Semantics without Semantic Content

Daniel W. Harris*

December 23, 2016

Abstract

Propositions are expressed by people, not by sentences, or sentences relative to contexts. The meaning of a sentence provides hearers with partial evidence of the contents that speakers express with it, but isn't itself a content. The notion of "context" appealed to in compositional semantics is unexplanatory. These ideas have been defended only in the abstract. My aim is to show how they may be embodied in a detailed compositional-semantic theory. This theory is conservative, in that I reach it by way of nondestructive changes to the influential framework of Heim & Kratzer (1998) and von Fintel & Heim (2011). The result has implications about the nature of variables, indexicality, polysemy, and the distinction between functional and non-functional vocabulary. I also argue that it does a better job of locating semantics within a well motivated picture of human cognitive architecture. Specifically, I argue that semantics is the study of a modular component of the human faculty of language whereas pragmatics is the study of one application of the mind's central system(s).

My aim is to answer two questions about compositional semantics and to explore the relationship between them. First, a methodological question: what kinds of semantic values should a semantic theory assign to expressions? Second, a foundational question: what is the subject matter of compositional semantics? My answer to the first question is that the semantic value of an expression should be thought of as a constraint on what speakers can say with the expression, rather than as the content of what they say. My answer to the second question is that semantics is the study of a modular component of the mind. I will defend each of these claims, in part by showing how well they work together (§2), and in part by showing how to construct a detailed compositional-semantic theory that embodies both ideas (§3).

*For helpful feedback, thanks to David Braun, Jonathan Cohen, Kai von Fintel, Graeme Forbes, Manuel García-Carpentiero, Justin Khoo, Eric Mandelbaum, Eliot Michaelson, Stephen Neale, David Pereplyotchik, Paul Pietroski, Brian Rabern, François Recanati, Kate Ritchie, Craige Roberts, Yael Sharvit, Richard Stillman, and Elmar Unnsteinsson.

1 Two Conceptions of Semantic Values

My first order of business is to compare two ideas about the nature of semantic values, by which I mean the entities assigned to expressions by semantic theories. The two options are what I'll call *content semantics* and *constraint semantics*.

1.1 Semantic Values as Contents

Content semantics is the standard view, on which an expression's semantic value, relative to a given context, is its semantic content. This idea is implemented in different ways by different theories, in part because of disagreement about what contents are. However, the most influential compositional-semantic framework identifies the semantic value of an expression φ, relative to an assignment g and context c, with an intension $[\![\varphi]\!]^{g,c}$ (or, equivalently, with a world-relativized extension $[\![\varphi]\!]^{w,g,c}$).[1] Idealizing away from tense, aspect, clause-type, and some other complications, intensions for the sentence 'Ann smokes' and its parts can be given as follows, for example.

(1) $[\![\text{Ann}]\!]^{g,c} = \lambda w\,.\,\text{Ann}$

(2) $[\![\text{smokes}]\!]^{g,c} = \lambda w\,.\,\lambda x_e\,.\,x \text{ smokes at } w$

(3) $[\![\text{Ann smokes}]\!]^{g,c} = \lambda w\,.\,\text{Ann smokes at } w$

The central idea here is that the semantic value of a declarative sentence is a proposition, here understood as a function from worlds to truth values, and the semantic values of sub-sentential expressions are their contributions to the propositional contents of the sentences in which they appear.

The assignment and context parameters on $[\![\cdot]\!]^{g,c}$ come into play because many natural-language expressions are *semantically underspecified*. Each of the following sentences can't have any single proposition as its semantic value, for example, because it expresses (or is used to express) different propositions on different occasions.

(4) I smoke.

[1] Throughout this essay, I refer to the framework codified by Heim & Kratzer (1998) and von Fintel & Heim (2011) as 'the standard framework', and I'll use their ideas and notation (or slightly simplified versions) as a starting point. There are, of course, alternative approaches—to name just two: Davidsonian truth-theoretic semantics (Larson and Segal, 1995; Lepore and Ludwig, 2007) and variable-free semantics (Jacobson, 2014). There are interesting questions about how the morals of the present essay apply to these other frameworks, but I don't have space to explore them adequately here.

(5) The man smokes.

(6) He smokes.

The standard solution is to assume that the assignment and context parameters do some of the work in determining these sentences' contents. For any context c, there will be a speaker, $\textsc{speaker}_c$, for example, which will be the semantic value of 'I' relative to c (Kaplan, 1989b). This yields the following value for (4).

(7) $[\![\text{I smoke}]\!]^{g,c} = \lambda w\,.\ \textsc{speaker}_c \text{ smokes at } w$

This strategy has been extended in various ways to handle other expressions. It is common to posit a hidden, context-sensitive domain restrictor in the syntax of every DP, so that (8) is understood to have an LF of the following form, for example (Stanley, 2000; Stanley and Szabó, 2000; Westerståhl, 1984).

(8) $[_{\text{DP}}$ The man *dom*] smokes.

On this simplified version of the proposal, *dom* functions as an unpronounced relative clause that somehow gets its value from context. Just how contexts supply semantic values is usually left unclear, though it is widely agreed that most context-dependent expressions depend in some way on facts about the mental states of whoever is involved in the conversation. One popular view is that context-sensitive expressions' semantic values are fixed by speakers' referential intentions.[2]

Finally, deictic pronouns, such as the occurrence of 'he' in (6), are standardly taken to be sensitive to the assignment parameter. It is assumed that each pronoun is subscripted with a numerical index, and that assignments are mappings from indices to semantic values. The semantic value of a pronoun x_i, relative to an assignment g, is $g(i)$, subject to constraints imposed by x's φ-features. This gives rise to assignment-dependent sentential semantic values, such as (10).

(9) $[\![\text{He}_1 \text{ smokes}]\!]^{g,c} = \lambda w\,.\ g(1) \text{ smokes at } w$ if $g(1)$ is male; else undefined.

The assignment relative to which an expression's content is fixed is itself, according to the standard account, given by context. So, although context and assignment are often treated as distinct parameters of the interpretation function, we can follow

[2] E.g. Heim (2008); Kaplan (1989a); King (2013, 2014); Perry (2009); Schoubye (2016). This view should be distinguished from the view that although the semantic value of a semantically underspecified expression φ is not itself a content, the content that speakers express when uttering φ is a matter of their intentions. The latter view, which is compatible with the nonexistence of semantic content, and which I think is roughly correct, has been defended by Bach (1987, 1992); Neale (2004, 2005, 2016); Schiffer (1994, 2003).

Heim & Kratzer in assuming that each context, c, somehow determines an assignment, g_c, thereby collapsing the two parameters, and allowing us to represent the semantic value of 'He smokes' as follows.

(10) $[\![\text{He}_1 \text{ smokes}]\!]^c = \lambda w \,.\, g_c(1)$ smokes at w if $g_c(1)$ is male; else undefined.

What facts about a context determine which assignment is operative in it? (An answer to this question is sometimes described as a theory of the *metasemantics* of variables.) Heim & Kratzer say only that "the physical and psychological circumstances that prevail when an LF is processed will (if the utterance is felicitous) determine an assignment to all the free variables occurring in this LF" (1998, 243). In a later work, Heim instead says that "the relevant assignment is given by the utterance context and represents the speaker's referential intentions" (2008). Although the details of such an account would have to be spelled out, the idea that a deictic pronoun's referent is somehow determined by the intentions with which they speak is highly intuitive, and several philosophers have defended the view.[3] Even those who eschew intentions substitute other facts about interlocutors' propositional attitudes as the determinants of pronominal reference. Some have taken contexts to be bodies of information that are intersubjectively constructed from the beliefs or belief-like states of the participants in a conversation, for example, and so have held that it is these states that ultimately determine a pronoun's reference.[4]

Why divide off a subcategory of expressions that are indirectly sensitive to context by being sensitive to assignments? The reason is that pronouns can also have bound occurrences. For example, (6) can occur embedded in (11).

(11) $[\text{Every doctor}]_1$ denies that he_1 smokes.

In the standard framework, binding is understood in terms of compositional operations that abstract over the assignment-sensitive components of an expression, turning them into argument places in a complex predicate that can compose with the binding quantifier. As Heim & Kratzer put it, "roughly, what is meant by 'variable binding' is any semantic operation which removes (or reduces) assignment dependency" (1998, 116). A treatment on which pronouns depend on contexts by depending on assignments is thus necessary for a unified treatment of their bound and deictic (i.e. free) occurrences. Assignment-dependent expressions are therefore *variables*, unlike directly context-dependent expressions. Just where to draw the line between variables and directly context-dependent expressions is not always clear, though the crucial issue is whether they can be bound. For example, Stanley and Szabó (2000) argue that the unpronounced context-dependent expressions

[3] E.g. Bach (1987, 1992); King (2013, 2014); Neale (2004); Schiffer (1981).

[4] E.g. Heim (1982, 1983); Roberts (2002, 2003); Stalnaker (1978, 2014).

involved in domain restriction are (or include) variables on the ground that they admit of bound readings. I stress this distinction because it will become important again later: one quirk of the semantic framework I'll build in §3 is that it classifies many more expressions as variables than is standard.

1.2 Semantic Values as Constraints

Constraint semantics, as I understand it, is the view that an expression's semantic value is not a context-relative content, but a context-neutral specification of any content that can be literally expressed using the expression.[5] Roughly, the semantic value of an expression φ is just what a competent speaker can know about what someone would be saying in uttering φ, assuming they were speaking literally, but without any knowledge about the context or the speaker's intentions. An expression is semantically underspecified insofar as there is slack in the constraint that its semantic value places on the content that it can be used to literally express. If a view like this is correct, there may no need for compositional semantics to traffic in a theoretical notion of "context" at all.

I am not the first to advocate a version of constraint semantics. For example, Sperber & Wilson say that a sentence's "semantic representation is a schema, which must be completed and integrated into an assumption about the speaker's informative intention" (1995, 175). Carston likewise argues that a sentence's semantic representation "is typically not fully propositional, so does not have a determinate truth condition, but consists of an incomplete conceptual representation which functions as a schema or template for the pragmatic construction of propositional forms" (Carston, 2006, 633). Bach argues that "…the semantics of an expression gives the information that a competent speaker can glean from it independently of any context of utterance" (1987, 5). Bach has sometimes fleshed out this idea by referring to the semantic value of a semantically underspecified sentence as a "propositional radical", which he describes as a structured proposition that "lacks at least one constituent needed for it to be true or false and to be the content of a thought or a statement" (2006, 437). Neale argues that "a semantic theory for a language *L* will provide, for each sentence *X* of *L*, a *blueprint* for…what someone will be taken to be saying when using *X* to say something" (2005, 189). Schiffer argues that the semantic value of a sentence (which he calls a 'character*') is "an ordered pair ⟨A,P⟩, where A is the kind of speech act that must be performed in a literal utterance of the sentence, and P is the kind of propositional content that speech act must have" (Schiffer, 2003, 112). As Schiffer makes clear, both the force and content of a speech act will normally be underspecified by the character* of the sentence with

[5] I abstract away from the component of sentence meaning that constrains illocutionary force.

which it is performed, so that, for example, P in the above formula needn't stand in for any particular proposition, but only a rough-grained property of one (Schiffer, 2003, §3.4).[6]

Each of these ideas may remind readers of Kaplan's idea that the context-neutral component of an expression's meaning can be represented as its *character*, which Kaplan defines as a function that maps contexts in which the expression might be uttered to the expression's content at that context (1989b). Kaplan even claims that the character of a complex expression "is a function of the character of the parts" (1989b, 507), thereby suggesting that compositional semantics should be concerned with the composition of characters in addition to contents. But Kaplan's own semantics composes contents rather than characters, and he does not show how to construct a compositional character theory.[7] Kaplan's theory also differs from what I'm calling constraint semantics in two crucial respects: on Kaplan's view, contents and characters are two different kinds of semantic values, each of which participates in semantic composition.[8] By contrast, the kind of constraint semantics I'll defend deals *only* in context-neutral semantic values. Moreover, Kaplan assumes that the role of an expression's character is to determine its semantic content in a context. By contrast, I think that the idea of semantic content is a category mistake, and that the role of an expression's semantic value is to provide the hearer with incomplete evidence of what the speaker is using the expression to say.

These differences between what Kaplan calls 'characters' and what I call 'constraints' forestall an objection to constraint semantics that has been formulated by King & Stanley (2005). They argue, for example, that a compositional theory of characters would be redundant, since "the job of character is to give us content, and I can assign contents to complex expressions in contexts using only the characters of the parts, and combining the contents they determine in those contexts" (2005, 127–8). But I simply deny King & Stanley's opening premise here: semantics needn't (and shouldn't) deal in contents of any kind, but only in constraints.[9]

The biggest weakness of previous defenses of constraint semantics, I think, is that they haven't followed through on their big-picture proposals by showing how

[6] For proposals that share some elements with what I'm calling 'constraint semantics', see Barwise and Perry (1983); Swanson (2016); von Fintel and Gillies (2008). For some recent defenses of the distinction between semantic values and the contents of speech acts—albeit, for reasons rather different than my own—see Ninan (2010); Rabern (2012); Stanley (1997); Yalcin (2007).

[7] Braun (1996) shows how to give a compositional character theory, but only for a limited fragment involving complex demonstratives.

[8] For others who distinguish context-neutral and context-relative semantic values, see Stalnaker (1978) and Perry (2001). García-Carpintero (2006) defends views of this kind against objections by King and Stanley (2005).

[9] King & Stanley's other argument (2005, 127) is based on assumptions about the form that a compositional character theory would take that do not correctly describe my view in §3.

to actually construct a compositional-semantic theory. Indeed, as far as I can tell, nobody has attempted to construct a detailed compositional-semantic theory that outputs constraints rather than context-relativized contents. It should be no surprise, then, if constraint semantics is viewed as a philosophical curiosity rather than as a serious methodological proposal. I will therefore tackle this problem in detail in §3. First, in §2, I'll say more about my motivations for favoring constraint semantics.

2 Semantics and Cognitive Architecture

What aspect of the world do semantic theories aim to model? The answer I wish to advocate—and that I'll use to provide support for constraint semantics—is that compositional semantics is the study of a modular component of the human mind. Specifically, semantics aims to model humans' *semantic competence*—a body of information that is more-or-less shared by normal speakers of a common language. When a speaker produces a linguistic utterance, a mechanism in their mind draws on their semantic competence in order to encode, in the form of a sentence, partial evidence about what the speaker is saying. When a hearer perceives a linguistic utterance, a mechanism in their mind draws on their semantic competence in order to decode the evidence that is encoded in the sentence they've perceived. These mechanisms reside within the human faculty of language. Although semantic competence is thus implicated in both the production and comprehension of linguistic utterances, I'll focus on its role in linguistic comprehension—a process whose overall architecture I represent in Figure 1. I thus construe semantics as a part of linguistic theory as it is understood within the tradition of generative grammar.

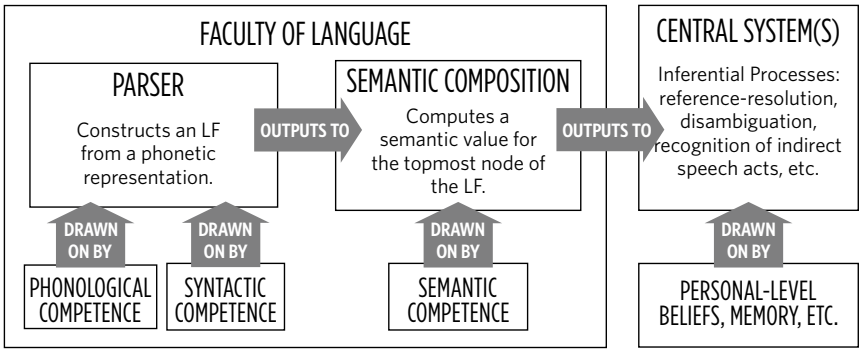

Figure 1: The Cognitive architecture of language comprehension.

Versions of psychologism about semantics view have often been assumed or

defended by semanticists and philosophers of language.[10] Indeed, although Heim & Kratzer and von Fintel & Heim devote little attention to foundational questions about the standard framework, it is relatively clear that they take some version of psychologism for granted. Heim & Kratzer name their book 'Semantics in Generative Grammar', after all, and they say that their object of study is "the semantic component" of "the grammar", thus using the terminology favored by generative grammarians to discuss the cognitive competences underlying the faculty of language (1998, ch.1). Further clues to the psychologism presupposed the standard framework are scattered throughout both textbooks. Anticipating one of the crucial premises (though not the conclusion) of my argument in this section, for example, Heim & Kratzer argue that their semantics cannot include "simple rules" governing deictic and anaphoric uses of pronouns because reference resolution is a "complex cognitive task", thus presupposing that the goal of semantics is to formulate the rules followed by the cognitive processes underlying linguistic comprehension (1998, 286).

The main alternative to psychologism is the view that the subject matter of semantics—and perhaps of linguistic theory in general—is "the structure of natural language, considered in abstraction from the cognitive mechanisms causally responsible for language acquisition and mastery" (Soames, 1984, 136). Languages, on this view, have been thought of either as abstract objects[11] or as vast collections of actual and potential token utterances.[12] The main argument for these views revolve around the idea that psychologism is needlessly speculative. Lewis argues that if psychologism is true, "whenever we resort to extrapolation to answer questions of syntax and semantics, we are engaged in risky speculation about the secret workings of the brain" (1992, 110, fn.6). A related objection proceeds from the claim that although there are, no doubt, good reasons to think that our language use must be explained by appeal to facts about language users' psychology, it does not follow that these facts line up with the kinds of grammars posited by linguistic theorists.[13]

My broad response is that while it is true that semantics is in the business of constructing models that describe the properties of human languages at a high level of abstraction, these models are be interesting only insofar as they attribute to linguistic entities the same properties to which the cognitive systems underlying language use are sensitive. This is because I take the overall goal of semantics to be a partial explanation of how language users actually do things—including how we

[10] E.g. Borg (2004, 2012); Carston (2002); Chierchia and McConnell-Ginet (1990); Glanzberg (2014); Larson and Segal (1995); Ludlow (2011); Pietroski (2008); Schiffer (2003); Sperber and Wilson (1995); Yalcin (2014).

[11] E.g., Katz (1981, 1984); Lewis (1975); Montague (1970a,b); Soames (1984, 1985, 1989).

[12] E.g., Devitt (2006, 2013); Devitt and Sterelny (1999).

[13] Cf. Devitt (2006); Devitt and Sterelny (1999); Soames (1984, 1985).

communicate—with language.[14] Moreover, I take it for granted that any such explanation will have to be given in terms of the psychological processes involved in language use and the cognitive competences on which they draw. Mere abstract description of linguistic structures won't help with this explanatory project.

A more direct form of support for psychologism comes from the fact that our best current accounts of the processes involved in language acquisition and comprehension rely on just the kind of grammatical principles posited by linguistic theory.[15] Admittedly, psycholinguistics has advanced further in the study of syntactic processing as compared to semantic processing. However, psycholinguists have recently turned their attention to semantics, and much of this work provides empirical support for the psychological reality of both semantic principles and semantic representations of linguistic inputs. Some examples: (i) MEG evidence and other neuroimaging data has been used to argue that homonymy must be represented differently than polysemy, both in the minds of language users and in semantic theories (Falkum and Vicente, 2015; Pylkkänen et al., 2006, 2002); (ii) data about the real-time sentence processing of antecedent-contained deletion constructions has been used to argue that hearers comprehend sentences by representing them as having LFs in which their quantifiers are raised (Hackl et al., 2012); (iii) a combination of corpus data and data from laboratory tasks have been used to argue that children acquire the meanings of attitude verbs in part by representing and generalizing from the syntactic and semantic properties of the verbs' compliment clauses (Hacquard, 2014; Lewis, 2013; Lewis et al., MS); and (iv) a variety of empirical evidence has been used to show that some semantic clauses for the determiner 'most' do a better job of capturing how speakers' represent its meaning than other, extensionally equivalent clauses (Hackl, 2009; Pietroski et al., 2009).

Other lines of attack on psychologism about semantics depend on the assumption that semantic values are contents. Soames' main anti-psychologistic arguments, for example, rest on the premise that "an adequate semantics for natural language must contain a theory of truth conditions that characterizes logical properties and relations such as logical truth, contradiction, entailment, and consistency" (1985, 204). From this starting point, Soames argues that semantics cannot be the study of a cognitive competence, because "a complete specification of the nonlinguistic conditions under which they are true will not follow from a specification of mental states and processes, or a description of the relationship between sensory input and behavioral output" (2009, 145). The motivation for this claim comes from semantic externalism: Soames takes Kripke and Putnam to have shown that "the se-

[14] Cf. Neale (2016) on 'the master question'.

[15] See Pereplyotchik (2017) for an overview of contemporary psycholinguistics in defense of the psychological reality of grammars and syntactic representations of linguistic inputs.

mantic values (intensions/extensions) of proper names and many general terms are determined not by descriptions speakers associate with them, but rather by socio-historical chains connecting speakers to one another and, ultimately, to entities in the nonlinguistic world" (2009, 161–2). (See also Partee (1981) for a similar argument.) But there is a relatively clear problem with this line of thought. As, e.g., Burge (1979; 1986) has stressed, arguments for externalism support externalism about the contents of intentional mental states just as well as they support externalism about semantic values. But we should presumably not therefore deny psychologism about mental states. Rather, we should conclude that facts about psychological states are themselves are partly grounded in facts about agents' physical and social environments. I would also reply that this argument assumes that semantic values are contents—a premise which I deny.

I anticipate a methodological objection to constraint semantics arising from the fact that two of the crucial kinds of evidence for semantic theories—namely, intuitions about entailment and consistency relations and intuitions about truth and falsity—appear to concern contents rather than constraints. My response to the point about entailment and consistency will come in the §3, where I will show how to define a consequence relation over my preferred sort of sentential semantic values that is equivalent to the standard consequence relation. My response to the point about truth conditions is that intuitions that appear to be about the truth and falsity of sentences are actually intuitions about the truth and falsity of what speakers are or would be *saying* in using the sentences literally. Although a sentence on its own normally encodes only a rough-grained type of proposition, hearers who are presented with a sentence make reflexive assumptions about the context in which it would have been uttered, and about the intentions with which it would have been uttered, and these assumptions bridge the epistemic gap from semantic value to content.

It may be asked: what makes it the case that a speaker says what they do in uttering a sentence, and in what sense does the meaning of a sentence constrain what a speaker can say with it? My approach is compatible with two ways of answering this question. The first is that what a speaker says in uttering a sentence is partly determined by the sentence's semantic value, which defines a space of options, and partly determined by the context of utterance or by the speaker's intentions, which select from among those options. On this view, the semantic value of a sentence *directly* constrains what can be said with it.

I am partial to a more radical view on which what a speaker says in uttering a sentence is wholly constituted by their communicative intentions. On this view, saying is just a special case of speaker meaning—namely, the case in which a speaker intends their addressee to recognize what they mean *directly*, and not by means of first recognizing something else that they mean. (The indirect case of meaning in-

cludes what Grice calls implicating.) On this view, the role of sentence meaning is to provide the hearer with incomplete and defeasible perceptual evidence about the content of the speaker's intentions. By playing this role, the meaning of a sentence *indirectly* constrains what a speaker can say with it, by constraining what they take it to be possible for the addressee to recognize them as intending. In normal circumstances, it will be rational for a speaker to believe that they can get their intention recognized only if they are providing adequate evidence of its content in the form of a meaningful sentence.[16] In this way, uttering a meaningful sentence both enables complex acts of speaker meaning and indirectly constrains their contents. I can't adequately defend this total-intentionalist view about saying in this paper, but I plan to do so elsewhere.

The role of a hearer's semantic competence, on this view, is to guide the psychological process that takes a representation of the LF of a sentence σ and turns it into a representation of the property possessed by any proposition that a speaker who used σ literally would be saying. A presumption embodied in this process is that it is occurring in the mind of an agent who has perceived a linguistic utterance, and that it is part of a larger process that aims to answer the question of what the utterer was saying. But a hearer's semantic competence cannot offer a complete answer to this question; it can only narrow the space of interpretive options. For a complete answer, a hearer must also rely on their *pragmatic* abilities—their capacity to infer the speaker's thoughts.

Semantic underspecification, on this view, is the gap between what a speaker says when they use an expression literally and what a hearer's semantic module can tell them about what the speaker has said. I'll call the process of bridging this gap *reference resolution*. (This is a wider-than-usual use of the term, since expressions other than referring expressions might be semantically underspecified.) If we grant that semantics is the study of semantic competence, then the central disagreement between content semantics and constraint semantics is about the order in which reference resolution and semantic composition take place within the overall process of interpreting what a speaker has said by uttering a sentence.

According to content semanticists, reference resolution has to happen *before* semantic composition. This is because the semantic values that compose with one another, according to the standard framework, are themselves contents. But before the contents of 'I' and 'smoke' can be composed to yield the content of 'I smoke', on this view, the contents of 'I' and 'smoke' must be identified. According to constraint semanticists, on the other hand, reference resolution comes *after* semantic compo-

[16] For some variations on the idea that our intentions are constrained by what we take to be possible, see, e.g., Bratman (1987); Broome (2013); Donnellan (1968); Grice (1971); Holton (2011); Neale (2004).

sition. On this view, reference resolution is a pragmatic process that is required to get from a representation of a sentence's semantic value, which is an incomplete specification of what the speaker said, to a full representation of what they said.

Why think that the second picture captures the cognitive architecture of semantics better than the first? My reason is based on independently motivated views about how the human mind is organized. My argument, in brief, runs as follows.

The Argument from Cognitive Architecture

(P1) Semantic composition is a modular process.

(P2) Reference-resolution is a central process.

(P3) Modular processes output to central processes, not the other way around.

(C) Therefore, reference-resolution cannot serve as an input to semantic composition.

I use the terms 'modular' and 'central' in Fodor's (1983) sense.[17] A module, in this sense, is a domain-specific perceptual input system that draws on a proprietary database of information in order to quickly, automatically, and algorithmically construct representations of perceptual inputs before passing them off to the central system(s). Central processes, such as belief-fixation, abductive (non-demonstrative) inference, and analogical reasoning, by contrast, are domain-neutral, sometimes (though not necessarily always) slow, effortful, and conscious, and have access, in principle if not always in practice, to all of an agent's personal-level beliefs and memories.[18]

The crucial properties of modular processes, for my purposes, is that they are *informationally encapsulated* and *centrally inaccessible*.[19] To say that they are informationally encapsulated is to say that modules have access only to a proprietary and finitely axiomatizable body of information on which to draw in doing their work. In particular, modules don't, as a rule, have access to the information that is available to

[17] Crucially, I *don't* use 'modular' in the sense of "massive modularity" (Carruthers, 2006; Sperber and Wilson, 2002), according to which there are special-purpose mental "modules" with distinctive evolutionary functions whose processes are isotropic and so not informationally encapsulated. I am neutral about whether the central system is divided into modules in this sense, which is why I sometimes refer to the 'central system(s)'.

[18] I won't defend of a Fodorian picture of mental architecture here. For the classic defense, see Fodor (1983). For recent defenses, see Firestone and Scholl (2015) and Mandelbaum (2017).

[19] Some other properties commonly attributed to modules—such as that they are implemented by characteristic neural architecture with characteristic breakdown patterns and common phylogenetic and ontogenetic origins—will not concern me here, and what I'll have to say will be neutral about whether semantic composition has these qualities.

central systems. This encapsulation needn't be hard and fast: central processes may sometimes influence the workings of some more-or-less modular systems to some extent, resulting in "top-down effects" (a.k.a. "cognitive penetration"). (Fodor takes modularity to be gradable (1983, 37).) But, in general, and as a rule, modular processes don't take inputs from central processes. To say that a modular process is centrally inaccessible is to say that central processes do not have access to either the proprietary database on which the module draws or to the intermediate stages of the representations computed by modular processes. Modules interface with central systems in only limited ways—most notably, by sending outputs to them. So, although modules' outputs have to be the sorts of representations that central systems know what to do with, modules may get to these outputs via intermediate representations that aren't centrally accessible, and that central systems wouldn't know what to do with if they were (for example, because they are framed in terms of concepts that central systems lack). A paradigmatic example of a modular process is syntactic parsing, which psycholinguists have modeled as a process that draws on syntactic competence, a proprietary database of centrally inaccessible principles that syntacticians have made impressive progress toward reverse-engineering, and that employs a battery of concepts (e.g. DETERMINER PHRASE) not generally possessed by the central systems of language users.

By contrast, the crucial property of central processes, for my purposes, is that they are *isotropic* (Fodor, 1983, 105). This is to say that any information available to the agent, including any of their beliefs or memories, may, in principle, be drawn on by a central process. Fodor illustrates this point by analogy to scientific confirmation, which he sees as a public and deliberate analog of central cognitive processes.

> By saying that confirmation is isotropic, I mean that the facts relevant to the confirmation of a scientific hypothesis may be drawn from anywhere in the field of previously established empirical (or, of course, demonstrative) truths. Crudely: everything that the scientist knows is, in principle, relevant to determining what else he ought to believe. In principle , our botany constrains our astronomy, if only we could think of ways to make them connect. (Fodor, 1983, 105)

Mindreading—the process of inferring other agents' intentional mental states—is a paradigmatic central process. In drawing conclusions about the state of another agent's mind in order to explain their behavior, any of my beliefs or memories could turn out to be relevant. What I knows about Dungeons and Dragons and interest rates constrains what I will (and should) think about why someone chooses to dance in the way that they do, if only I can think of ways to make them connect.

Semantic composition—the process by which hearers compute sentence meanings from the meanings of their parts—bears the hallmarks of a modular process.

It is domain specific, fast, involuntary, and unconscious. It is also centrally inaccessible: we have no personal-level access to the principles that guide it, or two the intermediate representations it computes. This is why constructing semantic theories is so hard: it requires reverse-engineering a black box whose inner workings cannot be introspected. There are also good reasons to think that semantic composition is informationally encapsulated, in that it is insensitive to the influence of the central system(s). We can't help but perceive sentence-like perceptual inputs as meaningful, even when we know that they do not have the right etiology to be genuinely meaningful. Consider, for example, the trusty old case of stones on a beach, arranged by a storm into the shape of a sentence. In this case, there is no utterance to interpret, and the agent may realize this, but the stones are sufficiently similar to the evidence usually left behind by genuine utterances that the faculty of language fires up and construes the "sentence" as having a meaning just the same. I take this phenomenon to be a genuine perceptual illusion wherein an agent's faculty of language persists in attributing properties to a stimulus that, the agent knows full well, the stimulus does not have. I'll say more about perceptual illusions of this kind in §3.

It is because semantic composition is a modular process—plausibly, the faculty of lannguage's last task before outputting to the central system(s)—that semantic competence can be represented as a finite collection of principles. Semantic competence turns LFs into semantic representations by drawing on a proprietary database of lexical rules and compositional principles.

Indeed, the only thing unmodular about semantic composition as it is standardly understood is that it depends on input from reference resolution. Reference resolution is an isotropic and, so, paradigmatically central process. Anything a hearer knows may turn out to be relevant to how they work out the content that a speaker has expressed by using a semantically underspecified expression literally. This is a well-worn point when it comes to demonstratives, domain restriction, and many other instances of semantic underspecification. But it is in fact true of *all* instances of semantic underspecification.

Suppose, for example, that Jim bursts into Carol's office and shouts, 'he's drinking again!'. In order to figure out whom Jim is talking about, it might help Carol to know any number of biographical details about Jim, his acquaintances, the celebrities he admires, and who knows what else. In particular, working out whom Jim is talking about will likely involve mindreading: information available to Carol will be relevant to interpreting Jim's utterance only insofar as this information bears on Jim's possible reasons for telling Carol about one person's drinking rather than another. If Carol believes that Jim would not wish to reveal compromising personal information about his family members, for example, this would give Carol good reason to assume that Jim is not referring to his brother. On the other hand, if Carol

thinks that Jim has no interest in the drinking habits of Hollywood's leading men, this gives her good reason to conclude that Jim isn't talking about Brad Pitt.

The same goes even for so-called 'automatic indexicals', such as 'I'. To see this, imagine a middle-school principal who finds 'I smoke' spray-painted on the school's bathroom wall, and consider the various information that they might have to draw on in order to identify the speaker. Again, this inference may involve mindreading: the principal's best strategy for identifying the speaker may be to reason about their motives—to answer the question, 'why would someone have written this here?'. Even when one can see an utterance being produced, knowing who the speaker is may require isotropoic inference—for example, if one knows that one of the candidate speakers in one's visual field is a skilled ventriloquist. Analogous points go for all other kinds of reference resolution.

I am not the first to give a version of the argument from cognitive architecture. One previous version of it is due to Emma Borg, who uses it to defend her version of semantic minimalism (Borg 2004, ch.2; 2012). Borg characterizes minimalism as the view that the semantic value of a sentence is a minimal proposition—a proposition derived with input from context only when that input is required by syntactically overt context-sensitive elements. There are several problems with Borg's version of the argument, as I see it. One is that there is excellent evidence for much more semantic underspecification than she recognizes. However, the fundamental problem with Borg's view, as I see it, is that reference resolution is *never* the sort of task that a modular input system can perform, even in the case of the simplest overt indexicals. And even her minimal version of content semantics requires that semantic composition get at least *some* inputs from reference resolution.

Borg is aware of this problem, and she attempts to solve it by showing how it is possible for a hearer's semantic module to identify the referent of a semantically underspecified expression without having to reason about the speaker's intentions. Suppose that S utters a sentence ⌜That is F⌝, and thereby expresses the proposition that A is F, for example. Borg argues that although the semantic value of 'that' is A in virtue of facts about the speaker's intentions, still, the hearer's semantic module may identify the semantic content of 'that' under an actualized metalinguistic description that can be formulated via a modular process. The hearer, according to Borg, "can entertain this content even if she is in no position to non-linguistically identify A; that is to say, if she is capable of thinking of A only as *the actual object referred to by the speaker with this token of 'that'* " (2012). Generalized, Borg's proposal would seem to be that reference resolution requires only formulating an actualized metalinguistic definite description of the bit of semantic content in question.

I see a clear problem with this idea.[20] In general, formulating a metalinguis-

[20] There are also other problems. For example, Soames (2002, 40–50) argues that referring expres-

tic description of the content that a speaker has expressed with an expression not sufficient for understanding a use of the expression. Consider two examples. First, suppose that I hear my teenage cousin utter (12).

(12) Kanye's kicks are on fleek.

I'm not familiar with the expression 'on fleek', and so I form a belief that I might express as follows:

(13) My cousin said that Kanye West's sneakers have the property that my cousin actually expressed with her utterance of 'on fleek'.

I think it's clear that I should not be counted as having understood what my cousin said, and that this is because formulating my actualized metalinguistic description of the content my cousin expressed with 'on fleek' is no substitute for grasping the property she was expressing with it on this occasion. Now consider an even more extreme example. Wandering into my favorite Polish restaurant, I overhear one waitress utter (14) to another:

(14) Wegetariańska jest tutaj.

I don't speak Polish, and so I form a belief that I would express as follows:

(15) The waitress said what she actually expressed with her utterance of 'Wegetariańska jest tutaj'.

In this case, I have formed an accurate metalinguistic description of the content of the waitress's utterance. And yet it is abundantly clear that I have not understood her utterance or grasped the content of what she said. But given that formulating actualized metalinguistic descriptions of contents is not sufficient for grasping those contents in cases like these, I can't see why formulating an actualized metalinguistic description of the content of a demonstrative would be sufficient for grasping its content, either.

The real source of these problems with Borg's position, as I see it, is her assumption that the outputs of semantics must be contents of some kind. This same assumption has led others who are sensitive to considerations about mental architecture to conclude either that compositional semantics, as usually conceived, is impossible (e.g. Fodor 1989; 1998; 2002), or to give up modularity and conclude that

sions cannot always be substituted for co-designating actualized definite descriptions without changing the meaning of the whole sentence.

content composition is actually an isotropic, central, and therefore *pragmatic* process (Recanati, 2010). But again, each of these views is rooted in the erroneous assumption that semantic composition must deal in contents rather than constraints.

The moral that I draw from the argument from cognitive architecture is, instead, that semantics should deliver constraints rather than contents. This view allows for a precise semantics–pragmatics interface that is grounded in an independently motivated theory about human cognitive architecture. So long as we confine our attention to utterance interpretation and bracket, for the moment, utterance production, the present view suggests that the semantics–pragmatics interface is just with the interface between the modular and central processes involved in interpreting linguistic utterances. The semantics–pragmatics interface is thus also an instance of the perception–cognition interface.

What remains, then, is to show how to work out the precise formal details of a constraint semantics. I turn to that task now.

3 Compositional Semantics

My goal in this remaining section is to formulate a compositional theory of constraints by making minimal adjustments to the standard framework without giving up on what makes this framework so successful.[21] I will proceed in stages, first showing how to give a compositional semantics and logic of the overall kind I am after, then showing how to refine this theory to handle a wide range of semantically underspecified expressions. I will also point out some explanatory bonuses along the way.

3.1 First Steps

The semantic value of a sentence, I have been arguing, is something that specifies the range of things that can be said by using the sentence literally. How should a semantic value of this kind be represented?

Before presenting my preferred answer to this question, I will consider an option that represents even less of a departure from the standard framework. Let's start with a slightly simplified standard semantic value for 'it stinks'.

(16) $[\![\text{it}_1 \text{ stinks}]\!]^c = \lambda w \,.\, g_c(1) \text{ stinks at } w$

[21] The framework presented here also inherits some limitations in its explanatory scope from the standard framework as regards, e.g., presupposition and unbound anaphora. There isn't space for a detailed discussion of these issues here; I plan to address them in future work.

(16) is a content only given a context c that determines an assignment g that in turn determines a referent for 'it_1'. But we can remove this context dependency by abstracting over c and removing it as a parameter of interpretation, as follows.

(17) $[\![\text{it}_1 \text{ stinks}]\!] = \lambda c \,.\, \lambda w \,.\, g_c(1)$ stinks at w

This can be thought of as the sentence's Kaplanian character—a function from contexts to contents. So, in keeping with this simple move, here is a simple hypothesis: the last step in the computation of any sentence's semantic value is to abstract over the context parameter in this way. The semantic module's output, given a sentence's LF as input, is the sentence's character. Having received a sentence's character from the semantic module, the central system must identify the context in which it was uttered and apply the character to it in order to identify the content. On this view, reference resolution is just the process of identifying the right context and plugging it into a sentence's character in this way—something that happens after the semantic module has done its work. Of course, the semantic module will have to work with representations of contexts and assignments at intermediate stages of its derivation, up until the point at which they are abstracted away. But these can be mere placeholders with dummy values, and so needn't incorporate any information to which the module lacks access.

This is an attractive picture, and my positive proposal will resemble it in several ways. However, I think it can't be quite correct. My reason is that this proposal requires language users' central systems to be capable of representing assignment functions and numerical indices, and I doubt that the central systems of those who haven't taken a semantics class can wield these concepts. For example, if my central system is given (17), and if contexts of utterance are the sorts of things that can be mapped to propositions by (17), then identifying a context of utterance for 'it stinks' and deriving a proposition from it will require me to have beliefs (or other personal-level propositional attitudes) about the assignment function g that is determined by that context and to work out what g assigns to the numerical index 1. But it would be bizarre to conclude that being a competent language user requires having beliefs about assignment functions and the indices in their domain. After all: the idea that human language processing traffics in assignments is at best a working hypothesis of contemporary linguistics—one that resulted from the sort of reverse engineering that is characteristic of black-box investigations, and one that is denied by proponents of competing views (e.g., Jacobson 2014). Assuming that representations of variable assignments play a role in semantic processing, then, we should conclude that they feature only in centrally inaccessible intermediate stages of representation within the language module.

It might be protested at this point that I am taking the appearance of assignment functions in semantic values like (17) too literally. After all: Heim says that

the assignment relative to which a sentence's content is determined "*represents* the speaker's referential intentions", not that speakers have intentions *about* either assignment functions or numerical indices (2008, emphasis added). Perhaps we should read the instance of $g(1)$ in (17) not as 'the value of the assignment function g for 1', but as a kind of shorthand for a description of the speaker's referential intentions with respect to the relevant occurrence of 'it'. We could say, for example, that resolving the referent of an occurrence of 'it' requires only that the hearer recognize that the speaker intended this occurrence of 'it' to refer to a particular referent x (cf. King 2013; 2014). Assignment functions are formal stand-ins for facts about the speakers' referential intentions, on this view, and numerical indices are formal stand-ins for facts about speakers' intentions about which expressions are anaphorically connected.

Articulating this view adequately would require saying more than I have said here, and perhaps that could be done. But I think this solution would raise a new and closely related problem, because it would require both speakers and hearers to have centrally accessible representations (intentions and/or beliefs) about every unbound variable in every sentence they use. But there are good reasons to think that we lack central access to representations of at least some variables—namely, phonologically null variables such as PRO, the domain-restriction variables in the nominals of DPs posited by, e.g., Stanley and Szabó (2000), and the degree variables in positive-form gradable-adjective constructions posited by, e.g., Kennedy (1999, 2007). If understanding speech does require representing these null variables, then the representations in question would seem to be centrally inaccessible. We have no ability to introspect these representations, and, if they exist, then even speakers who believe that they don't exist, such as proponents of variable-free semantics and skeptics about null categories, represent sentences as containing them nonetheless. A plausible explanation of this fact is that null variables are represented only at intermediate levels of representation within the language module, and do not feature in the module's outputs. But in that case, our strategy for reinterpreting (17) in terms of speakers' referential intentions leads us to yet another inadequate mixture of central and modular processing.

My positive proposal is an attempt to capture the virtues of the foregoing account while positing sentential semantic values of a kind that the central system(s) plausibly have the conceptual resources to handle. Specifically, I will assume that the outputs of semantics include representations of individuals, possible worlds, properties and relations of various semantic types, and propositions, but not of assignment functions, indices, or variables themselves. I will take sentential semantic values to be properties of propositions, which I represent as type-$\langle st, t\rangle$ functions. For example, a rough first shot at the semantic value of 'it stinks' can be given as

follows:[22]

(18) $[\![\text{it stinks}]\!] = \lambda p_{st} \,.\, (\exists x)(p = \lambda w \,.\, x \text{ stinks at } w)$

In English: the semantic value of 'it stinks' is a property possessed by any proposition p such that, for some x, p is the proposition that x stinks.[23]

My plan in this section is to show how to derive semantic values of this kind from sentential semantic values of the kind outputted by standard theories. To do this, I'll posit an extra layer of syntactic structure at the top of each complete LF, as follows.

The S node here is just the usual sort of sentential node, with the usual sort of semantic value, as understood by the standard framework. The SA node—'SA' is short for 'Sentence Abstract'—can be thought of as a capstone that goes at the top of every complete LF. Although I think there may be reasons to think that SA nodes sometimes embed, I will ignore this possibility here, and assume that SA nodes are never found anywhere lower down in phrase structures; they don't embed, even when S nodes do.[24] The λp node is a binding operator whose semantics I'll explain momentarily.

It will be helpful to work with an example. Consider, again, 'it stinks', plugging in its standard intension at the S node and the semantic value I wish to derive for it at the SA node in my schematic LF.

[22] Among the things that this semantic value abstracts away from are clause-type, tense, aspect, and the φ-features of 'it'. I'll show how to build the φ-features back in in §3.3.

[23] Some readers may find it more intuitive to think of the semantic value of a sentence as the set of propositions that can be literally expressed with it. The following is equivalent to (18) for example.

(18)* $[\![\text{it stinks}]\!] = \{p : (\exists x)(p = \lambda w \,.\, x \text{ stinks at } w)\}$

The intuitive idea here is that when I hear someone utter 'it stinks', my semantic module tells me that what they said is a proposition in this set (assuming they were speaking literally). Provided they are speaking literally, I have to pragmatically infer *which* of these propositions is the content of what they said.

[24] If SA nodes really never do embed, there may be no good reason to posit this extra layer of structure in object language LFs. In that case, we can implement the present proposal differently, as the final operation that the semantic module performs on each semantic representation before sending it to central systems, not because the LF demands it, but just because doing so is always part of the module's job. These implementation details won't matter for present purposes.

(19)

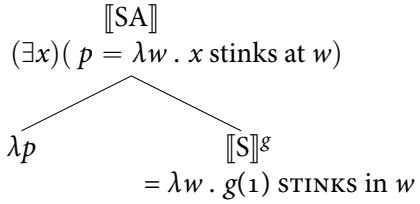

The question, then, is how to understand the semantic role of the λp node, which, by combining with $[\![\text{S}]\!]^g$, yields $[\![\text{SA}]\!]$. This operation has to take the intension of the daughter S node and add what I'll call a *prenex* to it. For any S node, the corresponding SA node's semantic value will take the following form.

(20) $\underbrace{\lambda p_{st} \,.\, (\exists x^1) \ldots (\exists x^n)\ p =}_{\text{prenex}} \varphi$

I'll use the term 'prenex quantifiers' for the existential quantifiers contained within a prenex. For each i to which at least one free variable in S is indexed, there will be a quantifier in the prenex of $[\![\text{SA}]\!]$ to bind variables indexed with i.

The desired result can be gotten by defining a new compositional principle for nodes of this kind.[25]

(21) *Proposition Abstraction* (PropA)
Let α be a branching node with daughters β and γ, where (a) β dominates only λp, and (b) γ contains unbound variables $v_i \ldots v_n$.
Then $[\![\alpha]\!] = \lambda p_{st} \,.\, (\exists x^i) \ldots (\exists x^n) \,.\, p = [\![\gamma]\!]^{g^{x^i/i \ldots x^n/n}}$

This principle, which is modeled after the Predicate Abstraction principles posited in the standard framework, is a one-stop method for turning the sentential semantic values of standard intensional semantics into semantic values of the kind I want to assign to SA nodes.

Already, it can be seen that the role of assignment functions has been substantially reduced and demystified. For one thing, SA nodes' semantic values aren't relativized to assignment functions, or contexts, or anything else. This means that "context" (however that notion is construed) doesn't play any role in this compositional semantics. Assignment functions still play a role at intermediate levels of representation. But since every variable is eventually bound, either within the S node or at the SA node, the role of assignment functions is merely to coordinate binding relations. In fact, it doesn't matter what is in the range of assignment functions, since they are always discharged at the top of a tree. It would be fine, for

[25] Notation: subscripts on variables are numerical indices, as usual. Superscripts on variables aren't indices but merely devices for disambiguating variables.

example, if all unmodified assignments were reflexive, simply mapping numerical indices to themselves. The values they assign are mere placeholders.

3.2 Consequence

Given what I've said so far, it is straightforward to define a notion of logical consequence that relates SA nodes that is equivalent to the standard framework's consequence relation. Let '$\vdash$' be the standard consequence relation and $\models$ be the consequence relation on SA nodes. The latter can be defined as follows:

(22) Let $SA^1 \ldots SA^m, SA^n$ be SA nodes whose S daughters are, respectively, $S^1 \ldots S^m, S^n$. Then:
$SA^1 \ldots SA^m \models SA^n =_{df.}$
$S^1 \ldots S^m \vdash S^n$ iff $(\forall g)(\forall w : [\![S^1]\!]^g(w) = 1 \ldots [\![S^m]\!]^g(w) = 1)\ [\![S^n]\!]^g(w) = 1$

3.3 The Semantics Refined

The semantics I've sketched so far is on the right track, but it is still underdeveloped in various ways. Most importantly, it remains unclear how to distinguish the semantic contributions of different semantically underspecified expressions. In this section, I'll make some progress on this problem.

I will begin with the relatively straightforward case study of 'he' and 'she'. As I illustrated in §1, the standard way of capturing the semantic difference between these two expressions is by restricting which assignment functions their semantic values may depend on: for any n and g, $[\![\text{he}_n]\!]^g$ is defined only if g maps n to a male, for example. But since the semantic import of binding a variable is to eliminate its assignment dependence, this semantic difference between 'he' and 'she' applies only when they are unbound. In the alternative semantic framework that I have been proposing, all variables are eventually bound, in one way or another. So I need a different way of distinguishing 'he' from 'she'.

I can get a sense of what is needed by considering what the semantics should say about sentences containing these expressions. If all I know about your utterance is that you used 'he smokes' literally, what can I know about the proposition you've expressed? Roughly, I know that you've expressed a proposition p with the following property: for some male individual x, p is the proposition that x smokes. *Mutatis mutandis* in the case of 'she smokes'. This idea can be captured by assigning the two sentences the following semantic values.

(23) $[\![\text{he smokes}]\!] = \lambda p_{st}\ .\ (\exists x_e : x \text{ is male})(p = \lambda w\ .\ x \text{ smokes at } w)$

(24) $[\![\text{she smokes}]\!] = \lambda p_{st}\ .\ (\exists x_e : x \text{ is female})(p = \lambda w\ .\ x \text{ smokes at } w)$

The trick, then is to find a way to distinguish the semantic values of 'he' and 'she' in a way that allows them to contribute different restrictions to the prenex quantifiers that bind them in the semantic values of SA nodes in which they occur. The prenex quantifier that binds a variable must be restricted by what I will call the variable's *constraint property*—the property that an entity has to have in order for a speaker to use the variable to refer to it. The constraint property of 'he' and 'she' are the properties of being male and female, respectively, for example. A variable's constraint property is what one can know, solely by virtue of semantic competence, about what someone who is speaking literally has used the variable to talk about. In effect, a variable's constraint property is how the variable's meaning is represented in the semantics.

My semantics is designed to formalize these ideas about variables' meanings. I'll say that, for any variable v, $\mu(v)$ is v's constraint property. Here are the constraint properties of 'he' and 'she', for example:

(25) $\mu(\text{he}) = \lambda x_e \,.\, x$ is male

(26) $\mu(\text{she}) = \lambda x_e \,.\, x$ is female

Because variables' constraint properties are part of what one must have internalized in order to be a competent speaker, I adopt the following compositional principle that expands on the role of the lexicon as it applies to variables.

(27) *Variables*
If v is a variable, $\mu(v)$ is specified in the lexicon.

This is not a radical departure from the standard framework, which builds roughly the same information into the semantic values of pronouns in the form of presuppositional requirements on admissible assignment functions. On this view, the information that is often taken to be encoded as existence, person, and number presuppositions of pronouns can be thought of, instead, as part of the instructions for interpreting their utterance that speakers give to addressees. When I say 'he smokes', it is not part of the content of what I say that the referent 'he' exists and is male; rather, this information is a clue I give my addressee in order to help them to recover my intended content.[26]

As I'll set things up, the constraint property of a variable v plays no role in determining $[\![v]\!]^g$, for any assignment g. Indeed, since I need no longer distinguish the semantic values of different variables in terms of different presuppositional constraints, the following lexical entry suffices for every variable:

[26] I remain neutral about the correct account of other presupposition triggers.

(28) For any variable v and any assignment g, $[\![v_n]\!]^g = g(n)$

My theory thus posits two different kinds of lexical entries. For each variable v and assignment g, the lexicon specifies an assignment-relative placeholder value of $[\![v]\!]^g$ and a substantive value for $\mu(v)$. For each non-variable φ, the lexicon specifies a substantive value for $[\![\varphi]\!]$. Think of $[\![\cdot]\!]$ and $\mu(\cdot)$ as yielding two different kinds of semantic values that play different roles within the overall system.

Variables' constraint properties enter into the semantics via an updated proposition-abstraction rule, which now installs each variable's constraint property as a restriction on the prenex quantifier that binds it.

(29) *Proposition Abstraction*
Let α be a branching node with daughters β and γ, where (a) β dominates only λp, and (b) γ contains unbound variables $u_i \ldots v_n$. Then $[\![\alpha]\!] = \lambda p_{st}$. $(\exists x^i : \mu(u)(x^i)) \ldots (\exists x^n : \mu(v)(x^n))(p = [\![\gamma]\!]^{g^{x^i/i \ldots x^n/n}})$.

This delivers the correct semantic values for the sentences (23) and (24), for example.

It is straightforward to handle a range of other variables within this framework. For any variable v, the semantics need only include a lexical entry assigning its constraint property. In the case of some variables, such as traces left behind by movement, the constraint property will be nearly vacuous, consisting only of a specification of the variable's semantic type.

(30) $\mu(t) = \lambda x_e$. x

It may also make sense to admit of some variables with semantic types other than e. The type-et domain restrictor variable posited in §1 can be given the following constraint property, for example.

(31) $\mu(dom) = \lambda x_{et}$. x

Nothing beyond the semantic type of *dom* is specified by its constraint property because, so far as I can tell, the only constraints on how quantifiers may be restricted are pragmatic rather than semantic.

3.4 Indexicals

What about other semantically underspecified expressions, such as the indexicals, 'I' and 'you', and the demonstratives, 'this' and 'that'? Following Kaplan (1989b), most philosophers take these expressions' semantic values to depend on the context parameter rather than the assignment parameter. The theory I've outlined so far

posits no context parameter, and for the reasons given earlier, I would like to keep it that way. I will therefore further revise the standard framework.

My strategy will be to treat indexicals as yet more variables, much like the pronouns I've discussed so far. Although this strategy departs from the standard practice among most philosophers, Heim & Kratzer treat first and second-person pronouns as variables that are distinguished from third-person pronouns only by their φ-features. I'll thus be agreeing with them in one respect, while departing from their view that variables' semantic values are assignment-dependent (and, so, indirectly context-dependent) contents.[27]

On my view, indexicals and demonstratives differ semantically from other variables only with respect to their constraint properties. The constraint properties of indexicals are special because they involve a kind of token-reflexivity, in the sense that they make reference to properties of "the utterance being interpreted". I can spell out the constraint properties of a variety of indexicals as follows, for example.

(32) μ(I) = λx_e . x is the speaker of the utterance being interpreted

(33) μ(you) = λx_e . x is the addressee of the utterance being interpreted[28]

(34) μ(today) = λx_e . x is the day on which the utterance being interpreted was produced

(35) μ(now) = λx_e . x is an interval of time during which the utterance being interpreted was produced

(36) μ(here) = λx_e . x is a location in space within which the utterance being interpreted was produced[29]

[27] One reason to treat indexicals as variables is that they have bound occurrences, as in the relevant readings of the following examples.

(i) [Each of you]$_1$ believes that you$_1$ are right.

(ii) Whenever a pianist comes to visit, we have to play duets. (Nunberg, 1993; Partee, 1989)

(iii) Only I got a question that I understood. (Kratzer, 1998; Rullmann, 2004)

In order to be bound by the usual mechanisms, these expressions must be variables with numerical indices. An alternative view, motivated in part by the grammaticality of phrases such as 'we Canadians' and 'you people', is to treat pronouns as the determiners of DPs whose nominals are often aphonic, and which themselves may contain bound variables (Abney, 1987; Elbourne, 2005; Neale, 2004; Postal, 1969). On this view, pronouns' semantic underspecification would turn out to be a special case of domain restriction.

[28] This entry deals only with singular 'you', and would need to be generalized in some way to handle plural senses of 'you'.

[29] There are uses of 'here' not captured by this constraint property, such as when one points to a map and says 'I'll meet you here'. It is unclear whether it is desirable to try to unify these senses.

It might be objected that, by including the description, 'the utterance being interpreted', in these constraint properties, I am sneaking reference to the context of utterance, in something like Kaplan's sense, back into my semantics. But the token-reflexive descriptions featured in indexicals' constraint properties function rather differently than the context-sensitive semantic values of Kaplan. In Kaplanian semantics, an indexical's semantic value is identical to its referent, and the context is a collection of entities that serve as the referents of indexicals. The semantic value of 'I', relative to a context c, is $\textsc{speaker}_c$, which is just the speaker themself. But, as I argued in §2, the semantic module doesn't "know" who the speaker is. Identifying the speaker of an utterance is an isotropic and pragmatic process, not an algorithmic and semantic one. Since the semantic module doesn't "know" who the speaker is, it can't plug the speaker in as the semantic value of an expression.

Instead, the occurrences of 'the utterance being interpreted' in (32)–(36) should be understood as attributive definite descriptions in the metalanguage in which the outputs of the semantic module are framed. When it operates, the semantic module presumes that there is a unique utterance that it is interpreting, and it uses 'the utterance being interpreted' to attributively pick out this utterance. It hands this description upstairs to the central system(s), one of whose jobs is to identify the description's denotation.

In some cases, the semantic module's presumption of a unique utterance that it is interpreting will turn out to be mistaken. I discussed the example of stones on a beach in §2, which, I argued, is a kind of perceptual illusion. Illusions of this kind take place because an agent's faculty of language, being relatively invulnerable to cognitive penetration, persists in attributing properties to a stimulus that, the agent knows full well, the stimulus does not have. One property that the semantic module might thus attribute is the property of standing in thus-and-such a relation to 'the utterance being interpreted' when there is, in fact, no such utterance, but only a stimulus that superficially resembles the product of an utterance.

Does this mean that indexicals aren't, as Kaplan thought, directly referential? If to say that an expression is directly referential is to say that its semantic value is its referent, then no, indexicals aren't directly referential. However, if I tweak the definition by saying that an expression is directly referential if and only if, whenever a speaker uses it literally, the content of what they say is an object-dependent proposition, then indexicals count as directly referential on the view I've given.

What about the demonstratives, 'this' and 'that'? Some have recently challenged the orthodoxy that demonstratives are unstructured referring expressions, instead claiming that they are determiners that sometimes appear with aphonic nominals (Elbourne, 2005; Hawthorne and Manley, 2012; Neale, 2004). If this view is correct, then the semantic underspecification of demonstratives can be treated as merely a special case of domain restriction, and the silent nominal of a bare use of 'that' works

just like the domain-restrictor variable *dom* hypothesized earlier.

If 'this' and 'that', in their "bare" uses, *are* unstructured referring expressions, however, then their constraint properties can be given roughly as follows.[30]

(37) μ(this) = λx_e . e is "proximal" from the perspective of the speaker of the utterance being interpreted

(38) μ(that) = λx_e . e is "distal" from the perspective of the speaker of the utterance being interpreted

I place 'proximal' and 'distal' in scare quotes here to signal that they are unexplicated technical terms—mere placeholders until someone comes up with a substantive account of the difference in these words' meanings. 'Distal' and 'proximal' can't be understood in purely spatial ways, for example, since we can use 'this' and 'that' to talk about abstracta. Nonetheless, there is *some* difference in the constraint properties of 'this' and 'that', and this difference seems to have *something* to do with proximity and distance in their ordinary senses. It is normally felicitous to say 'this is a nice apartment', but not 'that is a nice apartment', about an apartment in which one is sitting, for example.

3.5 Ambiguity

Suppose that you hear someone utter (39):

(39) John is at the bank.

Working out whether John is saying that he is at the financial institution or the riverside is an isotropic, central process, and not something that your semantic module can do for you.

There are three ways in which the framework I've been presenting might handle ambiguity. The first would be to assume that there are two distinct lexical items that are pronounced /bank/ in English. On this view, disambiguation is a job for the syntactic parser, which must choose which of the two gets plugged into an LF when it is given /bank/ as an input. This is a little bit like the job that the parser has to do in resolving syntactic ambiguities. In either case, since the parser is informationally encapsulated, it can't draw on centrally available information, and so its choice must be either arbitrary or based on a relatively unintelligent heuristic. If it chooses wrong the first time, it will be made to try again. (Alternatively, it may pursue both hypotheses in parallel.)

[30] For arguments that demonstratives ought to be thought of as variables, see Nowak (2016, MS).

There is psycholinguistic evidence that this is the right way to think about syntactic disambiguation. One well known piece of evidence involves the following two sentences.

(40) The spy saw the cop with a revolver, but the cop didn't see him.

(41) The spy saw the cop with the binoculars, but the cop didn't see him.

Rayner et al. (1983) found that subjects who read these two sentences fixated their gaze on the space immediately following the word 'revolver' in (40). The widely accepted explanation is that this pause in reading is due to a process of syntactic reanalysis: the subject's parser initially analyzes 'with a revolver' as a PP that modifies 'saw', but must halt and reanalyze it as modifying 'spy'. The parser initially attaches the PP to 'saw' because this analysis better satisfies the *minimal attachment* principle, which, very roughly, tells the parser to build phrase structures with as few nonterminal nodes as possible. The heuristics followed by the parser are widely thought to respect minimal attachment in constructing phrase structures for linguistic inputs (Frazier, 1979). The verb-modifying analysis is more likely in the case of (41), where it is given a default reading and subjects do not pause for reanalysis. But, in the case of (40), once the phrase structure is sent upstairs to the central system, the agent realizes that it doesn't make much sense to 'see…with a revolver', and the parser is forced to take a second pass.

The division of labor here is clear: the parser is in charge of building phrase structures for linguistic inputs—indeed, it is the only part of the mind that knows how to do so—but it doesn't "understand" those inputs, since it lacks the information that would be required to check if they make sense. I hypothesize a similar division of labor between semantic composition, on one hand, and pragmatic reasoning, including reference-resolution, on the other. The semantic module knows how to compose semantic values—indeed, it is the only part of the mind that does—but it doesn't know how to identify the contents of semantically underspecified expressions, or to disambiguate on the basis of rich world knowledge.

Should we conclude, then, that disambiguating lexical items works via fallible heuristics, like disambiguating (40)? Semantic garden-path effects, such as the difficulty we encounter in comprehending (42) and (43), bear on this issue (Hirst, 1987, 88).

(42) The astronomer married the star.

(43) The catcher filled the pitcher.

Whereas the syntactic garden-paths discussed above speak in favor of heuristic-based syntactic disambiguation, however, the upshot of these cases is less clear. This

is because the source of processing trouble in these cases appears to be a semantic-priming effect exerted by each sentence's first noun on our interpretation of its second noun. In (42), for example, the presence of 'Astronomer' early in the sentence biases us toward the incorrect, heavenly-body interpretation of 'star' further on. But it looks like this sort of semantic priming reveals the influence of central systems on the mind's first pass at disambiguation, since, for example, knowing what sort of pitchers typically go with catchers is a relatively rich sort of world knowledge, ignorance of which would not impugn one's linguistic competence. This line of thought seemingly threatens the idea that disambiguation is modular. But since nearly everyone assumes that at least first-pass disambiguation must happen prior to semantic composition—indeed, it seemingly must happen as part of syntactic parsing—this phenomenon threatens the modularity of language processing.

Fodor (1983) discusses this problem in connection to (44).

(44) Because he was afraid of electronic surveillance, the spy carefully searched the room for bugs.

Here, the reading of 'bugs' as 'microphones' is facilitated as compared to neutral contexts, in which the 'insects' reading is more prominent. But it is controversial whether this sort of priming can be accomplished within a module. Fodor draws on work by Swinney (1979) to argue that it might be possible to explain this phenomenon by appealing to relatively unsophisticated associative relations between lexical items to which the parser has access because they reside within the lexicon itself, rather than the agent's world knowledge (1983, 79–81). However, it is unclear whether all priming effects on lexical disambiguation could be explained in terms of module-internal heuristics in this way. The debate is, at present, unresolved.

Within the context of the semantic framework I've been developing, however, there is another theoretical option, and this brings me to my framework's second way of representing lexical ambiguity. On this view, ambiguity is treated as a kind of semantic underspecification. I could, for example, think of 'bank' and 'bug' as type-$\langle s, et\rangle$ variables with the following constraint properties:

(45) $\mu(\text{bank}) = \lambda\varphi_{\langle s,et\rangle}\ .\ (\varphi = \lambda w\ .\ \lambda x_e\ .\ x \text{ is a financial institution in } w) \vee (\varphi = \lambda w.\lambda x_e\ .\ x \text{ is a riverside in } w)$

(46) $\mu(\text{bug}) = \lambda\varphi_{\langle s,et\rangle}\ .\ (\varphi = \lambda w\ .\ \lambda x_e\ .\ x \text{ is an insect in } w) \vee (\varphi = \lambda w.\lambda x_e\ .\ x \text{ is a hidden microphone in } w)$

Here, /bank/ and /bug/ each correspond to a single lexical item with a single constraint property, but this constraint property specifies a disjunctive list of properties. Disambiguation, on this view, would not be a task for the parser, which could simply

install the single lexical item, 'bank', where it belongs in a phrase structure. Instead, disambiguation would be a kind of reference resolution—an isotropic, central process of choosing from among the disjuncts in a lexical item's constraint property. Semantic priming could affect this process by centrally activating one of the properties but not the other, or by giving the agent a reason to infer that one content rather than another is intended, all without constituting a case of cognitive penetration.

My point in raising this possibility is not to defend it, but merely to point out that it is a theoretical possibility that is opened up by my semantic framework. Content semanticists, by contrast, are forced to think of lexical disambiguation as something that happens, at least tentatively, before semantic composition. This is because semantic composition must be given one or another of an ambiguous expression's contents as an input. But the semantic framework I've developed offers the possibility (if not necessarily the actuality) that lexical disambiguation can be understood as a wholly post-semantic endeavor.

Although it is unclear whether this theoretical option should be taken up in the case of homonymous expressions such as those I've discussed so far, I think it is much more likely that polysemy should be understood as a form of form of semantic underspecification. The word 'paper' can be used with a variety of closely related senses, for example, and it likely makes sense to capture these meanings by building them all into a single lexical entry. One reason to think so comes from studies conducted by Pylkkänen et al. (2006, 2002), who found that whereas processing of the second of a pair of homonyms in close succession is delayed due to phonological inhibition induced by the first, processing of one sense of a polysemous expression is facilitated by having previously processed the same expression with a different sense. This suggests that the different meanings of a polysemous expression are stored "closer together", in some sense, within the lexicon.

One way to model the meanings of polysemous expressions would be to mimic the strategy applied to 'bank' and 'bug' above. For example, the constraint property of 'paper' can be represented as follows.

(47) $\mu(\text{paper}) = \lambda\varphi_{\langle s,et\rangle}\,.\,(\varphi = \lambda w.\lambda x_e\,.\,x \text{ is made of flattened wood-pulp in } w) \vee (\varphi = \lambda w.\lambda x_e\,.\,x \text{ is a newspaper in } w) \vee (\varphi = \lambda w.\lambda x_e\,.\,x \text{ is an essay in } w) \ldots$

However, some theorists have argued that the meanings of polysemous expressions cannot be captured by any finite list of meanings, but must instead be thought of as a single, thin constraint on the sorts of contents that can be expressed through a literal use of the expression.[31] On this view, 'paper' might be given a constraint property with the following form.

[31] See, for example, Carston (2002, 2012); Travis (2008). For a summary of the debate between

(48) $\mu(\text{paper}) = \lambda\varphi_{\langle s,et\rangle}\,.\,\Sigma(\varphi)$

Here, Σ stands in for a type-$\langle\langle s, et\rangle, t\rangle$ function that maps a property to 1 just in case it can be expressed by someone who uses 'paper' literally. One job of an interpreter—clearly an isotropic, richly inferential, and central job—is to work out which property satisfying Σ is being expressed on a particular occasion when a speaker utters 'paper'.

Just how Σ is represented in the lexicon is a pressing and difficult question for this approach—a question that I won't try to answer. Another interesting question is the proper extent of Σ. When I say that Σ constrains the properties that can be literally expressed using 'paper', it is initially tempting to think that 'literally' is being used in something like its colloquial sense. However, some relevance theorists have argued that contents that would normally be considered non-literal, including metaphorical, loose, and hyperbolic contents, are among the possible contents that an expression may be used to directly and explicitly express with a word on a given occasion, thus, in effect, downplaying or denying the theoretical significance of the literal/non-literal distinction (Carston, 2002, 2010, 2012; Sperber and Wilson, 2008; Wilson and Carston, 2007a,b). This view could be captured within the present framework by making Σ a very loose constraint.

Since polysemy is ubiquitous in natural language, these proposals, if accepted, would likely apply to all, or nearly all, non-functional lexical items, including nouns, verbs, adjectives, proper names, and so on.[32] Although I began my treatment of semantic underspecification by thinking of it as a theory of variables, then, it can also be adapted into a theory of non-functional vocabulary items, which are distinguished by the fact that their meanings are encoded in the lexicon as values of $\mu(\cdot)$ rather than $[\![\cdot]\!]$.[33] This allows me to implement the idea that whereas functional vocabulary items have fixed and precise meanings, the meanings of non-functional vocabulary allow for greater flexibility in what can be said with them.

I emphasize, once again, that my goal here is not to defend these views about polysemy and lexical semantics, but merely to show how they may be formulated within a semantic theory that differs minimally from the standard framework, and that loses little or none of that framework's explanatory power. I think that this constitutes an advance, since, although the view that some or all ambiguity is a kind

"sense-enumeration" and "thin semantics" theories, among other issues arising from polysemy, see Falkum and Vicente (2015).

[32] Applied to proper names, the result is a view according to which names are variables, along the lines of the proposal of Schoubye (2016). The thesis that names are variables and the thesis that names are polysemous converge on the present view.

[33] Cf. Glanzberg (2014, 287), who reaches a similar conception of the breakdown between "structural-functional" and "core conceptual" meaning, albeit with a different motivation and implementation.

of semantic underspecification is coherent and viable, it seemingly can't even be formulated within the standard framework.

Conclusion

The meaning of an expression provides hearers with perceptual evidence of the contents that speakers express in using it, but is not itself a kind of content. This is a good slogan, but although the foundational view that it conveys is not new, its consequences for semantics have remained unclear. My main goal here has been to remedy this situation by showing how to nondestructively modify the standard compositional-semantic framework in order to model semantic values as constraints rather than as contents. This revised framework makes better sense of the place of semantic composition within a well-motivated picture of the human mind. Moreover, the theory—especially when understood in terms of its psychological foundations—can help to make sense of perennially tricky semantic phenomena like indexicality and polysemy—all without giving up much or any of the explanatory power of the standard framework.

References

Abney, S. P. (1987). *The English Noun Phrase in its Sentential Aspect*. PhD thesis, MIT, Cambridge, MA.

Bach, K. (1987). *Thought and Reference*. Oxford University Press.

Bach, K. (1992). Intentions and demonstrations. *Analysis*, 52:140–146.

Bach, K. (2006). The excluded middle: Semantic minimalism without minimal propositions. *Philosophy and Phenomenological Research*, 73(2):435–442.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. MIT Press.

Borg, E. (2004). *Minimal Semantics*. Oxford University Press.

Borg, E. (2012). *Pursuing Meaning*. Oxford University Press.

Bratman, M. (1987). *Intention, plans, and practical reason*. Harvard University Press, Cambridge, Mass.

Braun, D. (1996). Demonstratives and their linguistic meanings. *Noûs*, 30(2):145–173.

Broome, J. (2013). *Rationality Through Reasoning*. Wiley-Blackwell.

Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4(1):73–121.

Burge, T. (1986). Individualism and psychology. *Philosophical Review*, 95(1):3–45.

Carruthers, P. (2006). *The Architecture of Mind: Massive Modularity and the Flexibility of Thought*. Oxford University Press.

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

Carston, R. (2006). Relevance theory and the saying/implying distinction. In Horn, L. R. and Ward, G., editors, *The Handbook of Pragmatics*, chapter 28, pages 633–656. Blackwell, Oxford.

Carston, R. (2010). Lexical pragmatics, ad hoc concepts and metaphor: from a relevance theory perspective. *Italian Journal of Linguistics*, 22(1):153–180.

Carston, R. (2012). Word meaning and concept expressed. *Linguistic Review*, 29(4):607–623.

Chierchia, G. and McConnell-Ginet, S. (1990). *Meaning and Grammar: An Introduction to Semantics*. MIT Press, Cambridge, MA.

Devitt, M. (2006). *Ignorance of Language*. Oxford University Press.

Devitt, M. (2013). What makes a property "semantic"? In Capone, A., Piparo, F. L., and Carapezza, M., editors, *Perspectives on Pragmatics and Philosophy*, pages 87–112. Springer.

Devitt, M. and Sterelny, K. (1999). *Language and Reality: An Introduction to the Philosophy of Language*. Blackwell, Oxford, UK, second edition.

Donnellan, K. S. (1968). Putting humpty dumpty together again. *Philosophical Review*, 77(2):203–215.

Elbourne, P. (2005). *Situations and Individuals*. MIT Press, Cambridge, MA.

Falkum, I. L. and Vicente, A. (2015). Polysemy: Current perspectives and approaches. *Lingua*, 157:1–15.

Firestone, C. and Scholl, B. J. (2015). Cognition does not affect perception: Evaluating the evidence for 'top-down' effects. *Behavioral and Brain Sciences*, Target Article Under Comment as of 2015.

Fodor, J. (1983). *The Modularity of Mind*. MIT Press, Cambridge, Mass.

Fodor, J. A. (1989). review of 'remnants of meaning' by stephen schiffer. *Philosophy and Phenomenological Research*, 50:409–423.

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

Fodor, J. A. (2002). Language, thought, and compositionality. *Mind and Language*, 16(1):1–15.

Frazier, L. (1979). *On Comprehending Sentences: Syntactic Parsing Strategies*. PhD thesis, University of Connecticut, Storrs, CT.

Garía-Carpintero, M. (2006). Recanati on the semantics/pragmatics distinction. *CRÍTICA, Revisita Hispanoamericana de Filosofia*, 38(112):35–68.

Glanzberg, M. (2014). Explanation and partiality in semantic theory. In Burgess, A. and Sherman, B., editors, *Metasemantics*, chapter 10, pages 259—292. Oxford University Press, Oxford, UK.

Grice, H. P. (1971). Intention and uncertainty. *Proceedings of the British Academy*, 57:263–279.

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

Hackl, M., Koster-Hale, J., and Varvoutis, J. (2012). Quantification and acd: Evidence from real-time sentence processing. *Journal of Semantics*, 29(2):145–206.

Hacquard, V. (2014). Bootstrapping attitudes. In Snider, T., editor, *Proceedings of SALT 24*, pages 330–352.

Hawthorne, J. and Manley, D. (2012). *The Reference Book*. Oxford University Press.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD Dissertation, University of Massachusetts at Amherst.

Heim, I. (1983). File change semantics and the familiarity theory of definiteness. In Bäuerle, R., Schwarze, C., and von Stechow, A., editors, *Meaning, Use and Interpretation of Language*, pages 164–189. de Gruyter, Berlin.

Heim, I. (2008). Features on bound pronouns. In Harbour, D., Adger, D., and Bejar, S., editors, *Phi-Theory: Phi-Features across Modules and Interfaces*. Oxford University Press, Oxford, UK.

Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*. Blackwell.

Hirst, G. (1987). *Semantic interpretation and the resolution of ambiguity*. Cambridge University Press.

Holton, R. (2011). *Willing, Wanting, Waiting*. Oxford University Press.

Jacobson, P. (2014). *Compositional Semantics: An Introduction to the Syntax/Semantics Interface*. Oxford University Press.

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

Kaplan, D. (1989b). Demonstratives. In Joseph Almog, J. P. and Wettstein, H., editors, *Themes from Kaplan*, pages 481–563. Oxford University Press.

Katz, J. J. (1981). *Language and Other Abstract Objects*. Rowman & Littlefield, Totowa, NJ.

Katz, J. J. (1984). An outline of a platonist grammar. In *The Philosophy of Linguistics*. Oxford University Press, Oxford, UK.

Kennedy, C. (1999). *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. Garland.

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

King, J. C. (2013). Supplementives, the coordination account, and conflicting intentions. *Philosophical Perspectives*, 27.

King, J. C. (2014). Speaker intentions in context. *Noûs*, 48(2):219–237.

King, J. C. and Stanley, J. (2005). Semantics, pragmatics, and the role of semantic content. In Szabó, Z. G., editor, *Semantics versus Pragmatics*, pages 111–164. Oxford University Press, Oxford.

Kratzer, A. (1998). More structural analogies between pronouns and tenses. In *SALT 8, MIT*.

Larson, R. and Segal, G. (1995). *Knowledge of Meaning*. MIT Press, Cambridge, MA.

Lepore, E. and Ludwig, K. (2007). *Donald Davidson's Truth-Theoretic Semantics*. Oxford University Press, Oxford.

Lewis, D. K. (1975). Languages and language. *Minnesota Studies in the Philosophy of Science*, VII:3–35.

Lewis, D. K. (1992). Meaning without use: Reply to hawthorne. *Australasian Journal of Philosophy*, 70(1):106–110.

Lewis, S. (2013). *Pragmatic Enrichment in Language Processing and Development*. PhD thesis, University of Maryland, College Park, Maryland.

Lewis, S., Hacquard, V., and Lidz, J. (MS). "think" pragmatically: Children's interpretation of belief reports. *Unpublished Manuscript*.

Ludlow, P. (2011). *The Philosophy of Generative Linguistics*. Oxford University Press, Oxford, UK.

Mandelbaum, E. (2017). Seeing and conceptualizing: Modularity and the shallow contents of perception.

Philosophy and Phenomenological Research, forthcoming.

Montague, R. (1970a). English as a formal language. In Visentini, B., editor, *Linguaggi nella società e nella tecnica*, pages 189–223. Mailand.

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

Neale, S. (2004). This, that, and the other. In Bezuidenhout, A. and Reimer, M., editors, *Descriptions and Beyond*, pages 68–182. Oxford University Press.

Neale, S. (2005). Pragmatism and binding. In Szabó, Z. G., editor, *Semantics versus Pragmatics*, pages 165–285. Oxford University Press.

Neale, S. (2016). Implicit and aphonic reference. In Ostertag, G., editor, *Meanings and Other Things: Essays in Honor of Stephen Schiffer*. Oxford University Press.

Ninan, D. (2010). Semantics and the objects of assertion. *Linguistics and Philosophy*, 33(5):355–380.

Nowak, E. (2016). *Two Dogmas about Demonstratives*. PhD thesis, University of California, Berkeley, Berkeley.

Nowak, E. (MS). No context, no content, no problem. *Unpublished Manuscript*.

Nunberg, G. (1993). Indexicality and deixis. *Linguistics and Philosophy*, 16(1):1–43.

Partee, B. (1981). Montague grammar, mental representations, and reality. In Kanger, S. and Öhman, S., editors, *Philosophy and Grammar*, volume 143 of *Synthese Library*, chapter 5, pages 59–78. Synthese.

Partee, B. (1989). Binding implicit variables in quantified contexts: Parasession on language in context. In Wiltshire, C., Music, B., and Graczyk., R., editors, *Papers from the 25th Regional Meeting, Chicago Linguistics Society*. Chicago Linguistics Society, Chicago.

Pereplyotchik, D. (2017). *The Psychological Import of Syntactic Theory*. Springer.

Perry, J. (2001). *Reference and Reflexivity*. CSLI.

Perry, J. (2009). Directing intentions. In Almog, J. and Leonardi, P., editors, *The Philosophy of David Kaplan*. Oxford University Press, Oxford.

Pietroski, P., Lidz, J., Hunter, T., and Halberda, J. (2009). The meaning of 'most': Semantics, numerosity and psychology. *Mind and Language*, 24(5):554–585.

Pietroski, P. M. (2008). Minimalist meaning, internalist interpretation. *Biolinguistics*, 2(4):317–341.

Postal, P. (1969). On so-called 'pronouns' in english. In Reibel, D. A. and Schane, S. A., editors, *Modern Studies in English: Readings in Transformational Grammar*. Prentice Hall, New Jersey.

Pylkkänen, L., Llinás, R., and Murphy, G. L. (2006). The representation of polysemy: Meg evidence. *Journal of Cognitive Neuroscience*', 18(1):97–109.

Pylkkänen, L., Strangefellow, A., and Marantz, A. (2002). Neuromagnetic evidence for the timing of lexical activation: An meg component sensitive to phonotactic probability but not to neighborhood density. *Brain and Language*, 81:666–678.

Rabern, B. (2012). Against the identification of assertoric content with compositional value. *Synthese*, 189(1):75–96.

Rayner, K., Carlson, M., and Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. *Journal of Verbal Learning and Verbal Behavior*, 22(3):358–374.

Recanati, F. (2010). *Truth-Conditional Pragmatics*. Oxford University Press.

Roberts, C. (2002). Demonstratives as definites. In van Deemyer, K. and Kibble, R., editors, *Information Sharing*, pages 1–48. CSLI.

Roberts, C. (2003). Uniqueness in definite noun phrases. *Linguistics and Philosophy*, 26:287–350.

Rullmann, H. (2004). First and second person pronouns as bound variables. *Linguistic Inquiry*, 35(1):159–168.

Schiffer, S. (1981). Indexicals and the theory of reference. *Synthese*, 49(1):43–100.

Schiffer, S. (1994). A paradox of meaning. *Noûs*, 28:279–324.

Schiffer, S. (2003). *The Things We Mean*. Oxford University Press.

Schoubye, A. (2016). Type-ambiguous names. *Mind*, forthcoming.

Soames, S. (1984). Linguistics and psychology. *Linguistics and Philosophy*, 7:155–79.

Soames, S. (1985). Semantics and psychology. In Katz, J. J., editor, *The Philosophy of Linguistics*, pages 204–226. Oxford University Press, Oxford, UK.

Soames, S. (1989). Semantics and semantic competence. In Tomberlin, J. E., editor, *Philosophocal Perspectives 3: Philosophy of Mind and Action Theory*, pages 575–596. Ridgeview, Atascadero, CA.

Soames, S. (2002). *Beyond Rigidity: The Unfinished Semantic Agenda of* Naming and Necessity. Oxford University Press, Oxford.

Soames, S. (2009). *Philosophical Essays, Volume 1: Natural Language: What it Means and How We Use It*. Princeton University Press.

Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*. Blackwell, Oxford, 2 edition.

Sperber, D. and Wilson, D. (2002). Pragmatics, modularity, and mindreading. *Mind and Language*, 17:3–23.

Sperber, D. and Wilson, D. (2008). A deflationary account of metaphors. In Gibbs, R., editor, *The Cambridge Handbook of Metaphor and Thought*, pages 84–105. Cambridge University Press.

Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Syntax and Semantics 9*, pages 315–332. Academic Press, New York.

Stalnaker, R. (2014). *Context*. Context and Content. Oxford University Press, Oxford.

Stanley, J. (1997). Rigidity and content. In Heck, R., editor, *Language, Thought, and Logic: Essays in honour of Michael Dummett*, pages 131–156. Oxford University Press, Oxford, UK.

Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23(4):391–434.

Stanley, J. and Szabó, Z. G. (2000). On quantifier domain restriction. *Mind and Language*, 15(2–3):219–261.

Swanson, E. (2016). The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*, 45(2):121–146.

Swinney, D. (1979). Lexical access during sentence comprehension: (re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18:645–660.

Travis, C. (2008). *Occasion-Sensitivity: Selected Essays*. Oxford University Press.

von Fintel, K. and Gillies, A. S. (2008). Cia leaks. *Philosophical Review*, 117(1):77–98.

von Fintel, K. and Heim, I. (2011). *Intensional Semantics*. Ubpublished Lecture Notes, online at http://web.mit.edu/fintel/fintel-heim-intensional.pdf, spring 2011 edition.

Westerståhl, D. (1984). Determiners and context sets. In van Bentham, J. and ter Meulen, A., editors, *Generalized Quantifiers in Natural Language*, pages 45–71. Foris, Dordrecht.

Wilson, D. and Carston, R. (2007a). Metaphor and the 'emergent property' problem: A relevance-theoretic approach. *The Baltic International Yearbook of Cognition, Logic, and Communication*, 3.

Wilson, D. and Carston, R. (2007b). A unitary approach to lexical pragmatics: relevance, inference and ad hoc concepts. In Burton-Roberts, N., editor, *Pragmatics*, pages 230–259. Palgrave Macmillan, Basingstoke and New York.

Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464):983–1026.

Yalcin, S. (2014). Semantics and metasemantics in the context of generative grammar. In Burgess, A. and Sherman, B., editors, *Metasemantics*, pages 17–54. Oxford University Press, Oxford, UK.