Explanation and Partiality in Semantic Theory*

Michael Glanzberg
Northwestern University

In this paper, I shall argue for a form of partiality in semantics. In particular, I shall argue that semantics, narrowly construed as part of our linguistic competence, is only a partial determinant of content. Likewise, semantic theories in linguistics function as partial theories of content. I shall go on to offer an account of where and how this partiality arises, which focuses on how lexical meaning combines elements of distinctively linguistic competence with elements from our broader cognitive resources. This account shows how we can accommodate some partiality in semantic theories without falling into skepticism about semantics or its place in linguistic theory.

In recent years, there have been a number of challenges to semantics. For instance, Chomsky (e.g. Chomsky, 2000) has in effect argued that semantics is not an aspect of linguistic competence, and so, is not on par with syntax or phonology (cf. Pietroski, 2005b). Also, a number of views sometimes grouped together as 'radical contextualism' have sought to shift much of the burden of explaining the contents of utterances from semantics onto pragmatics (e.g. Carston, 2002; Recanati, 2004, 2010; Sperber and Wilson, 1986; Travis, 1996).

*Earlier versions of this paper were presented at a number of places, including the Institute of Philosophy at the University of London, The Ohio State University, the Society for Exact Philosophy Annual Meeting in Kansas City, Tel Aviv University, the University of Connecticut, the University of Michigan, and the University of Toronto. I am grateful to all those audiences for helpful questions and stimulating discussions, which kept this paper changing over a long period of time. I am also very grateful to Jc Beall, Harvey Friedman, Phil Kremer, Michael Lynch, Mohan Matthen, Gurpreet Rattan, Craige Roberts, Kevin Scharp, and Stewart Shapiro for helpful discussions of this paper and the issues in it. Paul Pietroski's work, and conversations with him, were especially important to the development of the views I present here. Thanks also to Alexis Burgess, Sam Cumming, Magdalena Kaufmann, and Karen Lewis for insightful comments on earlier drafts of this paper, and again to Alexis and to Brett Sherman for support (and patience!) as the paper was being completed.

Such views consign semantics to a small and unexciting role, or no role at all. They cast doubt on the importance of semantic theory within the broader enterprise of linguistic theory, and they give it little to do in explaining how what we say gets its content. In contrast, the view I shall advance here is not skeptical about semantics or its role in linguistic theory. My view preserves an important place for semantics in linguistic competence, and in accounting for content. Indeed, recognizing the way semantics can be a partial determinant of content, I shall suggest, gives us a better appreciation of the robust role semantics plays within linguistic competence, and how there is a place for a substantial semantic theory within linguistic theory. A little modesty about the domain of semantic theory will give us a better understanding of what it does well. Such modesty, not skepticism, is what I shall propose here.

My argument in favor of partiality will proceed by examining where our semantic theories provide good explanations. I shall argue that semantic theories can and do offer good explanations, but they also typically lose their explanatory force at certain points. What this shows, I shall maintain, is that semantic theories, construed as theories of semantic competence, are indeed substantial theories, but they are only partial theories of semantic content. We may then conclude that semantic competence itself must be only a partial determinant of content. Semantics is thus partial. I show how we may make sense of this situation, by sketching a picture of lexical meaning that includes both elements that are genuinely part of linguistic competence and pointers to elements from our broader conceptual resources but outside of linguistic competence. The contents of our words and sentences are only provided by the combination of both. The partiality of semantics, I shall propose, is the result of meaning including such pointers, and thereby going beyond linguistic competence proper.

To show where semantic theories lose their explanatory force, I shall return to an old debate about the roles of model theory and disquotation in semantics. I shall argue that, as least as far as our current grip on lexical semantics goes, both have a place in our theories. (Thus, in effect, I shall defuse the old debate.) But, I shall claim, model theory provides one illustration of where and how semantic theories can provide good explanations of semantic competence. (Other applications of mathematics in semantics do as well.) In contrast, the use of disquotation in semantic theories precisely marks the places where they lose their explanatory force. Insofar as disquotation plays an ineliminable role in building theories of content, semantic theories can be at best partial theories of content. From there, as I said, we can conclude

that semantics is itself partial. Moreover, I shall propose, disquotation is a guide to where linguistic meaning contains pointers to extra-linguistic elements of content. It is thus no surprise that theories of linguistic competence fail to offer substantial explanations at those places.

The plan of this paper is as follows. In section 1, I shall revisit the old debate about the roles of model theory and disquotation in semantics, and argue that approaches that use either are still engaged in what is fundamentally the same project. With that background, I shall turn to the question of what model theory and disquotation do in semantic theories in section 2. I shall argue there that disquotation plays an ineliminable role, but that this shows us a way in which semantic theories are explanatorily partial. In section 3, I shall argue that this indicates a particular form of partiality in the subject-matter of semantics, which is generated by lexical entries including pointers to extra-linguistic concepts.

1 Approaches to Truth-Conditional Semantics

In this section, I shall set up the background we will need for our discussion, and revisit an old debate about the way semantic theories should be formulated, that pitted followers of Davidson (e.g. Davidson, 1967) against those of Montague (e.g. Montague, 1973). To set the stage for later sections, I shall here attempt to defuse that old debate, by arguing that at least when it comes to their current incarnations, the projects these two approaches are engaged in are not fundamentally different. This will allow us to revisit what roles disquotation (the hallmark of Davidsonian semantics) and model theory (the hallmark of Montagovian semantics) play in current semantic theories, starting in section 2. That in turn will set the stage for the discussion of partiality to come in section 3.

Before looking at the old debate, we should start by reviewing some basic assumptions that will guide the discussion to follow. First, our starting point will be linguistic competence. I shall assume that linguistic competence is a distinct aspect of our cognitive organization. This can be made most vivid if we assume, with Chomsky (e.g. Chomsky, 1965, 1980, 1986) that our cognitive architecture includes a language faculty. Not all of parts of Chomsky's view are required for the arguments to follow; indeed, any assumptions that

imply that linguistic competence is distinct from other cognitive abilities somehow will suffice. This will be most important in section 3, and I shall pause there to ask what weaker assumptions might suffice, but for the most part, I shall simply adopt a broadly Chomskian view and assume there is a language faculty.[1]

I shall also assume, as is generally assumed in the tradition of generative linguistics, that linguistic competence is the primary subject of study for linguistic theory. Thus, in effect, linguistic theory studies what knowledge—or more generally what cognitive states—underwrite our linguistic abilities. If we assume a distinct language faculty is responsible for our linguistic competence, then the language faculty is the primary object of study for linguistic theory. I shall thus use 'linguistic theory' to mean the theory that studies linguistic competence, presumed to be provided by a language faculty. This is somewhat stipulative. Though it has proved a useful way to approach syntax and phonology, and I shall in effect argue here, semantics, it does not really require us to ignore everything about performance systems, or other aspects of cognition related to linguistic abilities. It is just a way to remind ourselves what our primary focus is.

When we talk about what is provided by the language faculty (or whatever else might be responsible for our linguistic competence), we often talk about it in terms of what speakers know in virtue of knowing their languages. But, this way of talking is always accompanied by various provisos. At best, such knowledge is often highly tacit. More importantly, the cognitive state one is in in virtue of having a language faculty may not be a knowledge state at all.[2] Our assumptions about the language faculty guarantee there is some cognitive state or another that amounts to linguistic competence, but do not say just what that state is. Thus, 'knowledge' is used here as a place-holder for whatever the right cognitive state turns out to be. I shall continue to call that state knowledge, but the terminology should be taken with so many grains of salt.

Where does semantics fit into the project of studying linguistic competence? It studies what, in virtue of linguistic competence, we know about what our sentences, words, or phrases mean; i.e. what speakers know when they understand the words and sentences of their languages. A leading idea for how to pursue this project is to give truth conditions a central place. Part

[1]See also the discussion of Collins (2004).

[2]See Pettit (2002) for an argument that linguistic competence cannot be knowledge.

of what it is to understand a sentence, it is proposed, is to understand its truth conditions. So, at least, a semantic theory should account for knowledge of truth conditions. It may well need to do more, but it must do at least this much. Looking at subsentential constituents, we expect their contributions to truth conditions to be determined by facts about reference and satisfaction, among things. Thus, semantic competence amounts to knowledge of such properties as truth conditions, reference, and satisfaction.[3] This is a widely held view both in linguistics and philosophy, and it is one I shall take for granted here.[4] This sort of assumption is non-trivial. So, for instance, we are not exploring conceptual role or inferentialist semantics, cognitive semantics, or many others in the myriad of approaches to semantics the years have seen. But, the assumption is made widely enough to be warranted. More importantly, my focus in this paper is on how to understand the kinds of theories such an assumption leads to; hence, it is a starting-point for our discussion, rather than a conclusion. As we proceed, we will examine how much of knowledge of truth conditions really can be part of linguistic competence, but we will assume that if there is any job for semantics to do in linguistic theory, it must include articulating aspects of knowledge of truth conditions.[5]

With all this background in place, we can finally begin to approach the main issue of this paper: whether semantic competence, i.e. knowledge of truth conditions provided by the language faculty, suffices to determine truth-conditional content. Recall, I have promised to argue it only does so partially, and to do so by examining how semantic theories work, and where they offer good explanations. With that in mind, we should start by asking what a semantic theory should look like, given the assumptions we have just reviewed.

[3]There is thus a mild terminological issue. Sometimes 'semantics' is taken to mean facts about truth, reference, and satisfaction, and sometimes speakers' knowledge of them. From the perspective this paper is taking, the main object of study is speakers' knowledge, and we are concerned with additional facts about truth and reference where it helps us to understand speakers' cognitive states.

[4]Classic discussions of the place of truth conditions in semantics include Carnap (1947), Cresswell (1973), Davidson (1967), Lewis (1970), and Montague (1970). Discussions focusing on semantic competence include Higginbotham (1989b, 1992), Partee (1979), and Segal (2006). An opposing view is articulated by Soames (1992).

[5]Thus, at least for the moment, I am siding with Larson and Segal (1995) over Chomsky (2000) in assuming there will be some semantics in the language faculty. But, the main point of this paper is to ask how much, and in effect argue there is some, though a limited amount.

We will see that there are two competing ways of formulating such theories, and examining how they work will be our first step towards seeing where semantic theories provide good explanations.

What should a semantic theory do? With all our assumptions in place, it becomes entirely natural to assume that a semantic theory will produce statements of the truth conditions of sentences. We can expect more, of course. Such statements should be derived compositionally, and from the point of view we are adopting, we should expect the components of the theory to reflect the components of speakers' semantic competence. It should thus tell us what speakers know about meaning by telling us how truth conditions are determined from what they know. But when we write the theory down, we will be writing down truth conditions (and what determines them). Our starting point is thus the idea that statements of truth conditions for sentences should follow from a good semantic theory. As we proceed, we will ask whether this assumption really holds up, but it is natural place to start, and it is where traditionally semantic theories have started.

Indeed, this is where they did start, but they did so in two different ways. Beginning in about the 1960s, two different ways we might state such truth conditions came into focus. One, in the Davidsonian tradition (e.g. Davidson, 1967), anticipates deriving T-sentences like:

(1) *Ernie is happy* is true $\longleftrightarrow$ Ernie is happy

The other, in the Montagovian tradition (e.g. Montague, 1973), anticipates deriving statements like:

(2) For any model $\mathfrak{M}$, *Ernie is happy* is true in $\mathfrak{M}$ $\longleftrightarrow$ $[\![\text{Ernie}]\!]_{\mathfrak{M}} \in [\![\text{happy}]\!]_{\mathfrak{M}}$

It should be familiar how both of these present the truth conditions of the target sentence. The first does it by *disquotation*. Statement (1) provides truth conditions because *Ernie is happy* is used on the right hand side of the biconditional to state them. The second defines a range of models that represents the truth conditions of the target sentence. Statement (2) tells us that *Ernie is happy* is true in just the models $\mathfrak{M}$ where the referent $[\![\text{Ernie}]\!]_{\mathfrak{M}}$ of *Ernie* in $\mathfrak{M}$ is in the extension $[\![\text{happy}]\!]_{\mathfrak{M}}$ of *happy* in $\mathfrak{M}$.[6]

[6]This use of models appears, for instance, in Dowty *et al.* (1981, pp. 41–47). Models can be invoked in semantics in a number of ways, though they generally imply that truth is only relative to a model, which is the crucial issue we will explore. For some other ways of invoking models, see Lasersohn (2000) and Zimmermann (1999). Interestingly, we

We should note, as it will become important below, that disquotation clauses like (1) do state non-trivial facts about the truth conditions of sentences. From the perspective we are taking here, they are facts knowledge of which will, we presume, form a non-trivial aspect of speakers' linguistic competence. Whether this makes disquotation adequate for various purposes is an old question, and one we will return to in section 2.

The two ways of providing truth conditions have become associated with two approaches to truth-conditional semantics. The first, as in (1), is the hallmark of semantics in the Davidsonian tradition (e.g. Davidson, 1967), represented currently by the textbook of Larson and Segal (1995). The second, as in (2), is the hallmark of semantics in the Montagovian tradition (e.g. Montague, 1973), represented by such textbooks as Chierchia and McConnell-Ginet (2000), Dowty *et al.* (1981), and Heim and Kratzer (1998).[7] The existence of so many textbooks is an indication that truth-conditional semantics as a research program is alive and well, and has two distinct flavors.[8]

To fix some terminology, let us refer to these two sorts of semantic theories by the names of their fathers: Davidsonian for the sort that provides truth conditions in the manner of (1), and Montagovian for the sort that provides them in the manner of (2). I mean these names to pick out the theories that have emerged over the years and are represented in the textbooks just mentioned, not the particular views of Davidson and Montague themselves. As we go forward, we will be reconsidering what is fundamental to these approaches, but for now, we will assume it is the two ways of stating truth conditions represented by (1) and (2).

It is clear enough that the two distinctive ways of stating truth conditions, Davidsonian and Montagovian, give two different flavors of truth-conditional

see what appears to be a different use of models mentioned earlier in Dowty *et al.* (1981, pp. 10-11). For a more general discussion of model theory in semantics, see Zimmermann (2011).

[7]This way of dividing up the landscape in semantics sets up a conflict between two programs, and as such, it is fairly common way for philosophers to frame a foundational issue. But, to anticipate what is to come in a moment, it is not completely fair to how semantics has proceeded since the early proposals of Montague and Davidson. In particular, as I shall make much of in a moment, not all of these textbooks in the Montagovian tradition invoke models in this particular way.

[8]The textbooks I mention place truth-conditional semantics broadly within the program of generative linguistics. As such, they depart in important ways from the views of their primogenitors Davidson and Montague themselves.

semantics. But often, proponents of these two varieties have seen them as incompatible, and not as two variants of the same basic idea. Indeed, proponents of each have from time to time argued that the other is defective, or at least inadequate as a semantic theory. To begin our exploration of disquotation, I shall focus on one particular instance of this sort of argument: one due to Lepore (1983) that claims that the Montagovian approach to truth-conditional semantics is inadequate. Seeing what we can learn from this argument will help us to better understand the place of disquotation in truth-conditional semantic theories, and set up the arguments for partiality to come in the subsequent sections.[9]

Lepore argues that the way the Montagovian variant of truth-conditional semantics provides truth conditions is inadequate because it only provides what he labels *relative* truth conditions, and those are inadequate to capture what speakers know about the meanings of their sentences. The point is this. A statement like (2) only tells us that whatever the value of *Ernie* is in a model, and whatever the value of *happy* is in the model, the sentence *Ernie is happy* is true if the former is in the latter. Thus, relative to a model, we get information about what makes the sentence true. But models are allowed to vary in what values they assign to expressions, and predicates like *happy* can vary quite a bit in their extensions. In this way, (2) only delivers truth conditions relative to a model, i.e. relative truth conditions.

The problem, as Lepore points out, is that you can know all the information about the relative truth conditions of a sentence, and not know what it means. Indeed, you can know its relative truth conditions, and not really know its truth conditions at all. All you know is whatever a model assigns to the subject, it falls within the extension of the predicate in that model. You can know that and not know what the sentence means. Indeed, you can know that and know nothing more than that the sentence in question has a certain grammatical form (cf. Higginbotham, 1989b). You know no more about a sentence in virtue of this than you know about *The mome raths outgrage* from Carroll's "Jabberwocky" (Carroll, 1960).

In virtue of being only relative to models, the conditions delivered by (2) are too weak to capture the truth conditions of a sentence. They do

[9]I also discussed this argument in my (forthcoming). Though my focus shall be on Lepore, I pause to note that the path I shall go down starting with his article follows in the footsteps of a number of other authors, including Cresswell (1978), Higginbotham (1989b), Pietroski (2005b, 2010), and Zimmermann (1999). The conclusions I reach in the end share a close affinity with those of Pietroski, in particular.

provide a range of models, but far too large a one to represent what speakers understand about the sentence. In contrast, (1) states what speakers know about the truth conditions of the sentence quite well; viz., that it is true just in case Ernie is happy. These statements are not relative, and give the truth conditions in terms of Ernie and happiness, unrelativized. Such *absolute* statements of truth conditions appear to capture speakers' understanding much better. We could go further, and note that it is absolute assignments of reference and satisfaction properties to the parts of a sentence, and not assignments of values relative to a model, that allow us to derive the right truth conditions.[10]

Lepore concluded that absolute semantics, carried out in the Davidsonian tradition, makes a substantial step towards modeling speakers' semantic competence, while model-theoretic semantics in the Montagovian tradition fails, as it offers only relative truth conditions (cf. Higginbotham, 1988). There are a number of responses to this claim that have emerged over the years, focusing on various ways one might restrict the class of model involved, which could yield more accurate representations of truth conditions. Lepore considers some options along these lines as well. I shall not pursue them here, as the line of response I shall concentrate on takes a different course.

Though I doubt there is complete consensus about these issues among proponents of Montagovian semantics, a fairly typical reply has emerged since the early work of Montague. The response is actually quite simple. Semantics in the Montagovian 'model-theoretic' tradition can be absolute, and most importantly, it can be just as absolute as semantics in the Davidsonian tradition.

For instance, when we look at the widely-used textbook of Heim and Kratzer (1998) we find things like:[11]

(3) a. $[\![\text{Ernie}]\!] = \text{Ernie}$
b. $[\![\text{happy}]\!] = \lambda x \in D_e.\ x \text{ happy}$

These are just as *absolute* as the sorts of disquotation clauses we find in Davidsonian theories. They make no reference to a model to which semantic

[10]The terminology of 'absolute' and 'relative' appears in Davidson (1973).

[11]The particular notation I am using here is fairly standard, and a slight simplification of the one in Heim and Kratzer. The semantic value $[\![\text{happy}]\!]$ for *happy* is a function from individuals to truth values, which returns true if the individual is happy. This is the familiar way to define extensions in the λ-calculus.

values are relativized; rather, they pick out individuals and extensions, just as Davidsonian theories do.[12]

So, to a great extent, semantics in the Montagovian tradition is just as absolute as that in the Davidsonian tradition. Thus, we can simply accept Lepore's points about absolute semantics. Interestingly, we see a tendency to opt for absolute semantics even in works much closer to Montague's original papers. For instance, in the first textbook presentation of Montague semantics (Dowty *et al.*, 1981), there is an extensive discussion of the model theory of intensional logic in the early parts of the book. But when the linguistic analysis starts to get really interesting, reference to models usually drops out.[13]

This is not to say that there are no differences between Montagovian and Davidsonian approaches to semantics. There are many. Perhaps most important is the use in Montagovian semantics of the typed λ-calculus, which is rich in higher-order resources (e.g. Heim and Kratzer, 1998; Klein and Sag, 1985; Montague, 1973). Davidsonians tend to opt for first-order resources (e.g. Higginbotham, 1998). One important consequence of this is that Montagovians typically see semantic composition as function/argument composition, often with higher-order functions, while Davidsonians opt for composition by conjunction, relying heavily on the apparatus of thematic roles (e.g. Higginbotham, 1985, 1989a; Pietroski, 2005a).

I should pause to mention one difference I do not think is of great significance. The use of the λ-calculus allows for all constituents to be assigned objects as semantic values. Davidsonians typically do not do this for predicates or other phrases that would require higher-order elements. Instead, they state satisfaction conditions for such elements without providing objects to be semantic values. In my view, this is significant only insofar as it is

[12]We find something similar in the textbook of Chierchia and McConnell-Ginet (2000), especially in the early chapters. However, the situation there is somewhat more complicated. They do introduce models, in the context of explaining how quantifiers and variable assignments work. But in a telling (if perhaps off-hand) passage (p. 125), they note that when models are in play, we get truth conditions only after a model is fixed. That takes us most of the way to absolute truth conditions. They also rely on models in explaining how meaning postulates work, but my suspicion is that otherwise, models do relatively little work for them. For a critical discussion of meaning postulates and their relation to model theory, see Zimmermann (1999).

[13]To some extent this is an artifact of the Montagovian preference for translating first into the language of intensional logic. But, it is striking that the substance of the analyses does not invoke models even if they are working in the background.

a reflection of the issues I just mentioned. It often leads to distinct analyses, but it is not of great importance itself. I see no important difference between saying either of:[14]

(4) a. $[\![\text{smokes}]\!] = \lambda x \in D_e.\ x \text{ smokes}$
b. $Val(x, \text{smokes}) \longleftrightarrow x \text{ smokes}$

These do differ in what ontological demands are placed on their metatheories, but they in effect describe the way *smokes* contributes to truth conditions in the same way, by telling us it is a predicate that applies to things that smoke. They use different notation, and will figure in different theories, but they basically tell us the same thing, and attribute the same knowledge to speakers.

There are any number of other similarities and differences between Montagovian and Davidsonian approaches to semantics. They differ on how they approach intensionality, for instance, where Montagovians can make ready use of possible worlds, and Davidsonians typically look for other analyses.[15] My point is that in their basic goal of providing truth conditions, as part of an explanation of semantic competence, they are both doing essentially the same thing, and the Montagovian approach is not committed to giving relative truth conditions.

If the Davidsonian approach to semantics, relying on disquotational statements like (1), and the Montagovian approach, relying on the λ-calculus as illustrated in (3b), are basically in the same business of stating truth conditions as part of the project of describing semantic competence, does that mean there is nothing special about disquotation? I shall argue in the next section that there is something special about disquotation, though not what the debate over absolute semantics suggested. Seeing what is distinctive about it will help us to isolate where semantic theories provide substantial explanations, which in turn, I have promised, will reveal a way in which semantics can be partial but still a substantial component of a linguistic competence, as we will explore in section 3.

[14] For (4b) I follow the form used by Larson and Segal (1995).

[15] Though, I myself am not so sure they even differ all that much on this point, as I argued in my (2009).

2 Disquotation and Explanation

If it is right, as I argued in the preceding section, that Montagovian and Davidsonian semantic theories are on par in their goal of stating truth conditions, as part of a broader theory of semantic competence, then we face a new set of worries. Davidsonian theories, relying on disquotation as we see in (1), have raised questions about how such seemingly trivial statements can play a significant role in a theory. Though I noted above that disquotational statements are not genuinely trivial, I shall argue here that they threaten the explanatory value of theories, which in turn raises questions about how and whether they are really describing substantial aspects of linguistic competence.

In this section, I shall argue that disquotation really does reveal a species of partiality in the explanatory power of semantic theories, and I shall contrast this with ways that semantic theories can and do provide good explanations. Good explanations tend to appear where we apply model theory or other branches of mathematics to semantics, while mere disquotation signals explanatory weakness. The results of section 1 tell us that this is not a matter of any fundamental difference in how the two approaches provide truth conditions, but it is an important difference in where and how theories provide good explanations. The gap between explanatorily fruitful and weak aspects of semantic theories will show the way semantic theories are partial in explanatory power, as I shall argue in this section. This in turn will be the starting point for the argument in section 3 that there is a kind of partiality in the subject-matter of semantic theories.

Let us return to the sorts of disquotational statements, like (1), that are the hallmark of Davidsonian semantics:

(5) a. *Ernie is happy* is true $\longleftrightarrow$ Ernie is happy
b. *Ernie* refers to Ernie

It is a commonplace observation that these state non-trivial facts about truth conditions, or contributions to truth conditions. Indeed, that was a significant part of the argument from Lepore we considered above. But nonetheless, there is something seemingly weak about statements of these forms. Their weakness is nicely illustrated by the observation that they can be generated with only minimal knowledge of the grammatical category an expression falls under, so long as the metalanguage in which the theory is being given includes the object language (Higginbotham, 1989b).

We thus see disquotation statements as simultaneously non-trivial, but somehow weak. To better understand this, it will be useful to return to a classic discussion of related issues in the theory of truth, from Field (1972). In considering the value of a Tarskian theory of truth, which of course contains similar sorts of disquotation clauses to the ones we are considering, Field directs our attention to what kinds of explanations theories offer.[16] Disquotation clauses, Field points out, in effect *list* substantial facts about reference or truth conditions. For instance, the disquotation clauses for referring expressions amount to merely a list of their referents.

The problem with such lists, as Field pointed out, is that they often fail to provide good explanations of underlying phenomena. Field asks us to contrast two theories of valence (from chemistry): one which simply lists the valences of various elements, and the current quantum theory which explains how elements get their valences. The former does state non-trivial facts of chemistry, but in a way that fails to offer any kind of substantial explanation of the phenomena. (Field himself was more concerned with physicalist reduction, but I take it that would amount to a particular sort of explanation.) In practice, a list of values often tells us no more than that things have the values they have. A simple list of valences for elements helps us little more than just knowing that each element has whatever valence it has. Likewise, disquotational clauses offer little explanation beyond pointing out that expressions make whatever contributions to truth conditions they do.

When it comes to explaining semantic properties of expressions, we see this lack of explanation in the ease with which disquotation clauses can be generated. If you know the grammatical category of an expression, you can indicate that it has the semantic property that it does, and report that disquotationally. Of course, if you understand the object language, you will know what that semantic property is, and so, you can glean more from the disquotation clause. But, you gain nothing that was not already transparent to you as a speaker of the language. You thus fail to learn any non-trivial generalizations, make any non-trivial predictions, or do anything else that might figure into offering explanations in semantics. Just like lists, disquo-

[16] As has been much-discussed, somewhat different assumptions must be made to use the Tarskian apparatus as part of an empirical theory in semantics, or as part of a definition of truth. See Etchemendy (1988) and Soames (1984). I should also mention that the issue here is the status of disquotation, not the perhaps related question about whether deflationism about truth is incompatible with the kind of truth-conditional semantics we are exploring here. For this issue, see Burgess (2011), Collins (2002), and Horisk (2007).

tation fails to offer good explanations. To sum this up, we might say that though disquotation states non-trivial semantic facts, it is still *boring*!

One consequence of our examination of the role of model theory in semantics in section 1 is that Montagovian semantics is just as vulnerable to this sort of concern as the Davidsonian variety. The distinctive feature of Montagovian semantics, we saw above, is not so much the use of models as the use of higher-type objects from the λ-calculus as semantic values. But the problems with merely listing semantic facts can appear in that setting just as easily. There is no fundamental difference in explanatory power between:

(6) a. $[\![\text{smokes}]\!] = \lambda x \in D_e.\ x$ smokes
b. $Val(x, \text{smokes}) \longleftrightarrow x$ smokes

Each has the properties that make disquotation of limited explanatory value, as they fail to explain anything about the meaning of *smokes* that was not already transparent to a speaker beyond the grammatical information that is indicated by the type assignment or use of Val. Semantics of either flavor we have considered is vulnerable to the concerns about explanatory value Field raised for disquotation.

Reflecting on Field's point shows a genuine worry about semantic theories. To the extent that they wind up appealing to disquotation, they threaten to become explanatorily vacuous. They may be true, and non-trivial in some respects, but that is still a failing for any empirical theory. The worry applies, I have been arguing, equally to Montagovian or Davidsonian approaches to semantics. But, combined with Lepore's observation that semantic theories must be absolute to be able to provide an account of semantic competence, this worry becomes very acute. As far as we know, the only ways we have to provide absolute truth conditions rely on something like disquotation, at least at some points. I have argued it can appear in Montagovian or Davidsonian guises, but it does seem to be needed. But if that makes our theories explanatorily weak, we face a threat to the project of building theories of semantic competence that have any explanatory pay-off. That is a very serious concern.

We may conclude that the concern is very serious, but the general worry that disquotation threatens to make semantic theories explanatorily weak does not tell us where the problem arises, and how far it extends. And, to anticipate what is to come next, we of course know that in some cases our semantic theories do offer good explanations, so the problem cannot be totally pervasive. To see better how far it extends, it will be useful to consider

cases where semantic theories really do show their explanatory power, and see where disquotation and other techniques figure in them. We will look at one very clear case first; and then another, somewhat more multi-faceted case after that.

The example I shall start with is the semantics of determiners, where generalized quantifier theory has proved a rich and important tool. To briefly recall to mind some of the important aspects of this theory, remember that it treats determiners as relational expressions that compare cardinalities. For instance, we have:[17]

(7) $[\![\text{most}]\!](A, B) \longleftrightarrow |A \cap B| > |A \setminus B|$

The theory of which this is a representative statement makes non-trivial predictions, and offers some important generalizations. For instance, the well-known conservativity universal arose out of generalized quantifier theory, as did the Ladusaw-Fauconnier generalization on negative polarity items. These illustrate that the theory of generalized quantifiers does not suffer from the kind of explanatory vacuity Field worried about for disquotation. Whatever one's understanding of what makes a good explanation in science, it appears a safe conclusion that theories which make strong predictions and produce generalizations that were not available without the theory are good candidates to be solid explanatory theories.[18]

There is no known way to produce any of these sorts of results in purely disquotational form. Indeed, I know of no serious work on the semantics of

[17] As is common practice, I am putting the semantics of determiners in terms of relations rather than functions. The Montagovian tradition would want to 'Curry' these relations to fit them into the typed λ-calculus, where they come out of type $\langle\langle e, t\rangle\langle\langle e, t\rangle, t\rangle\rangle$. See Heim and Kratzer (1998) for the details. Likewise, the Davidsonian tradition would prefer to put determiner meanings in terms of the sort of Val function we saw above. See Larson and Segal (1995) for the details. The important point for us is that both approaches make use of the same core generalized quantifier theory, and it is a fairly routine matter to put that theory in whichever form is required.

[18] For those unfamiliar with these results from generalized quantifier, I ask you to take on faith that they provide examples of prediction and generalization. The generalized quantifier theory of determiner meanings arose out of seminal work of Barwise and Cooper (1981), Higginbotham and May (1981), and Keenan and Stavi (1986). See Keenan and Westerståhl (1997) and Westerståhl (1989) for surveys, where explanations of conservativity and the Ladusaw-Fauconnier generalization can be found. In claiming this is a very rich explanatory theory, I am by no means claiming it is a complete theory of all issues surrounding quantification in natural language. For some issues it does not solve, and approaches to them, see Szabolcsi (1997, 2010).

determiners that is disquotational. Both Montagovian and Davidsonian approaches to semantics at this point adopt the theory of generalized quantifiers in some form.[19]

The striking contrast we see with the semantics of determiners is between uses of mathematics and of disquotation.[20] But glancing back at the issue of the role of model theory in semantics our discussion in section 1 raised, it is worth noting that the use of mathematics in the semantics of determiners is rather different from the use of model theory as a *framework* that we saw with the *non*-absolute variant of semantics Lepore was criticizing. The mathematics of generalized quantifiers is often described as part of model theory, and indeed, it was developed originally in the setting of model theory and plays an important part in abstract model theory.[21] But, the semantic values for determiners do not require us to look at a whole range of models, as abstract model theory does, and as the target of Lepore's argument also did. At core, the mathematics involved is some elementary set theory, which makes cardinality comparisons between sets given as inputs to a generalized quantifier. In practice, a lot of model theory uses similar mathematics, sometimes focusing on particular models, and sometimes on classes of models. But for our application, it is not important that the models vary, and we can use generalized quantifiers in an absolute setting. The objects that form the sets whose sizes are compared can be real-world objects comprising the extensions of real-world properties, just as absolute semantics requires.[22]

[19]For instance, Higginbotham (1988, p. 44), speaking for the Davidsonian approach, describes model theory as playing a role in a "special department of lexicography concerned with the meanings of those expressions that remain invariant under various morphisms and permutations."

[20]This should not be overstated, as Kevin Scharp made clear to me. The non-disquotational feature of the semantics of determiners is that it uses the mathematics of generalized quantifiers to state the meanings of certain expressions, rather than simply repeating those expressions. The mathematics itself will ultimately need explaining, and that will raise a whole range of other issues about how explanations of foundational concepts in mathematics can be framed. But we need not worry about those here.

[21]See the early papers of Lindström (1966) and Mostowski (1957) and the papers in Barwise and Feferman (1985).

[22]This point is already made in generalized quantifier theory, which distinguishes *local* from *global* generalized quantifiers. The kind of semantic value given to a determiner is a local generalized quantifiers, while a global generalized quantifier is a function from domains to local quantifiers. Global quantifiers are, of course, the basic notion for studying model-theoretic logics, where the space of models is fundamental. According to the absolute approach to semantics, local quantifies are the right tool for semantic theory. There

So far, we have seen that disquotation can fail to be explanatory, and semantic theories that rely on disquotation can fail to be explanatory theories. As being explanatory is surely one of the important features of a good theory, we can conclude that too much disquotation can make for bad semantic theories. We also considered the example of the semantics of determiners and generalized quantifier theory, where semantic theory is clearly explanatory, by the rough-and-ready standard of making predictions and formulating generalizations. We have seen that in this case, some mathematics is used to build a more explanatory theory. It is sometimes described as model theory, but it is more apt to simply observe that a little bit of mathematics enables us to formulate generalizations and make predictions in this case. This example shows us that, at least in some cases, we have a contrast between disquotation, which is explanatorily very weak, and uses of mathematics, which can figure in good explanatory theories. (We will look at one more example of good explanation in semantics in a moment.)

In many ways this conclusion is not really so remarkable. It simply reaffirms the well-known fact that mathematics often plays an important role in empirical theories. We already knew from fields like physics that mathematics can help build rich explanatory theories, that go beyond merely listing features of the physical world. Indeed, as we see in this case too, mathematics developed independently of any particular empirical science can often have far-reaching effects when applied. We see that with model theory and set theory in semantics. This fact is a well known, if not so well understood point in the philosophy of mathematics.[23] I shall not try to explain why it is the case; I shall merely note that we can find good explanations in semantics, and that mathematics, rather than disquotation, is one resource that helps build them.

This sort of mathematically based explanation can be offered by semantics in either the Montagovian or Davidsonian tradition (and indeed, both happily make use of generalized quantifier theory), and as I mentioned, it is perfectly at home in an absolute semantics. The illustration from general-

are cases where global results are stronger, as we see, for instance, in the discussion of restricted quantifiers in Westerståhl (1989). In the spirit of simply helping ourselves to any mathematics that is useful, we might even find those results revealing about semantic properties of natural languages. But we can still keep our semantic values absolute if we wish.

[23]Famously, Wigner talked about the 'unreasonable effectiveness of mathematics'. See Steiner (1998).

ized quantifier theory shows that the right mathematics can form the core of an explanatorily rich empirical theory. The illustration also shows how such a theory can be drastically different from mere disquotation. Whereas disquotation threatens to provide mere lists, and so to be explanatorily too weak, a little mathematics can set the stage for deeper analyses that do far more than simply listing semantic facts.

In looking at determiners, we see a way that we can rely on some mathematics to get beyond mere disquotation, and how that can lead to more explanatory theories. The case of determiners, and the application of generalized quantifier theory, is an impressive one; but similar points could have been made with other expressions and other tools from mathematics. Tense could easily provide another example, as could mood and modals, focus, and so on. Generally, the so-called functional categories seem to lead us to rich theories where a little mathematics goes a long way in increasing explanatory power. It is sometimes said that we see the power of formal semantics in 'compositional semantics', where the ways the meanings of expressions combine is the primary focus. That may be correct, but we also see that compositional analyses are built out of the meanings of functional expressions, and we have at our disposal powerful theories of how those meanings work.

But there is a way in which these sorts of cases avoid some important issues about meaning. In many cases, the core meaning of what we say is not carried by the functional expressions and compositional structure. Rather, it is carried by the major lexical categories: nouns, verbs, and adjectives. We have seen how we can rely on some mathematics to get beyond disquotation for functional categories, but what about the lexical ones?

I shall argue that with lexical categories, we find something intermediate between mere disquotation and the kind of mathematical theory generalized quantifiers provides. We find some non-trivial applications of (at least a little) mathematics in stating the meanings of lexical items, and along with that mathematics goes some important explanatory force for our lexical semantic theories. But we also find that, unlike the generalized quantifier case, mathematical analysis gives out in most cases of lexical categories. When it does, we fall back on disquotation in some form, though not always the kind of pure disquotation we have encountered so far.

I shall make this argument by considering another case: that of gradable adjectives. These have been intensively studied in recent years, and clearly show the contributions of more mathematical and disquotation-like factors

in providing their meanings.

To begin, I shall sketch one common approach to the semantics of gradable predicates. (I shall follow the influential presentation of Kennedy (1997, 2007).) This semantics makes the core meaning of an adjective a function from individuals to degrees on a scale.[24] For instance, the meaning of *tall* is given by:

(8) $[\![\text{tall}]\!](x) = d$ a degree of tallness

Technically, a scale is an ordered collection of degrees, with a *dimension* specifying what the degrees represent. In this case, the dimension is tallness, or height.

Scale structure is useful for the compositional analysis of constructions in which adjectives occur. For instance, it helps analyze comparative constructions, and provides analyses of measure phrases and degree terms like *very*. But recent work has shown it can also have relevance to issues in the lexical semantics of gradable adjectives themselves. For instance, it is a common observation that gradable predicates frequently come in antonym pairs, like *short* and *tall* and *fast* and *slow*. If their meanings are based on scale structures, there is an easy explanation for why this should be: an ordered set of degrees always provides an inverse ordering, which generates the antonym meaning. Hence, we gain some insight into the ways gradable adjectives group into semantic classes, which is one of the phenomena lexical semantics often studies.

Recent work by Kennedy and McNally (2005) shows that stronger results are available. For example, scales can help explain which adjectives can combine in complex comparisons. Distinct adjectives which share a scale can combine in comparative constructions, while adjectives that do not share a scale cannot. If we suppose that *wide* and *tall* both have a scale of linear extent, while *flexible* does not, we predict (Kennedy, 1997; Kennedy and McNally, 2005):

(9) a. He is as wide as he is tall.
b. ?? He is as tall as he is flexible.

[24] Degree analyses are also developed by Barker (2002), Bartsch and Vennemann (1973), Cresswell (1977), Heim (1985), and von Stechow (1984). The main alternative, the partial predicate analysis, has been developed by Fine (1975), Kamp (1975), Klein (1980), McConnell-Ginet (1973), and Pinkal (1995).

This prediction appears correct, and we have the beginnings of an explanation of the fact.[25]

So far, we have simply exploited the most elementary structure of scales. But a little bit of mathematics can be applied to get further explanations. In particular, Kennedy and McNally (2005) note that the basic topological properties of scales can offer some explanations of semantic properties of adjectives. For instance, scales can be open or closed. This provides a typology of adjective meanings, which helps explain some of their interesting semantic properties. One is the difference between absolute adjectives including *wet* and *full*, and relative ones including *tall* and *large*. For instance, they show different entailment patterns with their antonyms:

(10) a. The door is not wet.
entails
The door is dry.

b. The door is not large.
does not entail
The door is small.

Furthermore, they differ in what proportional modifiers they allow:

(11) a. half full
b. * half tall

Kennedy and McNally (2005) offer explanations of these facts that start with the observation that adjectives like *wet* are associated with extrema of closed scales. Think, for instance, of measuring the amount of water on something, by values between 0 and 1, to form a scale for *wet*. Then for something to be wet is for it to have degree of wetness > 0, i.e. to have some water on it. This supposes we have a scale that is closed at 0. Topologically, it might look like $[0, 1)$. The antonym *dry* inverts the scale, so we can measure dryness by how much water is on something. But then to be dry is to have 0 degrees of dryness. That is the maximum amount, so we have a scale $(1, 0]$. In contrast, an adjective like *tall* is associated with a scale that is open at

[25]There is a little more to the story, of course. Just what counts as sharing a scale can be complicated to spell out. *Tall* might be linear extent on the vertical axis, while *wide* might be linear extent on the horizontal axis, for instance. If so, then just being scales of linear extent counts as being sufficiently similar.

both ends—there is a lower limit but it is outside the scale. So, topologically it might look like $(0, 1)$.

Generally, absolute adjectives are associated with closed scales, while relative adjectives are associated with open scales. This leads to a number of differences in their behavior, especially when it comes to their positive forms (i.e. non-comparative), as Kennedy and McNally observe. For instance, in positive form, relative, open-scale adjectives like *tall* look to the context for a standard, which says what degree of tallness you need to be tall. In contrast, absolute, closed-scale adjectives set their standards in terms of the extrema their scales, even in positive form, as we saw with *wet* and *dry*. They thus show markedly less context dependence than relative adjectives like *tall*, since they do not generally rely on context to fix a cut-off point for being dry or wet.

With this sort of analysis in hand, the entailments we see in (10) follow easily. If something is not wet, it does not have degree of wetness > 0, so its degree of wetness is 0. That is what it takes to be dry, and so the entailment follows. In contrast, to not be large is to not have more than a contextually determined degree of largeness, which is not sufficient to guarantee that something has more than the contextually determined degree of smallness (cf. Cruse, 1986; Rotstein and Winter, 2004). Adjectives like *full* seem to be associated with scales that are closed at both ends, and so can look something like $[0, 1]$. This gives us an indication of why we get the contrast in (11), as *half* can make mathematical sense on the appropriate scales, by requiring a degree value halfway between the two endpoints of a scale. If we are starting with a scale like $[0, 1]$, it might simply require the value $1/2$.

This is just a taste of how explanations in this theory go, and what they rely on. The point of these examples is that by applying a modest amount of mathematics (very modest indeed, by the standards of, say, mathematical physics), we get some substantial explanations of the lexical properties of gradable adjectives. We learn something about the nature of their meanings, which allows us to formulate generalizations, make predictions, and generally offer good explanations. We seem to have a good explanatory theory. Again, something similar could be observed for other lexical categories, notably verbs, where a huge amount of research has been carried out.

All the same, there remain some ways in which the kind of theory of adjective meaning we are working with relies on disquotation. In the particular theory I have presented here, it emerges with the *dimension*. Recall, an adjective meaning is built around a degree-valued function, where degrees

are degrees on some scale, with a specified dimension. Spelling this out, we have:

(12) a. $[\![[_{\mathrm{A}}\alpha]]\!](x) = d \in S$

b. S a scale: an ordered collection of degrees with a dimension

A *Dimension* is a specification of what the degrees represent, i.e. for *tall*, the dimension can be given as height, or perhaps linear extent along the vertical axis. For *flexible* we have a dimension of flexibility or bendability. In many cases, these are specified just this way: we just say that the dimension for *flexible* is flexibility. Hence, we in effect have disquotation. It is not quite the simple disquotation of saying *tall* applies to tall things, since we *also* have significant structure from the scale. But at the level of the dimension, we often really do have disquotation. We sometimes have something near to disquotation; if, for instance, we say the dimension for *flexible* is bendability. We have less when we say the dimension for *tall* might be linear extent on the vertical axis. Though not pure disquotation, these still have important features of disquotation. Though we do not simply repeat a phrase from the object language in the metalanguage, we do deploy the same concept, or a concept closely related to the one being expressed, and we do so by simply using the right expression in the metalanguage, without further analysis. In the linear extent case, we offer some modest analysis of the concept, but still rely on our prior understanding of concepts closely related to the one at work in the object language, and rely on the right phrases in the metalanguage to express them. We simply use words that express concepts like linear extent, and rely on our understanding that they capture the important aspect of tallness. Though not pure disquotation, this is very near to disquotation. Hence, I shall simply talk about disquotation to cover both the pure and impure cases. The way dimensions are specified shows that we do not fully avoid disquotation (pure or impure) in our theory of adjective meanings.

Though we have a rich explanatory theory of adjective meanings, building on a little mathematics, we in the end fall back on disquotation to fully specify adjective meanings. This contrasts with the case of determiners, where the mathematical side of our theory provided an essentially complete analysis of the concepts involved. The explanatorily fruitful, mathematically based side of the theory provides the full meanings of determiners, as generalized quantifies. Nothing more is needed. With genuine lexical categories, our example indicates, we do ultimately fall back on disquotation. The mathematical side of the theory helps to formulate generalizations and predictions,

and so helps give the theory explanatory force. But when it comes to really saying what the expressions mean, we make crucial use of disquotation (pure or impure). Knowing what ordering goes with *flexible* might help explain a number of its properties, but we do not specify that it means flexible just by displaying that ordering. We need to specify that the dimension is flexibility, which we do disquotationally, or nearly so. As we should expect from our discussion of the explanatory force of disquotation, our theory of adjectives loses its explanatory power just where disquotation enters to specify dimensions. Though we do not state the meaning of an adjective without specifying the dimension, disquotational specifications do not support generalizations or predictions. As we saw before, they merely list the relevant properties of expressions.

The morals drawn from this brief foray into gradable adjectives could be drawn from a great many areas of research in lexical semantics, including the lexical semantics of verbs which has been the focus of a huge amount of research in recent years. I find the case of adjectives a nice illustration, so I shall not go into further examples.

So far, we have been exploring the interplay between disquotational and non-disquotational aspects of semantics. We have seen that both current Montagovian and Davidsonian semantics state absolute truth conditions. Both can and do build explanatorily substantial theories. Though they differ in some details of analyses they propose, both offer interesting generalizations and make predictions, just as any good empirical theory should. When they do, they typically rely on at least a little bit of mathematics, and hence, we see another instance of the amazing way that mathematics can apply in empirical theories. The case of the semantics of determiners illustrates this vividly. The techniques used there come from a branch of model theory (or set theory), but there is nothing special about model theory in this regard, as we saw from the use of (a tiny bit of) topology in the case of gradable adjectives.

Yet, at the same time, each approach at some point or another falls back on disquotation, or near-disquotation, when it comes to the semantics of the lexical categories. This is not a special feature of Davidsonian semantics. We saw it equally for a theory of adjectives couched in Montagovian terms. Of course, I also claim, but shall not go into further details here, that we see the same thing for other theories of lexical meaning. And, we see, where our lexical semantic theories fall back on disquotation, they correspondingly lose explanatory force. Our theories offer interesting explanations and predic-

tions, often building on some mathematics, but complete their descriptions of lexical meanings in explanatorily weak ways by falling back on disquotation.

These conclusions are based on examples from current work in semantics. But assuming they are right, it shows that when it comes to lexical meaning, we typically find two aspects: a more explanatory side, relying on mathematics, and a less explanatory side, relying on disquotation. There are extremes, like the semantics of determiners which seems to be virtually free of disquotation. But the pattern is striking, and seems stable when it comes to the lexical categories.

This conclusion, if it indeed holds, already supports a particular kind of partiality in semantics. When it comes to explanatory force, our semantic theories are for the most part partial. Where they rely on mathematics, they are explanatorily substantial, but where they rely on disquotation they fail to be. In the cases of lexical categories, our semantics theories only provide partial explanations of semantic properties. At least as things currently look in semantics, this is unavoidable. We do not know how to build a semantic theory which avoids using disquotation at some point, outside of special cases like the determiners.

In the next section, I shall suggest a way to make sense of why this might be so. It will involve taking the step from partiality in explanation to partiality in subject-matter. If we are on the right track with our theories, we should infer from the pattern in explanatory force what our subject-matter is really like. Hence, I shall conclude, semantic competence is only a partial determinant of content. But, there must be more to the story than that, since disquotation, explanatory or not, is doing something in our lexical semantic theories. I shall address this is in the next section too. I shall propose a way to understand what it does, and why it is an indicator of partiality in semantics.

Before turning to this, I shall conclude this section by briefly noting one loose end. The examples of determiners and gradable adjectives show a contrast between mathematical and disquotational aspects of a semantic theory, and show that good explanations can go with the mathematical but not the disquotational side. But it is not clear if good explanation in semantics is exclusively the province of mathematics, or if other techniques can provide them as well. At this point, it seems to me that it is not clear one way or the other. Lexical semantics is full of non-mathematical formulations, and it is not simple to track where they offer substantial explanatory pay-off, and where are they are nearly disquotational. Mathematics plays the role of

what Higginbotham (1989a) calls 'elucidations of meaning', but there are at least proposals for such elucidations which are not couched in mathematical terms. One very well-known example is the proposed characterization of *cut* as a linear separation of the material integrity of something by an agent using an instrument (Hale and Keyser, 1987; Higginbotham, 1989a). This does seem to give some insight into the meaning of the verb, and so does not seem to show the explanatory weakness of disquotation or near disquotation. (It helps explain why *cut* is different from *crush* and *disintegrate*, for instance.) Likewise, a great deal in the lexical semantics of verbs posits various forms of decomposition of their meanings, but does little mathematical work.[26]

Would mathematical models provide more fruitful explanation than these sorts of analyses? Perhaps. Mature theories often get couched in more mathematical terms, and perhaps more fully developed lexical decomposition theories or elucidations might benefit from more thoroughly mathematical formulations. But it is not always easy to spot where the explanatory work is being done. So, it is not easy to tell. With this caveat, I shall rest with the observation that at least one important way we can get explanatory content from semantic theories is through mathematics, and not through disquotation.

3 Partiality in Semantic Competence

So far, we have concluded that our semantic theories show a particular sort of partiality: insofar as disquotation is explanatorily weak, and semantic theories fall back on disquotation at some point, they are explanatorily partial. We got here by a circuitous route. We have spent most of our time examining the various roles disquotation and model theory (and other branches of mathematics) play in semantic theories. We began by looking at the old question of absolute versus relative semantics, where we agreed that semantics should be absolute, but noted that this does not undermine a large body of work in the Montagovian tradition. Indeed, it leaves a role for model theory, as we illustrated with the case of determiners and generalized quantifier theory. But, we saw, that is not a use of model theory as a framework, and does not undermine absoluteness. It is an application of mathematics to an

[26]See Levin and Rappaport Hovav (2005) for a review. It is worth noting that the classic Dowty (1979) does seek to develop mathematically rich characterizations of the building blocks of decompositional analyses.

empirical subject. We saw, furthermore, that such applications tend to be the points where our semantic theories provide good explanations. We saw this with the determiners, and in part with gradable adjectives. In contrast, where disquotation appears, our semantic theories tend to lose explanatory power. Finally, we saw that with the lexical categories (at least in one example), even though we can build good explanatory theories that make use of various sorts of mathematics (and maybe other explanatory strategies), in the end we fall back on disquotation to fully fix contributions to truth conditions. At those points, our lexical theories indeed lose explanatory value. So, by a roundabout route, we reached our conclusion that semantic theories are explanatorily partial.

Let us assume that this not merely a reflection of some error in our current semantic theories, but a feature of them that is genuinely correct. Thus, let us assume it will not change fundamentally as our theories improve. Then we have to conclude that there are aspects of meaning that defy explanation by semantic theory. How are we to make sense of this? I shall suggest that we may do so best by seeing the subject-matter of semantics itself as partial, and I shall go on in this section to explore the nature of this partiality.

The idea I shall explore here is that the partiality of substantial explanations in semantics reveals that the elements of semantics that fall within the scope of linguistic competence—fall within the language faculty—only constitute a partial determination of truth conditions. To explain what this might mean, let me sketch a picture of how lexical meaning and the language faculty might relate that is compatible with it. The leading idea is that lexical meaning is an *interface phenomenon*. Let us, for the moment, think of this explicitly in terms of a language faculty. It is a common idea that the language faculty interfaces with other aspects of human cognition and action. At the very least, it must when it comes to actually speaking, where linguistic information is presumably passed on to the systems of articulation. Likewise, at some point, our sentences enter into our broader cognitive lives where we reason with them, assess them, and use them in countless other ways.[27]

The specific idea at issue here is about how lexical entries are structured. We shall take these to be part of the language faculty, containing, as always,

[27]Hence, as Chomsky has long-emphasized, we can expect two interface levels, PF (phonological form), which interaces with the articulatory system, and LF (logical form), which interfaces with the conceptual system (Chomsky, 1995, 2000).

syntactic, phonological, and semantic information. (In many recent theories, the first two are often thought of as provided by bundles of features.) When it comes to the semantics, we have good reason to think that some aspects of meaning are likewise coded into the lexicon by the language faculty. Why? Because of the sorts of explanatory successes we saw for some parts of semantic theory. Where we see good, explanatory, successful theories, we have reason to accept that they are describing some phenomenon. Insofar as these theories are part of the broader enterprise of linguistic theory—describing the language faculty—we have reason to infer that they show us something that is genuinely part of the language faculty.[28]

As I have been emphasizing, the good explanations that lead us to posit features of the language faculty for semantics eventually give out, and are replaced by disquotation. But what do we then do with this lack of explanation; especially, persisten lack of explanation we think is an unavoidable aspect of our theories? The proposal I want to explore is that lack of explanation should be reason to resist positing features. If substantial explanatory pay-off is a sign that our theory is describing something really in the language faculty, then lack if it should be taken as a sign that it is not. To be clear, the latter does not follow from the former, but it is a reasonable conclusion to reach nonetheless. If we persistently cannot get a good explanation out of theories describing some domain, then one reasonable conclusion is that there is nothing in that domain to describe. This, as we will see in moment, will get us to genuine partiality in semantic competence, as the determinants of truth conditions will not all lie within the language faculty.

I think this is generally the kind of conclusion we should draw from persisten lack of explanatory value. But in the particular case at hand, we face a problem in doing so. Showing how to solve the problem will be my main point in favor of the partiality proposal. Once it is solved, we will be able to make the generally sensible response to persisten partiality in explanation without further difficulties. The problem is that disquotation seems to be an ineliminable feature of our semantic theories. We cannot, and we think

[28]Recall, the broader enterprise of linguistic theory involves a range of evidence that we think indicates it is describing the language faculty, including evidence about acquisition, and more recently, neurological evidence. For a review of how these fit into semantics, see Krifka (2011). For some psycholinguistic work related to aspects of gradable adjectives we discussed in section 2, see, for example, Frazier *et al.* (2008), Syrett (2007), and Syrett *et al.* (2010). For work related to the semantics of determiners, see, for example, Hackl (2009) and Lidz *et al.* (2011).

will not be able to, formulate semantic theories without falling back on disquotation. It is explanatorily weak, so we would like to say there is nothing in the language faculty that it describes. But if there is nothing in the language faculty for disquotation to describe, it seems it should rather be an appendage to our theories we can eliminate. Why is it ineliminable if it does not describe anything in the target domain?

The way to resolve this problem, I propose, is by a more nuanced picture of the lexicon, and how it fits into the language faculty. I propose that what are in the lexicon corresponding to disquotation in our theories are simply pointers to other conceptual systems. If you like, they are pointers to *concepts* which are indicated by the non-quoted side of a disquotation clause. But I am not here insisting on a particular view of concepts.[29] All we need is that we have pointers to something outside of the language faculty proper which provides further content, sufficient to fix contributions to truth conditions, at least. If this is right, then it would be no surprise that substantial generalizations or explanations give out where we find disquotation, when it comes to theories that seek to describe the language faculty. All that are in the language faculty corresponding to those parts of the theory are pointers, with no further within-faculty content to be described or explained. But pointers are parts of lexical entries, so we should expect to see something in our theories corresponding to them.

One way to think of this general idea being implemented is to think of the specifically linguistic content, provided by the language faculty, as forming a kind of structural frame within which a pointer to a concept appears. This is in fact the way a great deal of work in lexical semantics proceeds. When we look at the kind of meaning we supposed for gradable adjectives in section 2, for instance, we see something like (8), repeated here:

(13) $[\![\text{tall}]\!](x) = d$ a degree of tallness

But unpacking this a little bit, recall that the codomain of the function is really a complex object $\langle D_\delta, <_\delta \rangle$, where D_δ is the set of degrees of dimension δ, ordered by $<_\delta$. As degrees are just abstract points, whose function is to

[29]For instance, there is nothing particularly internalist about this interface picture. If content is wide, it is just that pointers are to whatever provides wide content. Presumably, whatever pointers point to must provide extensions for predicative expressions, if they are to fix their contributions to truth conditions. If the targets of pointers are concepts, this does raise some complicated issues about how extensions relate to concepts, but I shall not pursue those here.

be ordered by $<_\delta$, it is the dimension δ that does the work of providing the specific content of any particular adjective. So, a more explicit entry would be something like:

(14) a. $S_{\text{tall}} = \langle D_{\delta_{\text{tall}}}, <_{\delta_{\text{tall}}} \rangle$
b. $[\![\text{tall}]\!] \colon D_e \to S_{\text{tall}}$

We might find reasons to restrict the domain of the function (see Kennedy, 2007), but this gives a sense of what the lexical entry might look like, spelled out more fully.

This illustrates the way the dimension is *packaged* by the lexicon. The dimension expresses tallness, or linear extent. This is then packaged into a scale structure, and then that is packaged into a degree-valued function, which is the appropriate meaning for an adjective.

Our discussion of adjectives in section 2 showed that, as at least as far as our best current theories go, the semantics makes use of the topological properties of $\langle D_\delta, <_\delta \rangle$, and so makes heavy use of the ordering relation $<_\delta$. It also makes use of the structure of degree-valued functions, e.g. in comparative constructions. The semantics 'sees' the packaging. But I know of no reason to think that it makes use of any internal structure of δ. As far as the semantics goes, δ is atomic. It marks whether or not two scales D_{δ_1} and D_{δ_2} are distinct, but does no other work in the theory. Sometimes this matters, as it allows us to note when two adjective share a dimension, and it distinguishes distinct adjective meanings in many cases. But it does so in a brute way, by simply marking a difference in scale. This is, as I have been arguing, reflected in the disquotational way that dimensions are often specified in our theories.

This is captured well by the idea of a pointer. All that is required, as far as the semantic theory is concerned, is whether or not it is the same pointer. That allows us to track sharing dimensions, and brute differences in meaning. As no internal structure of the dimension appears to be relevant in the theory, what the pointer points to is irrelevant. The pointer itself suffices for the theory. Of course, to really know what the adjective means, you have to know more, i.e. what it is to be tall or have linear extent. But semantic theories do not seem to make use of this. Hence, we think of that content as what the pointer points to: some concept in the broader conceptual system.

Let us call this view of the lexicon the *pointers conception*. It proposes that the lexicon contains, over and above important semantic information coded by the language faculty, also pointers towards features of meaning that are outside the language faculty. As such, they are atomic as far as se-

mantic theories that describe aspects of the language faculty are concerned. They are still important aspects of meaning, of course. They tell us what makes *tall* different from *fast*, for instance, which we will need to fix contributions to truth conditions. But they are not within the language faculty, and their structure plays no role in our theories. Hence, we just mark them disquotationally, or nearly disquotationally, within semantic theories. This amounts to merely listing these features of meaning, in a way that is explanatorily empty. As far as semantic theories go, this provides atomic markers of facets of meaning, which offer no interesting explanations in semantics. But the pointers so marked are pointers to something else, presumably the right sorts of concepts provided by other aspects of our cognition. Thus, their being atomic as far as linguistic theory goes does not mean there is no other theory which might describe their values in a rich explanatory way. We can happily hope that some other theory, from cognitive psychology, or wherever else is appropriate, might do this. But it will not be a theory which describes the language faculty.[30]

One of the virtues of the pointers conception is that according to it, we have no need to say that the semantic contents indicated by pointers are not there. Thus, we can do justice to the old observation that however weakly, disquotation statements do state substantial facts. These contents are just not part of the language faculty, and the facts stated are not facts about the language faculty itself (beyond the presence of pointers). Presumably they are parts of some other cognitive systems, to be explained by the appropriate theories from the wider realm of the cognitive sciences. Another virtue is that the pointers conception allows us to retain a rich role for semantics in linguistic theory, even if we are cautious about just how far our semantic theories will go. It provides rich lexical entries, with a broad range of linguistic properties which good linguistic theories can strive to explain. At the

[30]If this is right, then dimensions for gradable adjectives are their lexical roots. Moreover, I am endorsing the idea derived from Grimshaw (2005) that roots are atomic as far as semantic theory is concerned. A lexical root, on this picture, is *packaged* into a lexical entry, by combining it with distinctively linguistic structure. This is a fairly common idea in lexical semantics, notably, the semantics of verbs. I shall not go into details here. See Levin and Rappaport Hovav (2005) for a survey. I have discussed some of these issues with more of a focus on verbs in my (2011). I should mention that other approaches to lexical meaning, especially those relying on meaning postulates (Montague, 1973; Zimmermann, 1999), take a more top-down approach to specifying meanings, and do not really follow this sort of 'packaging' model. But, they still provide for partial speciations of meaning, and so lead to much the same kind of partiality as I am describing here.

same time, it does grant that when it comes to determining full contents, the language faculty might not be all that is needed to do the job.

The main argument I have offered for the pointers conception is that it accounts for the combination of explanatory and disquotational aspects we see in our semantic theories. Now, admittedly, this assumes that the current course of such theories is really on the right track. Even if the exact division between explanatory and disquotational elements changes as our theories improve, the argument assumes that there will be some such division, and it will have the same kinds of basic effects as the ones we see in current theories. Given the trajectory of theorizing in lexical semantics, I think this assumption is warranted. It relies not on the specific details of any one proposal in lexical semantics, but on a broad pattern we see across approaches and theories. But I do pause to note that lexical semantics is still very much in development, and we may not yet know just what a more fully developed form will look like. Hence, we should count the pointers conception as a hypothesis, which I think is justified given our current state of knowledge.[31]

[31]It is natural to ask whether the pointers conception falls on the side of descriptive or foundational semantics in the sense of Stalnaker (1997). According to Stalnaker's way characterizing these, descriptive semantics says what the semantic values of expressions are, while foundational semantics says what makes it the case that the expressions have the values they do. (A related distinction is drawn by Kaplan, 1989.) In light of this, the pointers proposal in first instance falls squarely on the descriptive side. It is a high-level proposal about what descriptive theories should look like, but it is about descriptive theories in semantics. If the pointers conception is right, then the official foundational semantics would have to explain how expressions get the packaging and pointers they do, and what makes it the case that the pointers point to what they do, etc. The pointers conception does not answer these questions. (One could imagine various stories, involving causal interactions, associationist style connections, learning mechanisms, and so on.) Officially, the pointers conception falls on the descriptive side, but there are, as always, complications. I have focused a great deal on explanations in semantics—in descriptive theories. We always require good explanations from good, empirical, descriptive, theories, which reminds us that merely listing what semantic properties words and phrases have is not enough to be a good descriptive theory. But also, the distinction tends to get muddied in practice, especially in the setting of generative linguistics, with its focus on psychology and acquisition. Explaining acquisition in some ways would explain how expressions get their meanings, for instance. A more specific way that the pointer conception muddies the distinction is that in positing two factors in meaning, one from the language faculty and one from the broader conceptual domain, the pointers conception does seem to have some foundational implications, even if it is a descriptive proposal. (Thanks to Karen Lewis for raising this question.)

Before going on to consider another option for accounting for the explanatory partiality of semantic theories, let me pause briefly to reconsider the role of the language faculty in the pointers conception. What is really required to make the pointers conception work is some domain-specificity for linguistic competence. That allows us to think of semantic theories as describing a specific domain, and allows us to have cross-domain pointers. Positing a language faculty provides that domain specificity. Other properties typically attributed to the language faculty are not important to the pointers conception. For instance, whether or not the language faculty is innate, or more generally, how it might interact with learning mechanisms, is not important. Nor is it important just what the distinctive processes within the language faculty are, or even if it is simply a computational system.

There is another option which might account for the combination of explanatory and disquotational features we see in semantics. Rather than supposing there are pointers to extra-linguistic aspects of cognition, we might simply suppose that the lexicon contains atomic elements of meaning (e.g. atomic concepts).[32] I am not going to argue against this sort of atomism directly;[33] rather, I shall just point out that the pointers conception makes it unnecessary. Pointers behave atomically as far as our linguistic theories are concerned, but there is no need to insist that some other branch of cognitive science (or some other field) will not have a great deal to say about the concepts to which they point, and perhaps reveal their internal natures. We can have elements that look atomic as far as semantic theory goes, without incurring the cost of insisting they are atomic in any stronger sense.

The more pressing worry for the pointers conception, to my mind, is one internal to linguistic theory. Among the primary data for semantic theories are truth value judgments and entailment judgments. Indeed, entailment judgments are often some of the most important data for semantics. Such data are one of the reasons that our semantic theories are *truth-conditional* theories. Truth conditions explain such data very well. The objection to the pointers conception is that it puts important determinants of truth value and of entailment outside of the scope of linguistic theory proper, by putting it outside the language faculty. This is a very serious worry, and I am not going to address it fully. But it is worth noting that we do not lose all entailments

[32]This sort of view is congenial to the positions of Cappelen and Lepore (2005), Fodor (1998) and Fodor and Lepore (2002).

[33]For the case of the lexicon, this has been done by Collins (2011) and Johnson (2004).

on the pointers conception, as many entailments will be predicted by the substantial (non-pointer) features of a theory. In fact, we often find that discovering such entailments is an important factor in building just those explanatory aspects of the theory. Truth value judgments, however, wind up being the result of two factors: both features of meaning represented within the language faculty, and features of extra-linguistic concepts, will be involved. This makes the status of truth value judgments delicate. Many practitioners of semantics have known that they can be delicate, but this is a topic that needs further investigation.[34]

The pointers conception has the result of making semantic competence, as described by substantial linguistic theories and as determined by the language faculty, only a partial determinant of truth conditions. We get full truth conditions only when we combine linguistic competence with the extra-linguistic concepts to which lexical items point. Thus, we only get the full content of an utterance from the combination of linguistic and extra-linguistic factors. We thus have the partiality of semantics I advertised at the beginning of this paper. I have argued that attending to where our semantic theories provide substantial explanations makes a case for semantic partiality. But as I also mentioned at the beginning of the paper, I see the particular variety of partiality I have advocated as not skeptical about semantics or its place in linguistic theory, as some other recent varieties have been. The pointers conception includes the idea that pointers are packaged by the lexicon, and this carves out a very substantial role for semantics in describing the packaging. Insofar as the packaging is part of the language faculty, this is a very substantial role for semantics in linguistic theory. Indeed, it makes a place for the whole contemporary enterprise of truth-conditional semantic theory, and is optimistic about progress in semantics. It merely puts that enterprise in a specific domain, and grants that some partiality ensues. This might identify limits to what certain theories might do, but it is not skepticism.[35]

I shall conclude by noting one other feature of meaning that the pointers conception suggests. The evidence for the pointers conception came from the combination of explanatory and disquotational factors we see in semantic theories. But where we find explanation, we observed, we also often find

[34]See, for instance, Crain and Thornton (1998) and Krifka (2011).

[35]In may respects, the position I have been defending has much in common with the one in Pietroski (2005b, 2010). Not that we agree entirely: we differ at least in how sanguine we are about work in truth-conditional semantics. But, we agree on a number of foundational points.

the use of mathematics. (I was cautious about how much explanation and mathematics go together, but the trend is striking.) Does the specific role of mathematics tell us anything about the aspects of meaning that might be coded up in the language faculty?

I believe it is too early to tell for sure, since the exact role of mathematics in lexical semantics is still unclear (as I noted in section 2), but I shall speculate a little on the matter. The most visible explanatory pay-offs of semantic theory as it currently stands tend to concern more abstract aspects of linguistic competence. They involve things like scale structure, or the structure of generalized quantifiers, and so on. In being abstract, these lend themselves to mathematical descriptions which yield interesting explanations and predictions. But in being abstract, they are also typically not things whose nature is transparent to speakers in virtue of their linguistic competence, and will thus be highly tacit aspects of competence. Above I described the substantial aspects of semantic theory as describing the packaging of pointers. Typically, that packaging is abstract, and less transparent to speakers.

The packaging of meaning typically involves the broadly functional and other 'structural' aspects of human language. We have already noted that in the case of one functional category, the determiners, we get a more or less full characterization of their meanings from the mathematics, without any need for pointers. Such functional categories, including also tenses, moods, comparative morphology, etc., are often described as the 'grammatical glue' that holds sentences together. It is not a great surprise that their semantic properties might turn out to be genuinely part of linguistic competence, and within the domain of linguistic theory. When it comes to packaging within the lexicon, we see features of meaning that have become heavily grammaticalized yielding to semantic theory, like scale structure. These provide the packaging of root meanings, to build lexical items that the grammar can make use of.[36]

[36]One might worry that, in the case of the determiners especially, we have put too much mathematics into linguistic competence. The worry is not that no mathematical structure could be in the language faculty, but that we have inadvertently crossed out of the language faculty into whatever is responsible for cognition of mathematics with our determiner meanings. (One could ask this about any use of mathematics, but it seems most pressing in this case.) After all, in many cases, entailments following from the determiner meanings we posited will be complicated mathematical facts, which only some people will be able to recognize. (Thanks to Sam Cumming for pressing this worry.) A full answer to this interesting question will have to wait for another occasion, but let

So, the substantial part of semantics within linguistic theory seems, as far as we know, to describe broadly functional or structural elements within the language faculty. What is left out by linguistic theory, and were we fall back on pointers, are the very basic elements of meaning, which determine specifically which concepts (or properties and objects, or whatever you prefer) our words express. Where we might look for fundamental explanations of just what concepts are expressed by *tall* or *fast*, for instance, we find pointers instead. As I mentioned, this does not mean there is no explanation to be found, but puts it outside the linguistic theory, and outside the language faculty.

It is thus temping to propose that there are two types of meaning at work. One is structural-functional. This is genuinely part of the language faculty, and so is within the domain of linguistic theory. It is generally abstract, and often yields to mathematical description. The other is core conceptual meaning, which tells us just which concepts (or properties, or whatever you prefer) our words express. These enter linguistic theory and the language faculty only through pointers, and are marked by the places our semantic theories fall back on disquotation. We thus have genuine partiality

me gesture towards a reply. First, it is important to remember that the case of the determiners is an extreme one, both involving unusually substantial mathematics, and making it fully determine contributions to truth conditions. So, we are not supposing this will happen across the board. The justification in this special case was based on the value of the resulting theory, which as I stressed, makes all sorts of interesting predictions and generalizations. These require the semantic to see the generalized quantifier meanings for determiners, and so, we need to put them into language faculty. But having those meanings in the language faculty, as highly tacit representations of some kind, does not mean that all the mathematical entailments that follow from them will also be included in the language faculty. It does not, as we have not supposed that all the machinery of set theory and logical consequence is included in it. Thus, there is no requirement that speakers must be particularly good at recognizing the mathematical consequences of a generalized quantifier. Likewise, though entailment data is always important, the data involved in our positing determiner meanings is not the full range of entailments logically provided by generalized quantifiers. Lots of those won't be transparent to speakers who are average-to-poor at set theory. This general answer, though I think on the right track, raises a number of questions, like how generalized quantifiers are represented in the language faculty to allow them to do what they are supposed to in semantics, but not predict speakers will make entailment judgments they generally do not; and also, how the mathematics in the language faculty will interact with other aspects of cognition, like our ability to make judgments of size, and our ability to do mathematics itself. There is some empirically informed discussion of these questions in Lidz *et al.* (2011).

of meaning within linguistic competence.[37]

In acquiring a lexical item, you acquire a packaged pointer. The pointer is there, and so, your word means what it does. When you learn the word, you have to learn to what the pointer points. But when it comes to core meaning, provided by the pointer, I have stressed that we should not expect semantic theory—within linguistics—to tell us anything more than that. To illustrate, if you want to explain how *tall* and *closed* or *most* and *some* differ in meaning, semantics will tell you a great deal, as we have already seen. But if you want to explain how *tall* and *fast* differ in meaning, semantics will tell you very little, beyond that one is based on a scale of tallness (or linear extent), and the other on fastness (or speed), and that you must know this if you understand the words. Anything else you might want to know about their core meanings, like what makes tallness and fastness the concepts they are, what they really pick out, what internal structure they might have, etc., you will have to find out from somewhere else. You will find out more if you ask a psychologist, or a physicist, or even a philosopher.[38]

[37]One might worry that the pointers conception, and the kind of partiality it brings with it, leaves open where semantics as part of linguistic theory should end, and some other field, perhaps in cognitive psychology, should step in. Practitioners of semantics, especially, might want a more robust delineation of the empirical domain of their subject-matter, rather than simply pushing things until they fail, and then saying 'oh, well, it must be someone else's problem'. (Thanks to Magdalena Kaufmann for pressing me on this issue.) The distinction here between structural-functional and core conceptual aspects of meaning gives a rough indication of where the boundary may lie, and is perhaps enough to assure us that the domain of semantics is well-defined. But I admit that it is not precise enough to tell us where the boundaries are in practice. I'm afraid that may just be how things are. We don't know enough about the extent of the language faculty proper to say anything much more precise, so we may have to work out the boundaries as we go. Indeed, given our current state of knowledge, it often seems that the language faculty is quite idiosyncratic in what it encodes. So, it is hard to see how we could expect a much more precise answer at this point.

[38]Throughout, I have been focusing on the meanings of predicative expressions, like adjectives. This is not an accident, as a great deal of work in linguistic semantics has focused on such expressions, especially, on verbs. But at some point, more will have to be said about how *reference* fits into the picture. So, for instance, we will have to sort out what, if any, kinds of core meanings different referring expressions have, and what kinds of packaging those meanings might receive. We likewise will have to examine how the combination of those fixes reference, and what other contributions to truth conditions referring expressions might make. As is well known, context dependence enters into the picture for many important referring expressions, especially demonstratives and indexicals. Obviously, addressing these matters would engage a wide range of issues that have been

References

Barker, C., 2002. The dynamics of vagueness. *Linguistics and Philosophy* **25**:1–36.

Bartsch, R. and T. Vennemann, 1973. *Semantic Structures: A Study in the Relation between Semantics and Syntax*. Frankfurt: Athenäum, 2nd edn.

Barwise, J. and R. Cooper, 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* **4**:159–219.

Barwise, J. and S. Feferman (eds.), 1985. *Model-Theoretic Logics*. Berlin: Springer-Verlag.

Burgess, A., 2011. Mainstream semantics + deflationary truth. *Linguistics and Philosophy* **34**:397–410.

Cappelen, H. and E. Lepore, 2005. *Insensitive Semantics*. Oxford: Blackwell.

Carnap, R., 1947. *Meaning and Necessity*. Chicago: University of Chicago Press.

Carroll, L., 1960. *The Annotated Alice*. Meridian. Introduction and Notes by Martin Gardner.

Carston, R., 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Chierchia, G. and S. McConnell-Ginet, 2000. *Meaning and Grammar: An Introduction to Semantics*. Cambridge: MIT Press, 2nd edn.

Chomsky, N., 1965. *Aspects of the Theory of Syntax*. Cambridge: MIT Press.

———, 1980. *Rules and Representations*. New York: Columbia University Press.

———, 1986. *Knowledge of Language*. New York: Preager.

———, 1995. *The Minimalist Program*. Cambridge: MIT Press.

the subject of extensive debate already, and a huge literature.

———, 2000. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.

Collins, J., 2002. Truth or meaning? A question of priority. *Philosophy and Phenomenological Research* **65**:497–536.

———, 2004. Faculty disputes. *Mind and Language* **19**:503–533.

———, 2011. Impossible words again: Or why beds break but not make. *Mind and Language* **26**:234–260.

Crain, S. and R. Thornton, 1998. *Investigations in Universal Grammar*. Cambridge: MIT Press.

Cresswell, M. J., 1973. *Logics and Languages*. London: Methuen.

———, 1977. The semantics of degree. In B. Partee (ed.), *Montague Grammar*, pp. 261–292. New York: Academic Press.

———, 1978. Semantics and logic. *Theoretical Linguistics* **5**:19–30.

Cruse, D. A., 1986. *Lexical Semantics*. Cambridge: Cambridge University Press.

Davidson, D., 1967. Truth and meaning. *Synthese* **17**:304–323. Reprinted in Davidson (2001).

———, 1973. In defense of convention T. In H. Leblanc (ed.), *Truth, Syntax and Modality*, pp. 76–86. Amsterdam: North-Holland. Reprinted in Davidson (2001).

———, 2001. *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press, 2nd edn.

Dowty, D. R., 1979. *Word Meaning and Montague Grammar*. Dordrecht: Reidel.

Dowty, D. R., R. E. Wall, and S. Peters, 1981. *Introduction to Montague Semantics*. Dordrecht: Reidel.

Etchemendy, J., 1988. Tarski on truth and logical consequence. *Journal of Symbolic Logic* **53**:51–79.

Field, H., 1972. Tarski's theory of truth. *Journal of Philosophy* **69**:347–375.

Fine, K., 1975. Vagueness, truth and logic. *Synthese* **30**:265–300.

Fodor, J. A., 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.

Fodor, J. A. and E. Lepore, 2002. *The Compositionality Papers*. Oxford: Clarendon Press.

Frazier, L., C. Clifton, Jr., and B. Stolterfoht, 2008. Scale structure: Processing minimum standard and maximum standard scalar adjectives. *Cognition* **106**:299–324.

Glanzberg, M., 2009. Semantics and truth relative to a world. *Synthese* **166**:281–307.

———, 2011. Meaning, concepts, and the lexicon. *Croatian Journal of Philosophy* **31**:1–29.

———, forthcoming. Logical consequence and natural language. In C. Caret and O. Hjortland (eds.), *Foundations of Logical Consequence*. Oxford: Oxford University Press.

Grimshaw, J., 2005. Semantic structure and semantic content in lexical representation. In *Words and Structure*, pp. 75–89. Stanford: CSLI Publications.

Hackl, M., 2009. On the grammar and processing of proportional quantifiers: *Most* versus *more than half*. *Natural Language Semantics* **17**:63–98.

Hale, K. and S. J. Keyser, 1987. A view from the middle. Lexicon Project Working Papers 10, Center for Cognitive Science, MIT.

Heim, I., 1985. Notes on comparatives and related matters. Unpublished ms, University of Texas, Austin, available at http://semanticsarchive.net/Archive/zc0ZjY0M/.

Heim, I. and A. Kratzer, 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.

Higginbotham, J., 1985. On semantics. *Linguistic Inquiry* **16**:547–593.

———, 1988. Contexts, models, and meanings: A note on the data of semantics. In R. Kempson (ed.), *Mental Representations: The Interface Between Language and Reality*, pp. 29–48. Cambridge: Cambridge University Press.

———, 1989a. Elucidations of meaning. *Linguistics and Philosophy* **12**:465–517.

———, 1989b. Knowledge of reference. In A. George (ed.), *Reflections on Chomsky*, pp. 153–174. Oxford: Basil Blackwell.

———, 1992. Truth and understanding. *Philosophical Studies* **65**:3–16.

———, 1998. On higher-order logic and natural language. *Proceedings of the British Academy* **95**:1–27.

Higginbotham, J. and R. May, 1981. Questions, quantifiers and crossing. *Linguistic Review* **1**:41–79.

Horisk, C., 2007. Truth, meaning, and circularity. *Philosophical Studies* **137**:269–300.

Johnson, K., 2004. From impossible words to conceptual structure: The role of structure and process in the lexicon. *Mind and Language* **19**:334–358.

Kamp, H., 1975. Two theories about adjectives. In E. L. Keenan (ed.), *Formal Semantics of Natural Language*, pp. 123–155. Cambridge: Cambridge University Press.

Kaplan, D., 1989. Afterthoughts. In J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan*, pp. 565–614. Oxford: Oxford University Press.

Keenan, E. L. and J. Stavi, 1986. A semantic characterization of natural language determiners. *Linguistics and Philosophy* **9**:253–326. Versions of this paper were circulated in the early 1980s.

Keenan, E. L. and D. Westerståhl, 1997. Generalized quantifiers in linguistics and logic. In J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*, pp. 837–893. Cambridge: MIT Press.

Kennedy, C., 1997. *Projecting the Adjective*. Ph.D. dissertation, University of California, Santa Cruz. Published by Garland, 1999.

———, 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* **30**:1–45.

Kennedy, C. and L. McNally, 2005. Scale structure, degree modification, and the semantics of gradable predicates. *Language* **81**:345–381.

Klein, E., 1980. A semantics for positive and comparative adjectives. *Linguistics and Philosophy* **4**:1–45.

Klein, E. and I. A. Sag, 1985. Type-driven translation. *Linguistics and Philosophy* **8**:163–201.

Krifka, M., 2011. Varieties of semantic evidence. In C. Maienborn, K. von Heusinger, and P. Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning*, vol. 1, pp. 242–268. Berlin: De Gruyter Mouton.

Larson, R. and G. Segal, 1995. *Knowledge of Meaning*. Cambridge: MIT Press.

Lasersohn, P., 2000. *Same*, models, and representation. In B. Jackson and T. Matthews (eds.), *Proceedings of Semantics and Linguistic Theory X*, pp. 83–97. Ithaca, NY: Cornell University.

Lepore, E., 1983. What model-theoretic semantics cannot do. *Synthese* **54**:167–187.

Levin, B. and M. Rappaport Hovav, 2005. *Argument Realization*. Cambridge: Cambridge University Press.

Lewis, D., 1970. General semantics. *Synthese* **22**:18–67. Reprinted in Lewis (1983).

———, 1983. *Philosophical Papers*, vol. 1. Oxford: Oxford University Press.

Lidz, J., P. Pietroski, J. Halberda, and T. Hunter, 2011. Interface transparency and the psychosemantics of *most*. *Natural Language Semantics* **19**:227–256.

Lindström, P., 1966. First order predicate logic with generalized quantifiers. *Theoria* **32**:186–195.

McConnell-Ginet, S., 1973. *Comparative Constructions in English: A Syntactic and Semantic Analysis*. Ph.D. dissertation, University of Rochester.

Montague, R., 1970. Universal grammar. *Theoria* **36**:373–398. Reprinted in Montague (1974).

———, 1973. The proper treatment of quantification in ordinary English. In J. Hintikka, J. Moravcsik, and P. Suppes (eds.), *Approaches to Natural Language*, pp. 221–242. Dordrecht: Reidel. Reprinted in Montague (1974).

———, 1974. *Formal Philosophy*. New Haven: Yale University Press. Edited by R. Thomason.

Mostowski, A., 1957. On a generalization of quantifiers. *Fundamenta Mathematicae* **44**:12–36.

Partee, B. H., 1979. Semantics—mathematics or psychology? In R. Bäuerle, U. Egli, and A. von Stechow (eds.), *Semantics from Different Points of View*, pp. 1–14. Berlin: Springer-Verlag.

Pettit, D., 2002. Why knowledge is unnecessary for understanding language. *Mind* **111**:519–550.

Pietroski, P. M., 2005a. *Events and Semantic Architecture*. Oxford: Oxford University Press.

———, 2005b. Meaning before truth. In G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*, pp. 255–302. Oxford: Oxford University Press.

———, 2010. Concepts, meanings and truth: First nature, second nature, and hard work. *Mind and Language* **25**:247–278.

Pinkal, M., 1995. *Logic and Lexicon*. Dordrecht: Kluwer.

Recanati, F., 2004. *Literal Meaning*. Cambridge: Cambridge University Press.

———, 2010. *Truth-Conditional Pragmatics*. Oxford: Oxford University Press.

Rotstein, C. and Y. Winter, 2004. Total adjectives vs. partial adjectives: Scale structure and higher-order modification. *Natural Language Semantics* **12**:259–288.

Segal, G., 2006. Truth and meaning. In E. Lepore and B. C. Smith (eds.), *The Oxford Handbook of Philosophy of Language*, pp. 189–212. Oxford: Oxford University Press.

Soames, S., 1984. What is a theory of truth? *Journal of Philosophy* **81**:411–429.

———, 1992. Truth, meaning, and understanding. *Philosophical Studies* **65**:17–35.

Sperber, D. and D. Wilson, 1986. *Relevance: Communication and Cognition*. Cambridge: Harvard University Press.

Stalnaker, R., 1997. Reference and necessity. In B. Hale and C. Wright (eds.), *A Companion to the Philosophy of Language*, pp. 534–554. Oxford: Blackwell.

von Stechow, A., 1984. Comparing semantic theories of comparison. *Journal of Semantics* **3**:1–77.

Steiner, M., 1998. *The Applicability of Mathematics as a Philosophical Problem*. Cambridge: Harvard University Press.

Syrett, K., 2007. *Learning about the Structure of Scales: Adverbial Modification and the Acquisition of the Semantics of Gradable Adjectives*. Ph.D. dissertation, Northwestern University.

Syrett, K., C. Kennedy, and J. Lidz, 2010. Meaning and context in children's understanding of gradable adjectives. *Journal of Semantics* **27**:1–35.

Szabolcsi, A., 1997. Strategies for scope taking. In A. Szabolcsi (ed.), *Ways of Scope Taking*, pp. 109–154. Dordrecht: Kluwer.

———, 2010. *Quantification*. Cambridge: Cambridge University Press.

Travis, C., 1996. Meaning's role in truth. *Mind* **105**:451–466.

Westerståhl, D., 1989. Quantifiers in formal and natural languages. In D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, vol. IV, pp. 1–131. Dordrecht: Kluwer.

Zimmermann, T. E., 1999. Meaning postulates and the model-theoretic approach to natural language semantics. *Linguistics and Philosophy* **22**:529–561.

———, 2011. Model-theoretic semantics. In C. Maienborn, K. von Heusinger, and P. Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning*, vol. 1, pp. 762–802. Berlin: De Gruyter Mouton.