the second question to be a primary question of metasemantics. To a first approximation (and subject to significant refinement below), metasemantic questions are foremost questions about what grounds the descriptive semantic facts. They are about what is for the descriptive semantic facts to obtain.

Analogously, concerning content, we can ask, for any given bearer of content *s*:

- (4) What is the content of *s*?
- (5) In virtue of what does *s* have the content that it has?

The first we could call a descriptive (or modeling) question about content; the second we could call a foundational question about content.¹³

Because the notions of semantic value and of content have tended to get conflated, metasemantic questions are sometimes conflated with foundational questions about content. But they should not be conflated. Absent some further argument, it is natural to expect these questions to receive quite different—perhaps radically different—answers. Moreover, having separated these two questions, we can observe that, in principle, the descriptive facts of one of these two domains might conceivably be appealed to in answering the foundational questions of the other domain. One might think compositional semantic value facts are at least partly grounded in content facts; or one might think the reverse.

Separating these questions helps us also to see that, in principle, one but not the other of these meta-questions might well fail to be interesting, because somehow based on a mistake. For example, one familiar kind of view is that the notion of content

¹³ Much that has been called ‘metasemantics’ is responsive to question (5), rather than question (3). I have elected to use ‘metasemantics’ so that it concerns a subclass of issues in the foundations of natural language semantics, namely questions of the (3)-variety. Of course, there is nothing at stake concerning which we call ‘metasemantic’, so long as the questions are separated.

is not a fruitful one for empirical theorizing into language and mind. Perhaps it simply will not appear in any robust scientific explanation. (Of anything.) On versions of such a view, (5) would have no interesting, non-deflationary answer. But (3) (or the further precisification of it I develop below) might yet have an interesting and substantive answer. Alternatively, moving in the other direction, future semantics may reveal that the notion of semantic value does not, after all, mark anything like the right joint in linguistic nature: the things we are now dimly trying to explain by appeal to semantic values are in fact properly explained using some other notion of a quite different character. If that were right, then (3) would fall away, compatible with (5) surviving as an interesting, substantive question.

2.6 *The role of content and communication in semantics*

The considerations reviewed so far suggest that we step back from those ways of describing the task of semantic theory that place the notion of content, and attendant questions about intentionality, at the center of the enterprise. They also recommend that we not casually assume that the primary task of semantic theory is to compositionally vindicate commonsensical descriptions of linguistic communication events, doing so in terms of some notion of content, or to map sentences in context to items of information communicated. In this section I will stress the gap between semantic theory and a theory of communication. (I return to the role of the notion of intentional content in semantic theory in Section 7 below.)

Scott Soames has described

... a conception of semantics in which the fundamental task of a semantic theory is to tell us what sentences say in various contexts of utterance. On this view, the meaning of a sentence can be thought of as a function from contexts to what is said by the sentence in those contexts. The crucial question then becomes: What sort of things are in the range of this function? In other words, what is semantic content, or information? (Soames 1989: 394)

This is a view of semantics one can naturally pull out of Kaplan (1977/1989), and it seems to be one that Soames himself favors in the paper cited above. It is a familiar and common conception of the goal of natural language semantics. One can easily find this kind of picture in introductory texts on the philosophy of language. Some would describe it as the ‘standard conception’ of semantics:

According to the standard conception of natural language semantics, its purpose is to give an account of the relation between a sentence, on the one hand, and the information about the world communicated by an utterance of it, on the other. (Kennedy and Stanley 2009: 583)

This conception of semantics is problematic. It underestimates the gap between the notion of compositional semantic value and the notion of content, and it overstates the role of facts of communication both in constraining compositional semantic theory and in stating its primary explanatory aim. It also leaves it unclear why semantic

theory should be constrained by syntactic theory, since in principle, a mapping from sentences to their ‘communicated content’ might be finitely stated in myriad ways.¹⁴

I do not mean to suggest that semantic theory does not bear on the question how the linguistic communication of information is possible; obviously it does, as noted above. But what precisely its role is in a full theory of linguistic information transfer is a complex matter. Relative to syntax, pragmatics, and other nonlinguistic aspects of human cognition, the explanatory load carried by compositional semantics *per se* in accounting for facts of communication is far from obvious in the present state of understanding, and the subject of continuing inquiry. More important, such facts may supply data for semantic theory, without it being the case that semantic theory is well-described as a theory of what is communicated by sentences. (Just as linguistic judgments may supply data for semantic theory, without it being the case that semantic theory is properly understood as a theory of linguistic intuitions.) From the perspective on semantic theory I am recommending here, communicative uses of language reveal aspects of the state of mind which consists in understanding and speaking the language. It is the state of mind, or the semantic aspect thereof (‘knowledge of meaning’, understood in the technical sense), that semantics is foremost concerned with modeling. The question how best to model this state of mind, and the capacity it underwrites, is quite distinct from the question how this capacity is deployed in communication to transfer information. Considerations of communicative import may constrain semantic theorizing insofar as they supply a window into the state whose structure we are aiming to characterize.

I certainly have no objection in principle to the idea of theorizing about semantic values from the point of view of various abstract formal models of linguistic communication.¹⁵ Such theorizing has its place,¹⁶ though one should not underestimate the high level of abstraction and idealization normally involved. Rather, I am questioning the idea that the communicative import of a sentence or utterance, whatever that is, is the property that semantic theory principally investigates, and recommending an alternative conception of the subject matter.

¹⁴ As Lewis (1975) recognized. For further discussion, see Section 6 below.

¹⁵ My own preoccupation in this domain has been with developing systems that do not assume that linguistic information transfer always involves expressing propositional contents; e.g., Yalcin (2011, 2012a, 2012b); also Rothschild and Yalcin (2012).

¹⁶ Arguably, the place is small: productivity facts and judgment facts (about acceptability, entailment, truth, etc.) play a more dominant role in constraining hypotheses about semantic values than communication facts, at least in present inquiry. Indeed, it is a bit misleading to speak of ‘communication facts’ on par with ‘judgment facts’, as descriptions of the communication facts are inevitably more theory-laden. There is no uncontroversial notion of the ‘information transferred’ in linguistic communication; no uncontroversial way of measuring or describing the way information propagates between agents in conversation; and no uncontroversial way of isolating those aspects of information transfer owing specifically to coordinated states of linguistic competence as opposed to other features. Theorizing may proceed, but only under weighty assumptions. By comparison, facts about, for example, the markedness of sentences are far more scrutable.

It is a basic fact that many of the properties of greatest interest for the compositional semantics of a sentence are properties revealed most clearly, not when the sentence is considered in isolation unembedded (the case most relevant to its ‘communicative potential’), but rather when it is placed in various embedded contexts. Embedded contexts draw out those aspects of expressions which constitute their generative semantic contribution—their stable semantic benefaction to larger constructions. Theorists might agree entirely on the nature of this generative contribution—on compositional semantic values—and yet differ radically on matters of communicative import.¹⁷ Natural language semantics is principally concerned with the former.

3 The question concerns a technical concept of linguistics

Above I mentioned (without endorsing) the idea that compositional semantic value facts might be at least partly grounded in mental content facts. This idea seems to be the leading view of the matter.¹⁸ But often this idea is framed in a questionable way. There is thought to be a single category of properties, the semantic properties, which (i) are employed in some central way in natural language semantic theory, and (ii) apply to a heterogenous domain consisting of (at least) mental states and linguistic expressions. The issue is then framed as one about which kind(s) of thing in this domain have semantic properties in a nonderivative way. What is questionable about this framing is the idea that there actually are properties satisfying both (i) and (ii). Absent some substantial argument, it is quite unclear whether the theoretical concepts employed in natural language semantics sensibly apply outside of semantics, to nonlinguistic entities.

I am not disputing the point that in ordinary parlance we apply a certain class of predicates—*is true*, *refers*, *is about*, *represents*, *means that*—variously in reference both to natural language expressions and to mental states (not to mention speech acts, persons, pictures, natural indicators, and more besides). But these are ordinary notions, playing virtually no role in natural language semantic theory beyond a superficial level. Like any other nontrivial explanatory enterprise, semantic theory employs its own theoretical notions and posits, and its chief questions and these are

¹⁷ Given some fragment of language, contextualism, relativism, and expressivism are typically conceived of as alternative semantic views. But more often, the differences between these views appear only at a post-semantic or pragmatic level. For example, the contextualist, the relativist, and the expressivist about epistemic modal talk could all agree entirely on the compositional semantics of epistemic modals (compare, e.g., Yalcin 2007 and MacFarlane 2011). Their disagreement concerns the question what role the sentences play in communication unembedded. The positions reflect differences on the question how one moves from the compositional semantics of the target sentences to their context-change potentials, and on the question what state of mind is characteristically expressed by the sentences unembedded. Stojanovic (2012) emphasizes related points.

¹⁸ See the so-called ‘Gricean program’. Notable alternative views are developed in, e.g., Davidson (1984), Brandom (1998).

stated in terms of these. Semantic theory is not beholden to ordinary notions any more than physics is beholden to folk contact mechanics. If metasemantic questions concern the aspect of reality that semantic theory seeks to characterize the structure of, they are best understood (at least in the present state of inquiry) as questions concerning the ground of properties stated largely in terms of the notion of semantic value. The notion of a semantic value is a technical concept not equivalent to, and not particularly resembling, any ordinary notion, and it is one not obviously applicable to nonlinguistic expressions.

In connection with the ways in which semantic theory, like other sciences, goes beyond ordinary notions, it is worthwhile to comment specifically on the notion of truth, since truth has often been thought to be central for delineating the subject matter of semantics.

Sometimes a serious explanatory enterprise begins with questions stated using ordinary notions and preconceptions. As inquiry progresses, ordinary notions are replaced with more appropriate technical notions. Sometimes, the technical notions of a mature theory bear a family resemblance to an ordinary language ancestor. So it is, for example, with the ordinary notion of truth and the technical notion of the truth of a sentence at a point of evaluation in a model, as the latter is employed in certain traditional intensional natural language semantic systems (e.g., Lewis 1970; and Montague 1974). But how these two notions of truth are related is a wide open question. Care should be taken here. The fact that the folk feel comfortable in describing sentences, things asserted, and things believed as true or false supplies little reason to think that the notion of truth at a point of evaluation in a model, as that notion emerges in natural language semantics, is one that particularly makes sense to apply to nonlinguistic entities. Moreover, it needn't be the case that a technical notion of truth postulated in semantics in order to model the structure of linguistic understanding must correspond, in some nice way, with any ordinary or folk notion of truth.

The semanticist does aim for semantic values which, *inter alia*, explain a range of folk judgments of truth. (As in (D5) above.) But there are ever so many imaginable ways one might connect a recursive definition of truth in a model, postulated primarily to reflect the compositional structure of a competent speaker's productive linguistic ability, with folk talk of truth. It should go without saying that the former notion need not slavishly follow the latter, or map onto it in some self-evident way.¹⁹

Further, granting some technical notion of truth is required by the best semantic theory, such a notion is not particularly required to describe the subject matter of semantics. Similar remarks could be made in connection with the notion of reference. (One could introduce a technical sense of 'refers' according to which expressions refer to the semantic values that they have, but it would remain to show that superadding

¹⁹ I have spoken, in the last two paragraphs, as though final semantic theory actually will involve a technical notion of truth in some important way. But I note that this is very much an open question. Dynamic approaches in semantics, for instance, tend not to centrally involve such a notion.

this to semantic theory adds any explanatory value. See Section 7 for further discussion.) The roles of these notions in characterizing semantic theory should not be overstated. Natural language semantics is not helpfully understood as the study of truth and reference, though it may (or may not) employ technical versions of these notions. Rather, it is concerned to explain a dimension of the ability to understand and speak a language, as described above.

Slurring over the semantic value-content distinction can make matters here yet more confusing. Perhaps a mature theory of content, understood as distinct from natural language semantic theory, will likewise employ some technical notion bearing a resemblance to the ordinary notion of truth, one applying to whatever the theory says contents are. Maybe it will employ a notion of truth applicable to structured n -tuples of senses, or to structured n -tuples of objects and properties, or to sets of centered worlds. If so, then care should be taken not to conflate this technical notion of truth with the technical notion of truth emerging in (some traditions of) natural language semantics. These are going to be different notions, and their relations may not be self-evident.

Summarizing: there is the notion of the truth or falsity of a sentence at a point of evaluation in a model, as it might emerge in natural language semantics; there is the notion of the truth or falsity of content, as the notion of content might be refined and modeled in a theory of content; and then there is ordinary loose thought and talk of truth and falsity. The three notions are distinct from one another. Maybe the technical notions will illuminate different aspects of ordinary thought and talk of truth and falsity. But maybe not. Meanwhile, one needn't speak of truth in order to say what semantics attempts to explain. Finally, if there is a central technical notion in contemporary natural language semantics, it is the notion of semantic value.

4 Not about languages qua abstract objects

We said that if descriptive semantics asks what the semantic value of e is, then metasemantics chiefly asks:

- (3) In virtue of what does e have the semantic value that it has?

But the content of this kind of question is far from obvious.

A first point to emphasize is that this is a broadly metaphysical 'in virtue of' question, one seeking a 'deeper explanation' of the relevant facts or properties in more fundamental terms. We are asking a version of the question what makes it the case that the descriptive semantic facts are true.²⁰ One could frame the question

²⁰ In a sense, we are asking what it is for the descriptive semantic facts to obtain. I sometimes engage in 'what it is' talk throughout. But 'what it is'-talk is sometimes taken to track questions about essence or 'real definition' rather than ground or metaphysical explanation (compare Pérez Carballo 2014). That is not what I have in mind. I would not want to suggest that the metasemantic question is one about the essence or the real definition of descriptive semantic facts.

as one about what the semantic facts or properties metaphysically or ontologically depend on (compare Rosen 2010), or as a question of what grounds the semantic facts or semantic properties (compare Fine 2001, 2012; and Rosen 2010) or as a question seeking a comparatively more fundamental characterization of the semantic value facts or properties (compare Sider 2011), or in some other way. I don't assume a particular approach to theorizing about questions of fundamentality *per se* here; I just take for granted some such notion is legitimate, as seems obvious. (I have tended to describe the metasemantic question in terms of ground and will continue to do so, but without the intention to presume a specific conception of fundamentality.) I take it that, generally, what is especially sought in metasemantic inquiry is some illuminating reduction: we should like to know whether and how we can be supplied with a nonsemantic explanatory basis for the descriptive semantic facts. We are not asking for a conceptual analysis, or an elucidation of the meaning of 'meaning' or 'semantic value'. Neither is the question about what the semantic facts supervene on, though answers to supervenience questions can be expected to follow from answers to questions of ground.

A second issue concerns the modal status of the facts that descriptive semantic inquiry is supposed to uncover, and that metasemantic inquiry is supposed to ground. It is natural to theorize with a use of 'expression' according to which expressions have semantic values necessarily—indeed, have the semantic values that they in fact have necessarily.²¹ This fits with the thought, standard in syntax and semantics, that expressions should be individuated in part by their semantic values, and that a language should be individuated in part by the way it pairs sounds with meanings.

This use of 'expression' (and of 'language') seems to me the theoretically appropriate one. But note that if expressions are assumed to have their semantic values necessarily, and the questions of descriptive semantics are taken to be of the form 'What is the semantic value of *e*?' (as in (2) above), then it follows that descriptive semantics is really in the business of uncovering a class of necessary truths about expressions. Correspondingly, it follows that metasemantics is really in the business of investigating the grounds of such necessities. Metasemantics, on this interpretation, would ask something like:

- (6) In virtue of what is having the semantic value *m* a necessary property of *e*?

I take it as obvious that this yields the wrong conception of descriptive semantics, and of metasemantics. Semantic theory is not helpfully understood as an inquiry into the necessary truths about expressions. The facts uncovered in descriptive semantic inquiry are largely empirical and contingent. Metasemantics is interested in the

²¹ A less restrictive use of 'expression' might be more appropriate in certain inquiries outside of the usual domain of natural language semantics—for instance, in theories of language change.

ground of those facts. Questions like (6) may be intelligible and worthwhile (or not), but metasemantics does not proceed in this direction.²²

What this shows, I suggest, is that our first attempt at stating the descriptive semantic question is too crude.²³ We should ask: in what sense exactly are the facts uncovered in descriptive semantic inquiry empirical?

5 The question concerns what it is to understand a language

I suggest we clarify the situation by borrowing variously from certain ideas in Chomsky (1956/1975, 1980, 1986) (certain related themes occur in Davidson 1965, 1967; Partee 1975; Larson and Segal 1995; and Pietroski 2003, among others). The basic structure of the picture is already in place. Human beings have the ability to speak and understand languages. Linguistic syntax and semantics are ultimately concerned to produce explanatory models of aspects of the knowledge that underwrites this capacity. They seek formal specifications of languages ('grammars', on a certain broad use) that can serve in explanatory models of the states of mind of the speakers of those languages. On the approach to semantics I take for granted here, languages are modeled as abstract objects which, *inter alia*, determine a mapping from some class of sounds—better, some natural class of syntactic objects (lexical items, or primitive morphemes) together with an associated class of phrase structures generated from that lexicon, in whatever style is recommended by the best syntactic theory for natural language—to semantic values. Semantic theory is not interested in the semantic value properties of these abstract objects *qua* abstract objects. Rather, it is interested in an aspect of the question which of these abstract objects well-models what it is one knows, when one knows a language. We might try putting the general shape of the sorts of descriptive questions it asks like this:

- (7) Given a linguistically competent subject *S*, what is the semantic value of *e* in the language *L* that *S* speaks and understands?

Contingency enters in, not because expressions are contingently connected to their semantic values (on the formal notion of language stipulatively in play, they are not), but because it will be a contingent matter which abstract object *L* is the one that does best in modeling the relevant linguistic abilities of the speaker (or sort of speaker)

²² Perhaps it would if we construed natural language semantics as a branch of mathematics, a proposal sometimes associated with Montague (see, e.g., Richmond Thomason's introduction to Montague 1974) and Katz (1981); but this idea is stupefying.

²³ One might think one could address the problem just described by using 'expression' in such a way that expressions are only contingently connected to their semantic values. My view is that this would raise spurious puzzles about how expressions are to be individuated, and would yield a distorted conception of the subject matter of semantics; but I lack the space to elaborate. Please read what follows, then, as simply exploring a different reaction to the problem.

in question—that is, which does best in modeling the productive character of their linguistic understanding, their capacity to recognize entailment relations between expressions in the language, and so on.²⁴

Naturally, what class of abstract objects are to be called ‘languages’ is not fixed once and for all in advance, but is rather uncovered as empirical inquiry into the state we informally describe as ‘understanding a language’ proceeds. In line with standard approaches in semantics, I assume a language incorporates at least (i) a finite lexicon associating elementary constituents with semantic values, (ii) a finite set of syntactic operations which build larger constituents by structuring smaller constituents, and (iii) some finite set of composition rules for determining the semantic values of complex expressions as a function of their elementary constituents and their syntactic structure.²⁵

If, as I want to suggest, something like (7) does better than (2) at getting at the descriptive questions of semantic theory, we should adjust our conception of metasemantics accordingly. Metasemantics is concerned with a part of the story about what it is to know a language, the part which clarifies in virtue of what one knows a language instantiating such-and-such semantic value properties. We could put it like this:

- (8) When a subject knows a language L wherein e has the semantic value m , in virtue of what is this the case?

The reductive ambition is to cash this out in a manner not itself ultimately appealing to the notion of semantic value.

It is occasionally useful to employ a technical term to refer to the mental relation one bears to a language one knows, in order to avoid potentially misleading associations of the ordinary word ‘knows’ and to emphasize the theoretical nature of the enterprise. I will use ‘cognize’ for this purpose (compare Chomsky 1980; additional discussion in Collins 2007a). Cognizing is not a relation to a proposition, or to a set of propositions; neither is it ‘know how’ in some ordinary sense. It is not presumed to be an intentional state. Rather, it is a relation to a language, in the technical sense: we might model it as a relation between a subject and some tuple including at least a lexicon, a system of generative rules, and some composition principles, together with whatever else we find a need to posit as theorizing into language proceeds.²⁶

²⁴ There may be other semantic properties of expressions, or of the language as a whole, not directly reflected in compositional semantic value, but which we should like semantic theory to help systematize and clarify. For example, type-shifting principles (e.g., Partee 1987), or principles of economy (e.g., Fox 2000), not to mention argument structure properties. Such features may give rise to further metasemantic questions beyond those I isolate.

²⁵ On some ways of using ‘language’, multiple grammars can determine the same language. As I hope is clear, I am not using ‘language’ in this way. (Substitute ‘grammar’ if you prefer.)

²⁶ Contrary to what is sometimes assumed in discussions of approaches along these lines (e.g., Devitt 2006), I am not taking it that the state of understanding a language ‘represents’ a language or grammar. Certainly it is not some kind of complex mental symbol for the grammar. (Ludlow’s (2011) preferred locution, ‘having a language’, seems to be roughly equivalent to what I intend by ‘cognize.’) See also Collins 2006.

(In fact it may be preferable to view the state as a relation to a set of (precise) languages, the languages ‘admissible’, or not ruled out, according to the state. This could assist in modeling vagueness, gradient acceptability data, and the process of acquisition (compare the acquisition model of Yang 2002). It could also facilitate accommodation of the basic point that individuals we informally describe as ‘speaking the same language’ are generally not coordinated on precisely the same language; rather their states overlap substantially in what languages they admit.)

A crude analogy to the postulation of the mental state of cognizing would be the postulation of states of credence in Bayesian decision theory. In that enterprise, agents are described as having doxastic states with a certain kind of overall structure—commonly, with the structure of probability spaces of some kind. Insofar as the explanatory emphasis is on agents and rational action, these projects are not interested in probability spaces (etc.) *per se*, but rather in probability spaces (etc.) *qua* models of (what we informally describe as) rational, uncertain states of mind. The probabilities in the formal theory model confidence; the probability measure is not characterizing some aspect of mind-independent reality putatively grasped by the agent being modeled. In a similar vein, we are interested in languages, abstractly understood as above, insofar as they can effectively serve to model the state of understanding a language, so that to theorize about the structure of a natural language just is, in part, to theorize about the structure of the state of cognizing that language. It is not that there are independently moored semantic value facts, and that we get into states of cognizing these facts; rather, talk of semantic value facts just is, indirectly, talk of cognizing states. Within this framework of assumptions, the metasemantic question is an aspect of the larger question of what it is to cognize a language. Thus it might also be put like this:

- (9) What is it to cognize a language wherein e has the semantic value m ?

(The analogy should not overextended. Insofar as a decision-theoretic representation of agents is sought as a regimentation of commonsense ideas about belief and desire (compare Lewis 1994), it is disanalogous to the postulated state of cognizing a language.)

As I intend it, ‘cognize’ does not definitionally incorporate any assumptions about whether the corresponding mental states are narrow or wide—this is just a further question (indeed, one aspect of the above question).²⁷ Neither does it prejudice the question whether the state is best understood as one individuated essentially by reference to a population of speakers, real or idealized. Below I will generally use ‘cognize’, ‘knows’, and ‘understands’ interchangeably to refer to this technical notion.

²⁷ Chomsky’s notion of an *I-language* (Chomsky 1986) assumes an internalist approach to the sort of state I am describing; so too the ‘conceptualist’ approach of Jackendoff (2002). Ludlow’s (2011) notion of a *ψ-language* is closer to the level of abstraction I have in mind.

To be clear, note that the question is, again, one about the explanatory ground of a state. It is not about the causal genesis of the state. We are not asking the question what caused a given agent to arrive at a state of knowing a language wherein e has the semantic value m . Questions of language acquisition are of course intelligible and interesting, but they are quite separable, and I take it not a part of metasemantics, though considerations of acquisition can be expected to strongly shape our conception of the properties of language-cognizing states, hence of their ground.²⁸

6 The question is heavily constrained by compositionality

Perhaps surprisingly, my gloss on model-theoretic semantics has more in common with the theoretical apparatus and disposition of Chomsky (1986) than with that of Lewis (1975). Despite appearances, the latter work was not strictly speaking concerned with a concept of semantic value of the sort described here, wherein productivity facts are among the central explananda. Lewis expressed no interest in accounting for productivity. On the contrary, he explicitly expressed skepticism that semantics could be supplied with determinate foundations at the level of subsentential expressions. He offered only to ground the notion of a *population's using a language*, where 'language' in his sense refers just to a function which pairs sentences with meanings, not the far richer object assumed above. This limited set of facts of 'semantic value' were to be grounded in certain conventions—on Lewis's analysis, in certain regularities in belief and action prevailing in a given population, owing to some common interest in communication. This story about the ground of sentential semantic value facts did not generally settle the choice between two compositional semantic theories agreeing on the semantic values of sentences, but differing at the subsentential level. At the subsentential level, we find substantial indeterminacy. Lewis (1975) was quite prepared to accept such indeterminacy.

Since Lewis had no evident interest in trying to explain, even in an abstract way, the productive character of linguistic competence, it is not totally clear why he took compositionality to be a desideratum in natural language semantic theory. Judging from Lewis (1980), his motive seems to have been that theorizing under the assumption of compositionality was the best route to a "concise" (21) and tractable "systematic restatement of our common knowledge about our practices of linguistic communication" (26). The following passages give some of the flavor of his view:

²⁸ As an illustration, it could be, for example, that recognition of intention is essential to some aspects of ordinary language acquisition (Baldwin 1991; Bloom 2002; and Sabbagh and Baldwin 2005), but that states of intention themselves are not constitutive of knowing a language in any way, hence irrelevant to the metasemantic question. If so, semantics would not be 'intention-based' in the sense relevant to metasemantics.

Different compositional grammars may assign different sorts of semantic values, yet succeed equally well in telling us the conditions of truth-in-English and therefore serve equally well in chapters in the systematic restatement of our common knowledge about language. Likewise, different but equally adequate grammars might parse sentences into different constituents, combined according to different rules.

More ambitious goals presumably would mean tighter constraints. Maybe a grammar that assigns one sort of semantic value could fit better into future psycholinguistics than one that assigns another sort. Thereof I shall not speculate. (Lewis 1980: 26)

Lewis was concerned to capture the meaning properties of the *sentences* of a language within a finitely storable theory, but not in order to illuminate the productive character of speaker competence. I conjecture his motive for embracing compositionality was simply that finitely storable theories are surely nicer to work with than the longer variety.²⁹

When in later work Lewis came to accept a metaphysical notion of naturalness (Lewis 1983), he allowed appeal to this notion to help fix, together with the convention facts at the level of sentences, the semantic value properties of subsentential expressions (Lewis 1992). Thereby he provided subsentential semantics a foundation missing in his earlier work. However, the resulting theory is still not directed at illuminating the productive character of language use. More important, it is not at all clear that it could even possibly explain productivity. As hinted above, I take it the productivity property is normally accommodated within standard approaches in natural language semantics by the assumption of something like the following thesis concerning the state of cognizing a language *L*:

Compositionality grounding principle. For a competent speaker of language *L*, knowledge of the semantic value of a complex expression *e* of *L* is (generally, largely) grounded in knowledge of the semantic values of the primitive constituents of *e*, together with knowledge of the syntax of *e* and the relevant composition rules of *L*.

This is a metaphysical claim about a class of mental states—not commonsense folk psychological states, but cognizing states. It has the status of a methodological assumption. If this reflects the order of explanation required to cover productivity—I am unaware of viable alternatives—then one cannot, as Lewis does, ground the state of knowing sentential semantic values in something other than the states of knowing the semantic values of the constituents of the sentence.³⁰ If this is correct, and states

²⁹ Compare Jackendoff's (2002) worry that on something like the Lewisian conception of language, "the subject matter of linguistics becomes limited to the mere description of languages, and the study of Universal Grammar becomes at best an exercise in statistical tendencies and/or formal elegance" (297).

³⁰ Compare Chomsky (1986), discussing Lewis-like views: "Despite appearances, the problem of accounting for the unbounded character of the E-language and the person's knowledge of language including this fundamental property is not squarely addressed in such approaches" (20). ("E-language" being Chomsky's name for the technical conception of language broadly in Lewis's style.) See also Chomsky (2003: 288).

of knowing the meanings of complex expressions are grounded largely in states of knowing the meanings of their parts, then Lewis's specific grounding project cannot be accepted.³¹ Knowledge of the meanings of complex expressions is simply not possessed in virtue of knowledge of the parts of those expressions, on this conception.

Since Lewis's metasemantic picture has been influential among those engaged in formal semantic inquiry, it is perhaps worth restating this point in a slightly different way. Lewis offers a theory of in virtue of what an agent counts as 'using' or 'knowing' a semantics wherein sentence *s* means *p* (perhaps relative to context *c*). What is required, says Lewis, is (crudely) that the agent belong to a population wherein certain conventions (of truthfulness and trust) prevail in connection with the (Lewisian) language that *s* belongs to. But this fact is not in turn partly grounded by Lewis in a further fact about the agent, namely his knowledge of meanings of the constituents of *s*. Rather, the prevalence of the relevant conventions is just a matter of certain regularities in belief, desire, and action obtaining in a population that the agent belongs to. Such regularities might prevail, whatever the compositional structure of *s*, or perhaps even if it had no compositional structure.³²

I am sometimes reassured by those sympathetic to Lewis that, of course, owing to the complexity of human language, one must be in a psychological state representable by grammar, or a language in the thicker sense I have worked with above, in order to participate in the sorts of linguistic conventions that characterize human populations. This kind of state is a prerequisite, that is, to participation in the convention.

I doubt that an agent which had the behavioral dispositions characteristic of the Lewisian linguistic conventions of some population of speakers must thereby be in a language-cognizing state of the sort I have been sketching. Anyway, that remains to be shown; it is not obvious. But supposing we granted this, it would miss the point. The question is: what is it to be a knower of the meanings of the expressions of the

³¹ Though it remains open whether another account employing some notion of convention can be shown to succeed.

³² I am granting for the sake of argument in this discussion that Lewis's way of grounding meaning in terms of a notion of trust, which is in turn articulated in terms of the probability that the sentences of the language are uttered (see, e.g., Lewis 1992), can at least roughly demarcate the class of declarative sentences belonging to the language used by a population. But it is worth noting that this assumption is implausible, for reasons noted long ago by Chomsky in reply to Quine:

On empirical grounds, the probability of my producing, say, "birds fly" or "Tuesday follows Monday," or whatever—is indistinguishable from the probability of my producing a given sentence of Japanese. Introduction of the notion of 'probability relative to a situation' changes nothing, at least if 'situations' are characterized on any known objective grounds (we can, of course, raise the conditional probability of any sentence as high as we like, say to unity, relative to 'situations' specified on *ad hoc*, invented grounds) In fact if the "complex of dispositions" is determined on grounds of empirical observation, then only a few conventional greetings, clichés, and so on, have much chance of being associated to the complex defining the language, since few other sentences are likely to have a non-null relative frequency, in the technical sense, in any reasonable corpus or set of observations—we would, for example, expect the attested frequency of any given sentence to decrease without limit as a corpus increases, under any but the most artificial conditions. (Chomsky 1968b: 57–8)

language, in the sense that compositional semantics is concerned with? Is it to be a participant in a certain set of communicative conventions in a population, or is it to be in a certain underlying language cognizing state of the sort envisaged above, a state which perhaps, *inter alia*, makes such participation possible? If an explanation of our productivity is part of what is desired, and if the compositionality of our semantics on psycholinguistically motivated syntactic structures is what we take ourselves to be constrained by in semantics, I see no escape from the latter kind of answer.³³

The population-level facts Lewis is concerned to systematize are (candidates for) the facts in virtue of which sentences can be deployed to transmit informational content. They could be described as postsemantic facts, or pragmatic facts. I am not at all denying there may be interest in a level of description of a 'language' that prescind from subsentential and syntactic structure, and which only characterizes the sorts of information flow the sentences of the language make possible in a population. My point is just that this enterprise is strictly speaking outside of the domain of compositional semantics proper. As we have already noted in detail above, one can state a compositional semantics for natural language compatible with saying very little about how the sentences of the language are deployed by speakers to transfer information.

So while we agree with Lewis (1980) about the importance of separating the notions of semantic value and content, we disagree with him about what, ultimately, to mean by 'semantic value'—that is, about what conception of the semantic value role is most theoretically fruitful. The *raison d'être* of the assumption of compositionality in natural language semantics is the aim of explicating the productive, generative character of the linguistic understanding of competent speakers. Semantics with no treatment of productivity is not semantics.

Given something like the compositionality grounding principle, it is natural to say that insofar as the central problem in metasemantics is to explain semantic features in nonsemantic terms, it is partly a problem about lexical primitives, and partly a problem about the generative apparatus that provides for the possibility of productivity. We could thus distinguish two questions:

- (10) **Knowledge of lexical semantic values.** What is it to cognize a language wherein primitive expression *e* has the semantic value *m*?
- (11) **Knowledge of the generative rules.** What is to cognize a language whose generative rules determine knowledge of the semantic values of complex expressions, given knowledge of its primitive constituents?

The answer to (11) is, *inter alia*, supposed to help clarify in virtue of what the compositionality grounding principle is correct. The answers to these two questions may

³³ For a different kind of response to Lewis, see Schiffer (2006). (I regret I lack the space to compare Schiffer's response to my own.)

be tightly connected, perhaps inseparable; or they may involve relatively distinct substates.

Reviewing: natural language semantic theory seeks to produce an explanatory model of the state of cognizing a language, with attention to the semantic aspects of this competence. In so doing it appeals to semantic value properties of expressions, and to the idea that cognizing a language is partly a matter of knowing the semantic values of the expressions of the language. Metasemantics asks for a more fundamental characterization of such states of mind. The assumption of compositionality, postulated principally to accommodate the productivity of language understanding and use, suggests an initial separation of two questions: one asking in virtue of what a competent speaker possesses knowledge of lexical semantic values, and one asking in virtue of what a competent speaker possesses knowledge of the generative rules of the language. Thus a central class of metasemantic questions come largely to questions of the (10)-variety, together with (11).

These questions seek more fundamental characterizations of the target semantic facts, but they do not specify the relevant ‘more fundamental level’, the level at which their answers are to be stated. This is because that is part of what is being asked. We should read the questions as implicitly asking what the nearest relevant illuminating reduction base is, if such there be.

There are various preliminary questions we might attempt to isolate and address *en route* to these questions. A notable preliminary concerns whether knowledge of semantic value properties *per se* is internally grounded. This is the question:

- (12) If one cognizes a language wherein primitive expression *e* has the semantic value *m*, is this entirely in virtue of the state of one’s brain?³⁴

It is not hard to read many philosophers engaged in metasemantic inquiry as concerned especially with this question. Within the present framework of assumptions, this is a way of stating the question of semantic externalism. This question should be distinguished from an analogous question one could ask about content:

- (13) If one is in mental state *s* with content *p*, is this entirely in virtue of the state of one’s brain?

This would be one way of stating the question of content externalism. Perhaps the answers to (12) and (13) are intimately related. But in light of the preceding considerations, it is more plausible to assume that their answers are farther apart than has been commonly supposed. It is not obvious, for instance, that semantic internalism could

³⁴ As I intend this question, the answer could be affirmative, compatible with the overall state of cognizing *L*’s being a wide state. Compare: the state *having 0.7 credence in the proposition that water is wet* could be a wide state, compatible with the thought that the credal dimension of the (wide) state is fixed by internal factors; that is, granting doxastic states are constitutively wide, one could yet be an internalist about the credal distribution *per se*. In a similar way, the present question is about the ground of an *aspect*—the semantic aspect—of a more general state, that of cognizing a language.

not be combined with content externalism.³⁵ Suffice to say, we should not conflate (12) and (13).

Where might the externalist insights of Kripke (1972/1980), Putnam (1975), Kaplan (1977/1989), and Burge (1979) fit into this picture? One can certainly see some clear ways of bringing this work to bear on questions about the nature of content (see, e.g., the enlightening developments of content externalism in Stalnaker (1984, 1989, 2008) and Williamson (2000, 2006)). It is less obvious how this work should be understood to bear on compositional semantic theory, once the gap between content and semantic value is recognized. In many ways, the bearing of this work on semantics is a function of the size of that gap, a matter still unclear. In the next section, I try to give a sense of why caution is advisable here, and I try to say more about why semantic theory does better to prescind insofar as possible from commonsense or intentional characterizations of its main notions.

7 Translation, interpretation, modeling

The overall theoretical orientation I have advanced will seem puzzling from certain more usual ways of glossing model-theoretic semantics, namely those according to which its primary task is to explain or model language-world relations. On the version of this approach I have in mind, expressions have model-theoretic interpretations as their ‘intentional contents.’ The foremost task of a semantics is to tell us what each symbol stands for—what each symbol represents.³⁶ Generally speaking, expressions will be taken to represent mind-independent objects, properties, or propositions (etc.) as formalized in the model, usually according to their syntactic category. This viewpoint is sometimes presented as virtually inevitable, the main alternative being some kind of non-explanatory translation procedure that simply maps linguistic expressions into some artificial technical language, the latter left uninterpreted or tacitly understood. Thus Lewis (1970), in a famous passage:

Semantics with no treatment of truth conditions is not semantics. Translation into Markerese is at best a substitute for real semantics, relying either on our tacit competence (at some future date) as speakers of Markerese or on our ability to do real semantics at least for the one language Markerese. (18)

Lewis is quite right to set aside Markerese. But his discussion presumes a certain very strong parallel between natural language semantics and the semantics of artificial

³⁵ To be clear, a package view of this sort should not be confused with the kind of two-dimensionalism about content which identifies narrow content with something like Kaplanian character (e.g., Fodor 1987). Semantic internalism is not a version of the view that there is a category of narrow content.

³⁶ Compare Devitt and Sterelny 1999, an introduction to the philosophy of language: “It is undoubtedly plausible to think that representing is the core of meaning in general, and hence that having a certain truth-condition is the core of a sentence’s meaning in particular . . . on this hypothesis, our main semantic task becomes that of explaining the representational properties of linguistic symbols. Applied to an indicative sentence, our task becomes that of explaining its property of being true if and only if a certain situation obtains” (20).

formal languages. The parallel is not justified, however, and as a result, his remarks present a false choice. It should not be assumed that natural language semantics aims at delivering the putative intentional content of expressions, or the referents of expressions in the technical sense appropriate to the relation between expressions in artificial formal languages and their stipulated interpretations (*pace* Montague 1970). Neither is it preoccupied with finding a translation procedure into Markerese, or Mentalese, or some other supposedly more fundamental language. Rather, it is aimed at modeling, in some predictive and explanatory way, a certain category of properties, namely the semantic properties of expressions.

It is a platitude that modeling a class of properties is, generally speaking, simply not the same as offering some kind of translation procedure operating on the bearers of those properties, and neither is it the same as giving an interpretation procedure in some intentional sense, or in the sense familiar from artificial languages. This platitude should be respected, even when what is to be modeled are the semantic properties of natural language, and even when the modeling proceeds using formal tools developed originally for stipulating model-theoretic interpretations for artificial languages. One can use model-theoretic tools to model meaning properties without assuming that in doing so, one must be associating expressions with their supposed referents or representational contents. Again, while that kind of gloss may be natural when one is using such tools to *stipulate* referents or meanings in an artificial formal language (for, say, logic or mathematics),³⁷ it is question-begging, and must be justified, when the tools are deployed in connection with the empirical study of natural language. Meanwhile, such a gloss is not a prerequisite to ‘real semantics’.

It is ironic that Lewis complains that Markerese translation would illicitly rely on our “our tacit competence,” in supposed contradistinction to the model-theoretic way. Unacknowledged reliance on tacit competence is precisely the locus of Chomsky’s criticisms of work in model-theoretic natural language semantics—or rather, his criticisms of what what I am calling the conventional way of glossing that work. Consider, for example, semantic theories that assume semantic values like the following (compare Kripke 1972/1980; Putnam 1975; and Heim and Kratzer 1998):

- (14) $\llbracket \text{London} \rrbracket^M = \text{London}$
- (15) $\llbracket \text{house} \rrbracket^{M,w} = \lambda x.x \text{ is a house in } w$
- (16) $\llbracket \text{water} \rrbracket^{M,w} = \lambda x.x \text{ is water in } w$

Chomsky has suggested such lexical entries register “a lack of interest in the problems” (Chomsky 2003: 290). This view seems correct. Save for their type-theoretic properties and some elementary interactions with modal operators, significant semantic dimensions of these words are left unexplicated by the above. The use of the object language

³⁷ On whether is accurate to even gloss what one is up to in model theory in this way, see Burgess (2008) for relevant discussion.

in the metalanguage largely flags the place where further explanation has not yet been seriously attempted, not where it has finally occurred. Chomsky emphasizes what is missing in numerous places (see also Pietroski 2003). For instance (from Chomsky 2000):

We can regard London with or without regard to its population: from one point of view, it is the same city if its people desert it; from another, we can say that London came to have a harsher feel to it through the Thatcher years, a comment on how people act and live. Referring to London, we can be talking about a location or area, people who sometimes live there, the air above it (but not too high), buildings, institutions, etc., in various combinations . . . (36)

If London is reduced to dust, it—that is, London—can be rebuilt elsewhere and be the same city, London. If my house is reduced to dust, it (my house) can be rebuilt elsewhere, but it won't be the same house. If the motor of my car is reduced to dust, it cannot be rebuilt, though if only partially damaged, it can be. Pronouns involve dependency of reference, but not necessarily to the same thing; and both referential dependence and the narrower notion of sameness involve roles in a highly intricate space of human interests and concerns. Judgments can be rather delicate, involving factors that have barely been explored. (126)

. . . whether something is water depends on special human interests and concerns. . . . If [a cup] contains pure H₂O into which a tea bag has been dipped, it is tea, not water, though it could have a higher concentration of H₂O molecules than what comes from the tap or is drawn from a river. (128)

Proceeding beyond the simplest cases, intricacies mount. I can paint the door to the kitchen brown, so it is plainly concrete; but I can walk through the door to the kitchen, switching figure and ground. The baby can finish the bottle and then break it, switching contents and container with fixed intended reference. (128)

These empirical observations do not show (and are not claimed to show) that natural language semantics is somehow not possible, or that model-theoretic tools are not useful in such work. Nor do they show that lexical entries such as (14)–(16) are not theoretically appropriate in many contexts—they often are, I would suggest.³⁸ Rather, the point is not to mistake the place where numerous semantic problems have been idealized away for the place where some deep explanation occurs, in the process using a notion of reference to characterize the project borrowed from the construction of stipulated formal languages (or worse, from common sense). To do that is to embrace a controversial, non-obvious, and potentially simplistic understanding of the way in which natural language semantics models semantic properties, and to risk obscuring just how much properly in the domain of semantics remains to be done.

Chomsky has stressed such points especially in connection with names and so-called natural kind terms, expressions whose meanings have been thought to be 'directly referential' in the philosophical literature since the emergence of externalist views of content. We could add that analogous points could be made in connection

³⁸ Chomsky appears to agree: they reflect idealizations which may be "entirely reasonable" since, after all, "no one seeks to study everything" (2003: 290).

with demonstratives, an important source of ‘direct reference’ talk in philosophy. These expressions give rise to all the same phenomena Chomsky highlights for determiner phrases in general, and exhibit some further curious features. A sample, for illustrative purposes:

As noted by Mikkelsen (2004), one cannot specify who Peter is by using *Peter is this* or *Peter is that*, and bare demonstratives cannot refer to persons in object position (compare *I’d like to take you to him/#this/#that*). Bare demonstratives in object position can be used to predicate something of Peter, but apparently *that* is generally preferred for this use (as in *We need a good programmer, and Peter is that/#this*). If Peter is glimpsed leaving, we can say *That was Peter*, but not *This was Peter*. Using the future tense, both demonstratives are sometimes fine (answering the door, one can say *This/that will be Mary*), and sometimes not (e.g., as a movie begins, we can say *This will be great*, but not *That will be great*—though the latter is fine if made in reference to one of the previews rather than the feature presentation). Shifting from agents to abilities and events, other problems emerge. Attempting but failing to mount a unicycle, (17) is preferred to (18):

(17) I don’t know how to do this.

(18) I don’t know how to do that.

Now on the sidelines observing a unicyclist, the reverse is the case. What is demonstrated here is obviously not any object; rather, ‘unicycle-riding’ is what is under discussion. Yet the demonstrative seems able to sustain ‘dependent reference’ for an event variable, as in:

(19) I know I can do this—I did it easily yesterday.

Like Chomsky’s examples, the relevant dependence here is not literal sameness of reference: the event (if it is an event) corresponding to *it* occurred yesterday, but the *this* it is anaphoric with clearly is not referring to that past event. Neither is it referring to an event (token or type) instantiated within the context of utterance, since it concerns successful unicycle riding. Rather, if it picks out an event at all, it seems to pick out a merely possible event, one occurring in one of the possibilities compatible with my abilities, relative to each of my epistemic alternatives. But it is not obvious that reference is the appropriate notion for analyzing the demonstrative and its anaphor here; one might alternatively argue that they are each semantically contributing properties, and playing a primarily predicational role (compare Mikkelsen 2004).

Of course, the possibility of demonstrating something that “isn’t there” extends to material objects. Pointing to the spot where the last puzzle piece goes, we have no problem understanding:

(20) The dog ate that piece.

(21) That piece was never made—manufacturer error.

Turning to complex demonstratives, further questions emerge. Complex demonstratives can take bound readings, as noted by King (2001), though the constraints on when this is felicitous are not well understood. Demonstratives do not like to combine with predicates requiring uniqueness (*Peter is that tall/#tallest guy in the class*). Curiously however, as noted by Nowak (2014), moving the superlative to a relative clause (as in *Peter is that guy who is tallest in the class*) fixes it. As Nowak shows, explaining this simple distributional fact is not trivial.

Such data, which appear to involve interactions between the perspectives of the interlocutors, tense, aspect, modality, argument structure, information structure, and syntax, seem to be the tip of the iceberg.³⁹ They manifestly involve the semantic properties of demonstratives. The formal system of Kaplan (1978, 1977/1989)—which, strictly speaking, doesn't actually contain any semantics for demonstratives (instead, 'dthat')—does not illuminate these dimensions of meaning and use. This is not to suggest Kaplan's formal system fails to contribute to understanding—with suitable further assumptions, it correctly predicts certain basic kinds of context-sensitivity and certain limited interactions between demonstratives and intensional environments, valuable contributions.⁴⁰ But this should not obscure the facts that (i) a great many of the semantic features of demonstratives are yet to be accounted for on that system; (ii) some of their features (such as their bound readings) are straightforwardly incompatible with that system (for further discussion, see Nowak 2014).

Further, stepping back: the data do not appear to justify the introduction of a technical notion of 'direct reference', understood as going beyond semantic value. Although some may "feel initially that in the use of a true demonstrative, not only is one trying to put the object itself into the proposition (direct reference), but that the connection between the demonstrative and the object, call this *reference*, is also extraordinarily direct . . ." (Kaplan 1989: 573), a closer look at the data disconfirms this impression; indeed, it apparently supports the opposite conclusion. There seems to be little reason to question Chomsky's observation that "the semantic properties of [such expressions] focus attention on selected aspects of the world as it is taken to be by other cognitive systems, and provide intricate and highly specialized perspectives from which to view them, crucially involving human interests and concerns even in the simplest cases" (Chomsky 2000: 125).

We seem far from a clear account of the sort of data reviewed in this section. The upshot is not that this constitutes some devastating objection to model-theoretic

³⁹ See Fillmore (1997) (a collection of lectures from 1971) for additional problematic data, including further subtle interactions with tense. See also Partee 2000; King 2001; Mikkelsen 2004; Elbourne 2008; Heller and Wolter 2008; Wolter 2009; Sherman 2009; Kayne and Pollock 2010; Moltmann 2013; and Stojanovic 2013.

⁴⁰ Though the broader significance of these results has been misconstrued. The working features of Kaplan's semantics do not license a character-content distinction, nor a prohibition on monsters. See Rabern (2012b) for discussion.

approaches in natural language semantics. Nothing we have said shows that these data are not within the ken of such approaches. Rather, the lesson is simply that much really does remain to be explained—that failure to be puzzled is a mistake (compare Chomsky 2011: 277)—and that we should avoid construals of the model-theoretic approach which tend obscure this fact.

8 Semantics as the road to metasemantics

If (10) and (11) approach two of the main questions to be addressed in metasemantic inquiry from the viewpoint of present semantics, much remains to be clarified. First, it is not obvious that we can reasonably hope for anything like reduction, as opposed merely to some kind of intelligible unification, with other sciences. Second, even if reduction is possible, it is unclear just what sort of facts semantic value facts are to be reduced to: it is unclear what the relevant more fundamental science is. Third, unification or reduction may be possible only after adjustments to the more fundamental science. Fourth and perhaps most obvious, we do not know whether we have hit upon the appropriate categories to seek to unify. Future inquiry may yet dissolve our present technical notion of semantic value. Natural language semantics is in an early stage of development, with fruitful but ongoing debate on fairly elementary questions concerning the character of the semantic properties of most fragments of discourse. Direct pursuit of questions of the form of (10) and (11) is not obviously sensible in the current state of understanding. At present, the questions of metasemantics seem best pursued simply by doing more first-order semantics, thereby further clarifying the structure and nature of the properties we hope, eventually, to reduce, or intelligibly unify, with nonsemantic properties of some other sort. Premature emphasis on reduction or unification in the context of an emerging science tends to be counter-productive, at least if history is any guide.

The idea that philosophers can pursue metasemantic inquiry in some way largely autonomous from ongoing debates in ‘descriptive’ semantics is a puzzling one. To suggest that first-order semantics is the sensible route into metasemantic questions is not at all to deny their interest; rather, it is just to acknowledge that metasemantic inquiry is not discontinuous from natural language semantics proper, and that the latter is in an early stage of development.⁴¹

⁴¹ For helpful conversations on the topics of this chapter, I am indebted to Joshua Armstrong, Cian Dorr, Andy Egan, Anil Gupta, Wes Holliday, Jeff King, Richard Lawrence, John MacFarlane, Dilip Ninan, Ethan Nowak, Brian Rabern, François Récanati, Michael Rieppel, Daniel Rothschild, Jeff Russell, Philippe Schlenker, James Shaw, Robert Stalnaker, Isidora Stojanovic, Mark Wilson, and audiences at the University of Pittsburgh, Oxford, Rutgers, the University of Barcelona, Arché, and the Institut Jean Nicod. Thanks especially to Alexis Burgess, Alejandro Pérez Carballo and Brett Sherman for detailed comments on earlier drafts.

References

- Bach, K. (1994). "Conversational Implicature". *Mind & Language* 9(2): 124–62.
- Baldwin, D.A. (1991). "Infants' Contribution to the Achievement of Joint Reference". *Child Development* 62(5): 874–90.
- Bezuidenhout, A. (2002). "Truth-conditional Pragmatics". *Nous* 36: 105–34.
- Bloom, P. (2002). *How Children Learn the Meanings of Words*. Cambridge, MA: MIT Press.
- Braddon-Mitchell, D. and Jackson, F. (2007). *The Philosophy of Mind and Cognition: An Introduction*. 2nd edn. Malden, MA: Blackwell Publishers.
- Brandom, R. (1998). *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.
- Burge, T. (1979). "Individualism and the Mental". In P.A. French, T. Euling, and H. Wettstein (eds), *Midwest Studies in Philosophy*. Vol. 4. Minneapolis, MN: University of Minnesota Press, pp. 73–121.
- Burgess, J. (2008). "Tarski's Tort". In J. Burgess, *Mathematics, Models, and Modality: Selected Philosophical Essays*. Cambridge: Cambridge University Press, pp. 149–68.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Malden, MA: Wiley-Blackwell.
- Chomsky, N. (1956/1975). *The Logical Structure of Linguistic Theory*. Cambridge, MA: Harvard Libraries and Chicago: Chicago University Press. Available in microfilm at Harvard Libraries in 1956. Published in book form by Chicago University Press in 1975.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (1968a). "Linguistics and Philosophy". Talk delivered at New York University. Reprinted in N. Chomsky, *Language and Mind*. Cambridge: Cambridge University Press, 2006, pp. 143–72.
- Chomsky, N. (1968b). "Quine's Empirical Assumptions". *Synthese* 21: 53–68.
- Chomsky, N. (1980). *Rules and Representations*. New York: Columbia University Press, 1980.
- Chomsky, N. (1986). *Knowledge of Language: Its Nature, Origin, and Use*. Westport, CT: Praeger.
- Chomsky, N. (1992). "Explaining Language Use". *Philosophical Topics* 20: 205–31.
- Chomsky, N. (1995). "Language and Nature". *Mind* 104(413): 1–61.
- Chomsky, N. (2000). *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Chomsky, N. (2003). "Replies". In LM. Antony and N. Hornstein (eds), *Chomsky and His Critics*. Malden, MA: Blackwell Publishing, pp. 255–328.
- Chomsky, N. (2006). *Language and Mind*. Cambridge: Cambridge University Press.
- Chomsky, N. (2011). "Language and Other Cognitive Systems. What is Special About Language?". *Language Learning and Development* 7(4): 263–78.
- Collins, J. (2006). "Between a Rock and a Hard Place: A Dialogue on the Philosophy and Methodology of Generative Linguistics". *Croatian Journal of Philosophy* 18: 469–503.
- Collins, J. (2007a). "Linguistic Competence Without Knowledge of Language". *Philosophy Compass* 2(6): 880–95.
- Collins, J. (2007b). "Syntax, More or Less". *Mind* 116(464): 805–50.
- Cumming, S. (2008). "Variabilism". *The Philosophical Review* 117(4): 525.
- Davidson, D. (1965). "Theories of Meaning and Learnable Languages". In Y. Bar-Hillel (ed.), *Proceedings of the Third International Congress for Logic, Methodology and Philosophy of Science*. Vol. 3. Amsterdam: North-Holland Publishing Company, pp. 3–17.

- Davidson, D. (1967). "Truth and Meaning". *Synthese* 17: 304–23.
- Davidson, D. (1984). *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
- Devitt, M. (2006). *Ignorance of Language*. Oxford: Clarendon Press.
- Devitt, M. and Sterelny, K. (1999). *Language and Reality: An Introduction to the Philosophy of Language*. 2nd edn. Cambridge, MA: MIT Press.
- Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: MIT Press.
- Dummett, M. (1973). *Frege: Philosophy of Language*. 2nd edn. Cambridge, MA: Harvard University Press.
- Dummett, M. (1993). *The Logical Basis of Metaphysics*. Cambridge, MA: Harvard University Press.
- Elbourne, P. (2008). "Demonstratives as Individual Concepts". *Linguistics and Philosophy* 31(4): 409–66.
- Fillmore, C.J. (1997). *Lectures on Deixis*. Stanford, CA: CSLI Publications.
- Fine, K. (2001). "The Question of Realism". *Philosophers' Imprint* 1(1): 1–30.
- Fine, K. (2012). "Guide to Ground". In F Correia and B Schnieder (eds), *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge University Press, pp. 37–80.
- Fodor, J. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.
- Fox, D. (2000). *Economy and Semantic Interpretation*. Cambridge, MA: MIT Press.
- Groenendijk, J. and Stokhof, M. (1991). "Two Theories of Dynamic Semantics". In *Logics in AI* (Lectures Notes in Computer Science Series. Vol. 478). Berlin: Springer, pp. 55–64.
- Heim, I. (1982). "The Semantics of Definite and Indefinite Noun Phrases". PhD Thesis, University of Massachusetts.
- Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*. Armhrest, MA: Blackwell.
- Heller, D. and Wolter, L. (2008). "That is Rosa: Identificational Sentences as Intensional Predication". In *Proceedings of Sinn und Bedeutung*, vol 12. Oslo: Department of Literature, Area Studies and European Languages, University of Oslo, pp. 226–40.
- Jackendoff, R. (2002). *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: Oxford University Press.
- Kaplan, D. (1977/1989). "Demonstratives". In J. Almog, J. Perry, and H. Wettstein (eds), *Themes from Kaplan*. Oxford: Oxford University Press, pp. 481–563.
- Kaplan, D. (1978). "Dthat". *Syntax and Semantics* 9: 221–43.
- Kaplan, D. (1989). "Afterthoughts". In J. Almog, J. Perry, and H. Wettstein (eds), *Themes from Kaplan*. Oxford: Oxford University Press, pp. 565–614.
- Katz, J. (1981). *Language and Other Abstract Objects*. Totowa, NJ: Rowman and Littlefield.
- Kayne, R.S. and Pollock, J.-Y. (2010). "Notes on French and English Demonstratives". In J.-W. Zwart and M. de Vries (eds), *Structure Preserved: Studies in Syntax for Jan Koster*. Amsterdam: Linguistik Aktuell/Linguistics Today, pp. 215–28.
- Kennedy, C. and Stanley, J. (2009). "On Average". *Mind* 118(471): 583–646.
- Kepser, S. and Reis, M. (eds). (2005). *Linguistic Evidence: Empirical, Theoretical and Computational Perspectives* (Studies in Generative Grammar Series. Vol. 85). Berlin: De Gruyter Mouton.
- King, J. (2001). *Complex Demonstratives*. Cambridge, MA: MIT Press.
- Klinedinst, N. and Rothschild, D. (2012). "Connectives Without Truth Tables". *Natural Language Semantics* 20(2): 137–75.

- Kolodny, N. and MacFarlane, J. (2010). "Ifs and Oughts". *Journal of Philosophy* 107(3): 115–43.
- Kratzer, A. (2012). *Modals and Conditionals: New and Revised Perspectives*. Oxford: Oxford University Press.
- Kripke, S. (1972/1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Larson, R.K. and Segal, G. (1995). *Knowledge of Meaning*. Cambridge, MA: MIT Press.
- Lewis, D.K. (1970). "General Semantics". *Synthese* 22: 18–67.
- Lewis, D.K. (1975). "Languages and Language". In K. Gunderson (ed.), *Language, Mind, and Knowledge* (Minnesota Studies in the Philosophy of Science Series. Vol. 7). Minneapolis, MN: University of Minnesota Press, pp. 3–35.
- Lewis, D.K. (1980). "Index, Context, and Content". In S. Kanger and S. Ohman (eds), *Philosophy and Grammar*. Amsterdam: D. Reidel, pp. 79–100.
- Lewis, D.K. (1983). "New Work for a Theory of Universals". *Australasian Journal of Philosophy* 61(4): 343–77.
- Lewis, D.K. (1992). "Meaning without use: Reply to Hawthorne". *Australasian Journal of Philosophy* 70(1): 106–10.
- Lewis, D.K. (1994). "Reduction of Mind". In S. Guttenplan (ed.), *A Companion to the Philosophy of Mind*. Malden, MA: Blackwell, pp. 412–31.
- Ludlow, P. (2011). *The Philosophy of Generative Linguistics*. Oxford: Oxford University Press.
- MacFarlane, J. (2011). "Epistemic Modals are Assessment-Sensitive". In A. Egan and B. Weatherson (eds), *Epistemic Modality*. Oxford: Oxford University Press, pp. 144–78.
- Mikkelsen, L. (2004). "Specifying Who: On the Structure, Meaning, and Use of Specificational Copular Clauses". PhD thesis, University of California, Santa Cruz.
- Moltmann, F. (2013). "Identificational Sentences". *Natural Language Semantics* 21(1): 43–77.
- Montague, R. (1970). "English as a Formal Language". In B. Visentini et al. (eds), *Linguaggi nella societ a nella tecnica*. Milan: Edizioni de Comunita, pp. 189–224.
- Montague, R. (1974). *Formal Philosophy: Selected Papers of Richard Montague*. Ed. and Introd. by R.H. Thomason. New Haven, CT: Yale University Press.
- Morgan, J. (1970). "On the Criterion of NP Deletion". *Chicago Linguistics Society* 6: 380–389.
- Ninan, D. (2010a). "De Se Attitudes: Ascription and Communication". *Philosophy Compass* 5(7): 551–67.
- Ninan, D. (2010b). "Semantics and the Objects of Assertion". *Linguistics and Philosophy* 33(5): 355–80.
- Ninan, D. (2012a). "Counterfactual Attitudes and Multi-centered Worlds". *Semantics and Pragmatics* 5: 1–57.
- Ninan, D. (2012b). "Propositions, Semantic Values, and Rigidity". *Philosophical Studies*, 158(3): 401–13.
- Nowak, E. (2013). "Demonstratives are not Rigid Designators". Unpublished manuscript, University of California, Berkeley.
- Nowak, E. (2014). "Demonstratives without Rigidity or Ambiguity". *Linguistics & Philosophy*, forthcoming.
- Partee, B. (1975). "Montague Grammar and Transformational Grammar". *Linguistic Inquiry* 6(2): 203–300.
- Partee, B. (1987). "Noun Phrase Interpretation and Type-Shifting Principles". *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers* 8: 115–43.
- Partee, B. (2000). "Copula Inversion Puzzles in English and Russian". *Issues in Semantics* 23: 183–208.

- Pérez Carballo, A. (2014). "Semantic Hermeneutics". In A. Burgess and B. Sherman (eds), *Metasemantics*. Oxford: Oxford University Press. (Also Chapter 4, this volume.)
- Pietroski, P. (2003). "The Character of Natural Language Semantics". In A. Barber (ed.), *Epistemology of Language*. Oxford: Oxford University Press, pp. 217–56.
- Putnam, H. (1975). "The Meaning of 'Meaning'". In H. Putnam, *Philosophical Papers*. Vol. 2: *Mind, Language, and Reality*. Cambridge: Cambridge University Press.
- Rabern, B. (2012a). "Against the Identification of Assertoric Content with Compositional Value". *Synthese* 189(1): 75–96.
- Rabern, B. (2012b). "Monsters and Communication: The Semantics of Contextual Shifting and Sensitivity". PhD thesis, Australia National University, Canberra.
- Rabern, B. (2013). "Monsters in Kaplan's Logic of Demonstratives". *Philosophical Studies* 164(2): 393–403.
- Récánati, F. (1989). "The Pragmatics of What is Said". *Mind and Language* 4(4): 295–329.
- Récánati, F. (1993). *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- Récánati, F. (2010). *Truth-Conditional Pragmatics*. Oxford: Oxford University Press.
- Rosen, G. (2010). "Metaphysical Dependence: Grounding and Reduction". In B. Hale and A. Hoffmann (eds), *Modality: Metaphysics, Logic, and Epistemology*, Oxford: Oxford University Press, pp. 109–36.
- Rothschild, D. and Yalcin, S. (2012). "On the Dynamics of Conversation". Manuscript, All Souls College, Oxford, and the University of California, Berkeley.
- Sabbagh, M. and Baldwin, D. (2005). "Understanding the Role of Communicative Intentions in Word Learning". In N. Eilan et al. (eds), *Joint Attention: Communication and Other Minds*. Oxford: Oxford University Press, pp. 165–84.
- Schiffer, S. (2006). "Two Perspectives on Knowledge of Language". *Philosophical Issues* 16(1): 275–87.
- Schütze, C.T. (1996). *The Empirical Base of Linguistics: Grammaticality Judgments and Linguistic Methodology*. Chicago: University of Chicago Press.
- Schütze, C.T. (2010). "Linguistic Evidence and Grammatical Theory". *Wiley Interdisciplinary Reviews: Cognitive Science* 2(2): 206–21.
- Schütze, C.T. and Sprouse, J. (2014). "Judgment Data". In D. Sharma and R. Podesva (eds), *Research Methods in Linguistics*. Cambridge: Cambridge University Press, pp. 27–50.
- Sherman, B. (2009). "Indexicality". PhD thesis, Princeton University, Princeton, NJ.
- Sider, T. (2011). *Writing the Book of the World*. Oxford: Oxford University Press.
- Soames, S. (1989). "Direct Reference and Propositional Attitudes". In J. Almog, J. Perry, and H.K. Wettstein (eds), *Themes from Kaplan*. Oxford: Oxford University Press, pp. 394–419.
- Soames, S. (2008). "The Gap Between Meaning and Assertion". In S. Soames, *Philosophical Essays*, Vol. 1: *Natural Language: What It Means and How We Use It*. Princeton, NJ: Princeton University Press, pp. 278–97.
- Soames, S. (2010). *What is Meaning?* Princeton, NJ: Princeton University Press.
- Sperber, D. and Wilson, D. (1986). *Relevance: Communication and Cognition*. Cambridge, MA: Harvard University Press.
- Sprouse, J. and Almeida, D. (2012). "Assessing the Reliability of Textbook Data in Syntax: Adger's Core Syntax". *Journal of Linguistics* 48: 609–52.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Stalnaker, R. (1989). "On What's in the Head". *Philosophical Perspectives* 3: 287–319.

- Stalnaker, R. (1997). "Reference and Necessity". In B. Hale and C. Wright (eds), *A Companion to the Philosophy of Language*. Malden, MA: Blackwell, pp. 534–54.
- Stalnaker, R. (2008). *Our Knowledge of the Internal World*. Oxford: Oxford University Press.
- Stanley, J. (1997a). "Rigidity and Content". In R.G. Heck Jr. (ed.), *Language, Thought, and Logic: Essays in Honor of Michael Dummett*. New York: Oxford University Press, pp. 131–56.
- Stanley, J. (1997b). "Names and Rigid Designation". In B. Hale and C. Wright (eds), *A Companion to the Philosophy of Language*. Malden, MA: Blackwell, pp. 555–85.
- Stanley, J. (2000). "Context and Logical Form". *Linguistics and Philosophy* 23(4): 391–434.
- Stephenson, T.C. (2007). "Judge Dependence, Epistemic Modals, and Predicates of Personal Taste". *Linguistics and Philosophy* 30(4): 487–525.
- Stojanovic, I. (2012). "On Value-Attributions: Semantics and Beyond". *The Southern Journal of Philosophy* 50(4): 621–38.
- Stojanovic, I. (2013). "Prepragmatics: Widening the Semantics-pragmatics Boundary". In A. Burgess and B. Sherman (eds), *Metasemantics*. Oxford: Oxford University Press.
- Stokke, A. (2014). "Truth and Context Change". *Journal of Philosophical Logic* 43(1): 33–51.
- Van Rooij, R. (2005). "A Modal Analysis of Presupposition and Modal Subordination". *Journal of Semantics* 22(3): 281–305.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2006). "Can Cognition be Factorised into Internal and External Components?". In R. Stainton (ed.), *Contemporary Debates in Cognitive Science*. Malden, MA: Blackwell, pp. 291–306.
- Wilson, D. and Sperber, D. (2012). *Meaning and Relevance*. Cambridge: Cambridge University Press.
- Wolter, L. (2009). "Demonstratives in Philosophy and Linguistics". *Philosophy Compass* 4(3): 451–68.
- Yalcin, S. (2007). "Epistemic Modals". *Mind* 116(464): 983–1026.
- Yalcin, S. (2011). "Nonfactualism About Epistemic Modality". In A. Egan and B. Weatherson (eds), *Epistemic Modality*. Oxford: Oxford University Press, pp. 295–332.
- Yalcin, S. (2012a). "Bayesian Expressivism". *Proceedings of the Aristotelian Society* 112(2): 123–60.
- Yalcin, S. (2012b). "Context Probabilism". In M. Aloni et al. (eds), *Proceedings of the Eighteenth Amsterdam Colloquium* (Lecture Notes in Computer Science Series. Vol. 7218). Amsterdam: Springer, pp. 12–21.
- Yang, C. (2002). *Knowledge and Learning in Natural Language*. Oxford: Oxford University Press.