

Semantics without Semantic Content

Daniel W. Harris*

October 1, 2018

Abstract

I argue that semantics is the study of the proprietary database of a modular input–output system. The role of this semantic module is to encode and decode partial and defeasible evidence of what speakers are saying. In particular, I argue that the semantic module is both centrally inaccessible and informationally encapsulated. This has consequences for how we do semantics, since it places information about nonlinguistic context outside the purview of semantic processing. We must therefore adopt the nonstandard view that a sentence’s semantic value is not its content but is rather a partial and defeasible constraint on what it can be used to say. Unlike others who have defended similar views, I show how to translate this thesis into a detailed compositional-semantic theory. I do so by nondestructively modifying the influential framework of Heim & Kratzer (1998) and von Stechow & Heim (2011). The resulting approach situates semantics within an independently motivated account of human cognitive architecture and reveals the semantics–pragmatics interface to be grounded in the underlying interface between modular and central systems.

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

*Forthcoming in *Mind & Language*. This is a penultimate draft. For helpful conversation and feedback on earlier versions, thanks to David Braun, Jonathan Cohen, Kai von Stechow, Manuel García-Carpentiero, Justin Khoo, Eric Mandelbaum, Eliot Michaelson, Stephen Neale, Gary Ostertag, David Pereplyotchik, Jake Quilty-Dunn, Brian Rabern, François Recanati, Kate Ritchie, Craige Roberts, Yael Sharvit, Richard Stillman, Elmar Unnsteinsson, and two anonymous referees for *Mind & Language*.

part by defending the latter view and showing that it entails the former (§2), and then by showing how to modify a standard compositional-semantic framework to make it fit with both ideas (§3).

1 Two Conceptions of Semantic Values

My first order of business is to distinguish two theories of the nature of semantic values, which I will 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 is its content, and this content may vary in a way that depends on extralinguistic context. 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 , indexed to a world w , assignment g , and context c , with an extension $\llbracket \varphi \rrbracket^{w,g,c}$ (or, by λ -abstracting over the world parameter, with an intension $\llbracket \varphi \rrbracket^{g,c}$).¹ Idealizing away from tense, aspect, clause-type, and some other complications, the semantic value of ‘Ann smokes’ can be specified in either of the following ways, for example.

- (1) $\llbracket \text{Ann smokes} \rrbracket^{w,g,c} = 1$ iff Ann smokes at w
- (2) $\llbracket \text{Ann smokes} \rrbracket^{g,c} = \lambda w . \text{Ann smokes at } w$

The sentence’s semantic content—the proposition that is true at all and only worlds where Ann smokes—is given by (2). Sentential semantic values like these can be derived by positing contents for lexical items—e.g., a referent for ‘Ann’ as in (3) and the property of being a smoker to ‘smokes’ as in (4)—and formulating principles of content composition, such as (5).

- (3) $\llbracket \text{Ann} \rrbracket^{w,g,c} = \text{Ann}$
- (4) $\llbracket \text{smokes} \rrbracket^{w,g,c} = \lambda x_e . x \text{ smokes at } w$

¹Throughout this essay, I will treat the framework codified by Heim & Kratzer (1998) and von Stechow & Heim (2011) as my principle foil, sometimes referring to it as ‘the standard framework’, using its claims and notation (or slightly simplified versions) as my starting point. There are, of course, alternative approaches that have matured enough to have their own textbooks—to name just three: Davidsonian truth-theoretic semantics (Larson and Segal, 1995; Lepore and Ludwig, 2007), variable-free semantics (Jacobson, 2014), and discourse representation theory (Kamp and Reyle, 1993). There are interesting questions about how the morals of the present essay would apply to these other frameworks, but I don’t have space to explore these questions here.

(5) FUNCTIONAL APPLICATION

If α is a branching node and $\{\beta, \gamma\}$ the set of its daughters, then, for any assignment g , α is in the domain of $\llbracket \cdot \rrbracket^{w,g,c}$ if both β and γ are, and $\llbracket \beta \rrbracket^{w,g,c}$ is a function whose domain contains $\llbracket \gamma \rrbracket^{w,g,c}$. In this case, $\llbracket \alpha \rrbracket^{w,g,c} = \llbracket \beta \rrbracket^{w,g,c}(\llbracket \gamma \rrbracket^{w,g,c})$.

The overarching idea of clauses like (1)–(5) is that the semantic value of a declarative sentence is a propositional content, here understood as a function from worlds to truth values, and that propositions can be computed by positing contents for sub-sentential expressions along with principles of content composition.

The assignment parameter g and context parameter 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.

(6) I smoke.

(7) The man smokes.

(8) 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_c who is the semantic value of 'I' relative to c , for example (Kaplan, 1989b). This yields the following value for (6).

(9) $\llbracket \text{I smoke} \rrbracket^{g,c} = \lambda w . \text{SPEAKER}_c \text{ smokes at } w$

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 LF of every DP, so that (10) is understood to have an LF of the following form, for example (Stanley, 2000; Stanley and Szabó, 2000; Westerstahl, 1984).

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

On this simplified version of the proposal, *dom* functions as an unpronounced restrictive 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 (8), 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 presuppositional constraints imposed by x 's φ -features. This gives rise to assignment-dependent sentential semantic values, such as (12).

$$(11) \llbracket \text{He}_1 \text{ smokes} \rrbracket^{g,c} = \lambda w . g(1) \text{ smokes at } w \quad \text{if } g(1) \text{ is male; else undefined.}$$

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 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.

$$(12) \llbracket \text{He}_1 \text{ smokes} \rrbracket^c = \lambda w . g_c(1) \text{ smokes at } w \quad \text{if } g_c(1) \text{ is male; else undefined.}$$

What facts about a context determine which assignment is operative? (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 it is uttered is an intuitive and popular 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,

²E.g. Heim (2008); Kaplan (1989a); King (2013, 2014); Perry (2009). 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 of what speakers say by 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).

³See, e.g., King (2013, 2014).

for example, and so have held that it is these states that ultimately determine a pronoun's semantic value.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, (8) can occur embedded in (13).

(13) [Every doctor]₁ denies that he₁ 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 expression. 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 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 more expressions as variables than is sometimes assumed.

1.2 Semantic Values as Constraints

Constraint semantics, as I understand it, is the view that an expression's semantic value is not its context-relative content, but rather something context-neutral and therefore less informative. 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. For an expression to be semantically underspecified is for there to be slack in the constraint that its semantic value places on what it can be used to say. The semantic value of ‘he’ tells a hearer that the speaker is referring to a male (if they are speaking literally), but not *which* male, for example. If a view like this is correct, there may no need for compositional semantics to traffic in a theoretical notion of context at all.⁵

⁴E.g. Clark and Marshall (1992); Heim (1982, 1983); Roberts (2002, 2003); Stalnaker (1978, 2014).

⁵Note that constraint semantics, as I conceive it, does *not* give up the idea that propositions are the things we assert, say, or mean. For some recent defenses of the distinction between semantic values

I am not the first to advocate a version of constraint semantics. For example, Sperber & Wilson (1995, 175) 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”. Carston (2006, 633) 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”. According to Pietroski (Pietroski, 2006, 34), the meanings of declarative sentences “constrain without determining truth/reference/satisfaction conditions”. Bach (1987, 5) argues that “...the semantics of an expression gives the information that a competent speaker can glean from it independently of any context of utterance”. 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” (Bach, 2006, 437). Neale (2005, 189) 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”. Schiffer (2003, 112) argues that the semantic value of a sentence (which he calls a ‘character*’) is “an ordered pair $\langle A, P \rangle$, 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”. As Schiffer makes clear, both the force and content of a speech act will normally be underspecified by the character* of the sentence with 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).^{6,7}

and the contents of speech acts—albeit, for reasons rather different than my own—see Ninan (2010); Rabern (2012); Stanley (1997); Yalcin (2007).

⁶See also Garía-Carpintero (2006) for a further defense of a Schiffer-influenced view against objections by King and Stanley (2005) and Recanati (2004).

⁷Some have defended views that are similar to constraint semantics, but which include commitments that I do not wish to incur.

One is that semantics should compute *both* context-relativized contents and context-neutral, constraint-like semantic values. For example, Kaplan (1989b, 507) suggests that semantics should include a compositional theory of characters as well as a compositional theory of contents, though he only spells out the latter. Similarly, Perry (2001) posits context-neutral “reflexive contents” in addition to context-specific “referential contents”, and Stalnaker (1978) posits context-neutral “propositional concepts” in addition to semantic contents. In §2, I will argue that semantics should deal *only* in context-neutral semantic values.

Gross (2005) argues that at least some context-sensitive expressions should be handled by translating them directly into context-sensitive expressions in the metalanguage. However, a problem with this view that Gross acknowledges is that some expressions will have to be translated in a way that accounts for the language user’s perspective. Whether ‘you’ is translated as ‘I’, ‘you’, or ‘they’, for ex-

Many of the proposals I've just mentioned have been motivated by foundational considerations that go roughly as follows. There is a deep distinction between the kinds of psychological faculties underlying semantics and pragmatics. Pragmatics is the study of an application of the human capacity for mindreading: to communicate is to have one's communicative intention recognized by the intended addressee. Semantics, by contrast, is the study of a psychological mechanism that encodes and decodes partial and defeasible evidence of speakers' communicative intentions. Since working out what a speaker says with a semantically underspecified expression typically involves mindreading—and therefore pragmatic reasoning—semantics can be in the business of specifying only constraints, rather than contents. My aim in §2 will be to argue that this line of thought identifies a real joint in nature that is robust enough to ground the semantics–pragmatics distinction.

A significant weakness of nearly all prior defenses of constraint semantics is that their foundational claims have not been accompanied by compositional-semantic implementation details.⁸ Nearly all of the authors mentioned above have eschewed the task of constructing a detailed compositional-semantic theory that outputs constraints rather than context-relativized contents. I will therefore demonstrate one way of tackling this problem in §3.

ample, will have to depend on whether the translator resides in the mind of the addressee, the speaker, or a third party. In §2, I argue that semantics lacks access to this kind of information.

Some other authors who have defended the view that semantics (sometimes) delivers something less specific than propositional contents include Alonso-Ovalle (2006); Barwise and Perry (1983); Ciardelli et al. (2013, 2017); von Fintel and Gillies (2008); Kratzer and Shimoyama (2002); Swanson (2016), though each of these views is quite different in its motivations and implementation from the theory I defend here.

⁸An important exception is Pietroski, who has fleshed out the formal-semantic details of his foundational vision—most fully in a recent monograph (Pietroski, 2018). Some of Pietroski's foundational motivations parallel those that I will raise in §2, and I am sympathetic to many aspects of his view. However, Pietroski departs from orthodoxy much more than I will here. For example, Pietroski commits himself to internalism about semantic representations, and uses that position as a premise in arguments for his conception of semantics (e.g., 2005; 2006; 2018; 2008), whereas none of my arguments for constraint semantics require me to take a stand on that issue. Pietroski also argues that we must abandon many of the compositional details of the standard framework (2018), but I will argue that constraint semantics needn't force us to abandon the usual compositional principles, only augment them. Ultimately, Pietroski's more radical position may be justified, but my defense and implementation of constraint semantics needn't push us that far.

2 Semantics and Cognitive Architecture

2.1 The Modularity of Semantics

One could take the content–constraint issue to be an idle one. This is the position of David Lewis (1980), who shows how to distinguish a kind of thin, context-neutral meanings from the usual context-dependent intensions, but argues that the question of which should be counted as semantic values is uninteresting: both “equally well deserve the name of semantic values because they equally well do the jobs of semantic values” (1980, 92). What, then, is the job of semantic values? Lewis assumes that semantic values are whatever it is that a good grammar assigns to expressions, and “a good grammar is one suited to play a certain role in a systematic restatement of our common knowledge about language” (Lewis, 1980, 81).

My aim in this section is to argue that Lewis was wrong to think that the content–constraint debate is idle, and that we have good reason to prefer constraint semantics. My argument rests on the claim that whereas pragmatics is (or at least includes) the study of how language use depends on central-cognitive inferential capacities, including mindreading, semantics is by contrast the study of a modular input-output system in Fodor’s (1983) sense.⁹ Although this system outputs representations to central cognition in language perception and takes inputs from central cognition in language production, it draws only on a proprietary database of information in carrying out its computations. I will refer to the system in question as the “semantic module”, to the task it performs in speech perception—computing sentential semantic values from LFs—as “semantic composition”, and to its proprietary database as “semantic competence”. On the view I wish to defend, the aim of compositional semantics is to reverse-engineer language-users’ semantic competence, in part by isolating it from the information available to their central-cognitive system(s), such as their beliefs and memories.¹⁰

Such a picture gives us a good reason to reject Lewis’s assumption that a good grammar models speakers’ “common knowledge about language”, since this assumption wrongly lumps together at least two functionally distinct bodies of information that play different roles in language use. On one hand, there is semantic competence, which informs the operations of the semantic module. On the other

⁹I am *not* advocating a massively modular approach to cognitive architecture, on which even what Fodor labels ‘central cognition’ is divided into unencapsulated subsystems with distinct evolutionary origins and that are, in some senses, domain specific (Carruthers, 2006; Cosmides and Tooby, 1992; Pinker, 1997; Sperber and Wilson, 2002). My argument is compatible with this possibility, but does not depend on it.

¹⁰Note that understanding the actual process that I call ‘semantic composition’ is a further task for psycholinguistics, and that I am not assuming that semantic composition closely resembles the derivations of sentential semantic values that are given by semanticists.

hand, there are our beliefs and other intentional mental states, which are drawn upon in pragmatic reasoning, but to which the semantic module lacks access. My point in this section is that an empirically adequate account of the psychology underlying language use mustn't conflate these two bodies of information.^{11,12}

If contemporary semantic theories are on the right track, then semantic composition draws on at least two kinds of information—lexical semantic values and composition principles—in order to perform this task. Why should we think that semantic composition is a modular process whose proprietary database includes information of these kinds?¹³

As a preliminary matter, we can note that semantic composition has many of what Fodor (1983) takes to be the stereotypical features of modular processes. It is fast. It is mandatory, in that we cannot voluntarily choose not to perceive linguistic utterances as meaningful. Although semantic composition is not domain-specific in the sense of being driven by the outputs of a proprietary sensory transducer, it is domain-specific in the sense that it can act only on a specific kind of inputs (LFs). As semanticists have shown us, semantic composition resembles typical modular processes in that it is amenable to computational modeling. And by contrast with many central-cognitive processes, semantic composition is not amenable to folk-psychological explanation.

Most importantly, semantic composition is performed by a system that is, in large part, informationally isolated from central cognition. The semantic module sends sentential semantic values as outputs to central cognition as one step in language comprehension, and takes instructions from central cognition as one step in language production, but there are good reasons to think that the semantic module and central cognition lack access to one another's inner workings and proprietary databases.

One half of this informational isolation is that the operations and database of the semantic module are *centrally inaccessible*. Language users are never consciously

¹¹To be fair to Lewis, he anticipated that “a grammar that assigns one sort of semantic value could fit better into future psycholinguistics than one that assigns another sort” (Lewis, 1980, 83), but claimed that it would be speculative to draw any such conclusion at the time. My point in this section is that this is not a speculative project at present: contemporary cognitive science gives us good reasons to prefer constraints over contents.

¹²Some have argued that semantics should not be understood as the study of any aspect of human psychology, but has some other subject matter—e.g. Devitt (2006); Devitt and Sterelny (1999); Katz (1981); Montague (1970); Soames (1984, 1985). I will bracket these views for lack of space, and take it for granted that semantics is the study of a body of mentally represented information that plays a causal role in the perception and production of linguistic utterances.

¹³For some previous defenses of the modularity of semantics, see Borg (2004, ch.3) and Sperber and Wilson (1995, ch.4). I will discuss Borg's view later in §2.2. I agree, in essence, with Sperber and Wilson's view, but my arguments differ from theirs in appealing to specific features of contemporary semantic theory.

aware of any aspect of the process of semantic composition, and are unable to report facts about the intermediate stages of the process or the database on which it draws. Theorists have, at various times, disagreed about nearly every detail of the principles underlying semantic composition, and none of these debates have ever been resolved by direct access to the semantic module. Instead, the inner workings and informational resources of the semantic module must be inferred via a painstaking process of reverse engineering. These points are particularly clear if we focus on the aspects of semantic theories on which research in semantics is focused. For example: Are the semantic-type mismatches created by object-position DPs resolved through quantifier raising (e.g. May 1985; Heim and Kratzer 1998), in-situ type-shifting (e.g. Szabolcsi 1987; Jacobson 2014), or by some other mechanism? Are proper names devices of direct reference (Kaplan, 1989b), variables (Cumming, 2008; Schoubye, 2016), predicates (Fara, 2015), quantifiers (Montague, 1973), or type-flexible (Partee, 1986)? Questions like these are the bread and butter of contemporary semantics, and yet we utterly lack the ability to answer them by introspecting or otherwise directly accessing the semantic principles or processes that allow us to use the expressions concerned.

Indeed, most language users lack the conceptual resources to even consider questions like these. If contemporary semanticists are on the right track, then speakers' semantic modules routinely deal in representations that are framed in terms of concepts that are sequestered from central cognition—concepts such as SEMANTIC TYPE and NUMERICAL INDEX. I will say that these are 'centrally inaccessible concepts'. Although different semantic frameworks posit semantic representations involving different centrally inaccessible concepts, every framework posits some such concepts. For example, variable-free semantics eschews numerical indices but posits copious type-shifting operations that are sensitive to semantic values' semantic types (Jacobson, 2014). But it is deeply implausible that ordinary language-users have the capacity to have beliefs, intentions, or other central-cognitive representations about things like numerical indices or semantic types—at least not without training in semantics.

The second half of the language module's isolation is its *informational encapsulation*. Semantic composition is insensitive to language-users' beliefs, intentions, and other central-cognitive states and information. The main reason for thinking so is that hearers sometimes have beliefs and other states that would be relevant to semantic processing, but that have no effect on it.

Pettit (2002; 2005) has illustrated this point with respect to hearers' beliefs about word meanings. In the case he imagines, a group of devious neuroscientists have convinced a subject to falsely believe that his "brain has been altered so as to produce an aphasia the effect of which is that all of the mass nouns in [his] vocabulary will seem to [him] to mean something they do not in fact mean" (Pettit, 2002, 544). The

subject is then presented with example sentences such as (14).

(14) Let's have mud for lunch.

Although the subject falsely believes, both before and after encountering this sentence, that 'mud' does not mean what it in fact does, Pettit argues that these beliefs do not interfere with the subject's ability to understand utterances of sentences like (14). According to Pettit, this can be seen from the fact that an utterance of (14) will still *seem* to the subject to have the correct meaning, even if the subject treats this impression as unreliable evidence about what the speaker is saying with it on this occasion. This case is analogous to illusions that have driven some of the most influential arguments for the modularity of visual perception. What makes the Müller-Lyer illusion such a striking illustration of informational encapsulation, for example, is that its two central lines continue to *seem* to have different lengths, no matter how strongly one believes that they are the same length.¹⁴

Here is a potential objection to the foregoing argument. If I tell you that I will use 'dog' to mean what people normally mean by 'cat' for the rest of the present conversation, you will be able to understand me on these terms. Doesn't this show that your beliefs about the meaning of 'dog' can play a role in how you will understand utterances of that word? The problem with this objection is that it equivocates two senses in which one can 'understand' an utterance. The kind of understanding that is relevant to my argument in this section is successful semantic composition—i.e., the successful computation of a sentential semantic value. The second kind of understanding is understanding a speaker, which entails forming a veridical representation (perhaps a true belief) about what they said or meant by their utterance. On the view I wish to defend here, what a speaker says or means is never identical to the semantic value of the sentence they utter, and so central-cognitive work is needed to bridge the gap between these two forms of understanding. This is to say that a sentence's semantic value gives a hearer only *partial* evidence of what the speaker means. I also maintain that the output of semantic composition gives a hearer only *defeasible* evidence of what a speaker says or means. Consider a case in which a hearer understands a speaker in spite of the fact that the speaker has uttered a malapropism. Hearing Yogi Berra utter the sentence, 'Texas has a lot of electrical votes,' a hearer may come to recognize that Berra meant that Texas has a lot of electoral votes. On the modularist view I am here defending, this process will

¹⁴The point about the Müller-Lyer illusion is originally due to (Fodor, 1983, 66). For a recent defense of the modularity of visual perception that emphasizes this and similar data for informational encapsulation, see Firestone and Scholl (2015). Although Pettit (2002) does not present this case as an argument for modularity, he briefly makes that implication explicit and mentions the analogy to visual-illusion cases in a later follow-up article (Pettit, 2005, 74).

go roughly as follows: the hearer's semantic module outputs a representation that gives the hearer evidence that Berra said, of something called 'Texas', that it has a lot of electrical votes. Bringing to bear the resources of central cognition, the hearer infers that this evidence mustn't be accurate and, noting the similarity of 'electrical' and 'electoral', concludes that Berra meant (and perhaps *said*¹⁵) that Texas has a lot of electoral votes.¹⁶ Similarly, in my 'dog'/'cat' scenario, the prediction of my view is that when I utter 'dog' it will initially seem to my hearer that I am saying something about dogs, since this is what their semantic module indicates, but that they will be capable of correcting for this misleading evidence with a little extra cognitive work. We might expect this extra work to manifest itself as extra time or cognitive load spent on understanding utterances in which words are being used with known but nonstandard meanings.

My Pettit-inspired argument focuses on the impenetrability of semantic composition by our beliefs about the most vague and general semantic features of words, such as the belief that the word 'mud' refers mud rather than something else. These are semantic features that are rough and obvious enough to be tracked into ordinary speakers' folk theories. Matters are much more cut and dried when it comes to the semantic properties of expressions that determine how they compose—i.e., the properties about which theorists tend to debate.

Suppose, for example, that two semanticists—call them 'Robert' and 'Pauline'—are having a conversation. Robert obstinately believes that object-position DPs are raised at LF and interpreted via a predicate-abstraction principle. On the other hand, Pauline vehemently denies the existence of raising, and believes that object-position DPs are interpreted in situ with the help of type-shifting operations. At least one of our two semanticists must be wrong, which is to say that their belief is inconsistent with the actual compositional principles that guide the processes that allows them to produce and perceive language. If semantic composition were not an informationally encapsulated process—that is, if it were sensitive to agents' beliefs—then we should expect strongly avowed false beliefs about semantic matters to have a deleterious effect on an agent's ability to create and understand meaningful utterances. But there is no evidence for top-down effects of this kind. Indeed, it would be

¹⁵See Unnsteinsson (2016) for a defense of the claim that what a speaker says with a malapropism needn't match the usual meaning of the word they utter, and may instead be their intended meaning.

¹⁶Now would be a good time to emphasize that central-cognitive processes needn't always be either slow, conscious, or reportable. We have plenty of unconscious beliefs and we draw plenty of unconscious inferences, for example, and these are among the denizens of central cognition. The fact that hearers wouldn't necessarily report themselves as going through inferential steps like those described here is therefore not on its own a good reason to think that they don't do so. Nonetheless, as I will emphasize below, the workings of central cognition are somewhat mysterious, and so I don't wish to present my rough story as anything more than a loose conjecture about how the inferences in question transpire.

truly shocking if we were to find evidence that believing certain semantic theories either enhanced or detracted from one's linguistic capacities, even a little bit. The explanation for this fact is that our beliefs and other central-cognitive representations about compositional-semantic matters have no impact on the activities of the semantic module. This is a powerful reason to think that the grammatical principles that guide semantic composition inhabit the proprietary database of a cognitively impenetrable system, and aren't central-cognitive representations.

One complication for the idea that semantics is the study of a modular system arises from the fact that this system must presumably play a role in language production as well as comprehension. This means that the system must be capable of taking input from central cognition in the course of speech production. Nonetheless, I see no reason to think that this fact undermines the informational encapsulation of the language module in its role as an input system. Moreover, there is a clear sense in which the semantic module *qua* output system is informationally encapsulated as well. The representations at intermediate stages of semantic processing and many of the concepts in terms of which these representations are framed are just as centrally inaccessible on the way out as they are on the way in. Speakers don't recognize themselves as designing particular semantic representations in the course of producing an utterance, and they can't, in part because their central system(s) lack some of the concepts that would be needed in order to do so. Speakers thus aren't in a position to interfere with semantic processing at these stages, or to otherwise integrate central information with intermediate semantic representations in language production. A speaker simply can't choose whether to represent proper names as type-*e* terms or as type-*et* predicates, for example. Likewise, our disagreeing semanticists, Robert and Pauline, can't bring their theoretical beliefs to bear on whether to represent the meanings of object-position DPs by means of quantifier raising or type shifting. And so, although my focus here is on language perception, in part because it is less clear how the cognition-*semantics* interface might work in production, there are some preliminary reasons to adopt a modular account of that process as well.

A second complication with the idea of semantics as the study of a modular system is that the semantic module needs to acquire (and then maintain and update) some of the information in its proprietary database. We aren't born representing the semantic value of 'dog'. How do we get it? I won't pretend to have an answer to this question; the study of semantic acquisition is nascent compared to the study of, e.g., syntactic acquisition. However, there are good reasons to expect that any plausible explanation of semantic acquisition will take the form of a computational theory of the kind offered by psycholinguists working within the framework of gen-

erative grammar.¹⁷ Theories in this tradition typically posit domain-specific learning mechanisms that are, to a significant extent, dissociated from the operations of central systems (Newport, 1990). One reason to anticipate that such an explanation will be required is that acquiring semantic competence with any vocabulary item requires coming to represent it in terms of centrally inaccessible concepts. In order to acquire the verb ‘to run,’ one must discern not only that it is used to talk about running; one must also represent its argument structure, its semantic type, etc. (Pinker and Jackendoff, 2005, §2.1). Since these are centrally inaccessible concepts, semantic acquisition cannot be a simple transfer of information from central cognition to the semantic module’s database. Still, it seems likely that there is *some* such transfer during acquisition.¹⁸ However, from the fact that information flows from central cognition into the semantic module via the acquisition mechanism, it does not follow that this information can affect real-time processing by the semantic module in its role as an input system, and my arguments in this section suggest that it does not.

2.2 Modularity and Constraint Semantics

What does the modularity of semantics have to do with the content–constraint debate that I outlined in §1? My argument, in brief, goes as follows. Content semantics presupposes that all of the information needed to identify the content of a sentence in a context of utterance is available to semantic composition. It is the contents of ‘he’ and ‘smokes’ that must be composed in order to deliver the content of ‘he smokes,’ for example. And so semantic composition must have access to whatever information is needed in order to identify these expressions’ contents. However, if semantic composition is a modular process—in particular, if it is an informationally encapsulated process—then it does not have access to all of the information that it would need in order to identify lexical expressions’ contents. The reason is that identifying the contents of semantically underspecified expressions—a process that I will call ‘content resolution’—is not the sort of thing that a cognitively impenetrable input system can accomplish for us, since any system that can accomplish it would need to draw on all manner of beliefs, memories, and other information that is proprietary to central cognition. If I am right that semantic composition is a modular process, and that content resolution is a central-cognitive process, then semantics can only deliver constraints, not contents.

Why think that content resolution depends on information that is proprietary to central cognition? In short, the reason is that understanding what someone has

¹⁷For a current survey of the kind of work I have in mind, see Guasti (2017).

¹⁸Bloom (2000, ch.3) and Hacquard and Lidz (2018) both present evidence that the semantic acquisition process recruits mindreading, for example.

said with a semantically underspecified expression—what they have referred to with a deictic use of a pronoun, or whether they are using ‘may’ in a deontic or epistemic sense, for example—requires integrating grammatically encoded information with information from the extralinguistic context, including information about other agents’ mental states. This kind of information is outside the purview of a modular input system and may need to be derived on the fly by means of complex, informationally unencapsulated inferences.

The claim that content resolution depends on a wide range of information and inference, and so can’t be done by an encapsulated system, is well rehearsed in the literature, and so I won’t belabor it here.¹⁹ Still, several points bear emphasis. First, notice that the claim applies even to the most straightforward indexical expressions, such as ‘I’ and ‘today’. In order to identify the content conveyed by an anonymous letter that reads, ‘today, I love you’, I would need to somehow infer the identity of the sender and the day on which it was sent, for example. Second, the claim does not depend on any particular view of what content resolution consists in. Some have argued that recognizing what a speaker says with semantically underspecified expressions is a matter of recognizing the speaker’s communicative or referential intentions.²⁰ Others have emphasized the importance of integrating grammatical information with the information that is shared by the interlocutors, such as the set of propositions that they commonly know or commonly accept for the purposes of the conversation.²¹ Either way, content resolution is a process that relies heavily on mindreading, which is a paradigmatic example of an unencapsulated, central-cognitive task.²²

I am not the first to argue that semantics is the study of a modular system, and that this has important implications for the nature of semantic values. One previ-

¹⁹See, for example, Bach (1987); Borg (2004, 2012); Carston (2002, 2012); Fodor (1983, 2002); Grice (1975); Neale (2004, 2005, 2007); Recanati (2004, 2010); Sperber and Wilson (1995).

²⁰Bach (1987, 1992); Heim (2008); Kaplan (1989a); King (2013, 2014, 2017); Michaelson (2013); Neale (2004, 2005, 2007).

²¹See, e.g., Clark (1996); Clark and Brennan (1991); Clark and Marshall (1992); Heim (1982, 1983); Hunter (2013); King (2013, 2014, 2017); Lepore and Stone (2015); Lewis (1979); Roberts (2002, 2003, 2005, 2012); Stalnaker (1973, 1978, 2014); Stojnić et al. (2017); Thomason (1990).

²²The alternative is to think that content resolution relies on information from a discourse context that is fully under grammatical control—see, e.g., Lepore and Stone (2015); Stojnić et al. (2017). This opens up the possibility that content resolution relies only on information available to a semantic module. But the cost of this grammaticalization is the proliferation of ambiguities whenever the prior discourse does not fully determine a discourse context that in turn determines a unique content for each semantically underspecified expression. In effect, advocates of this view treat semantic underspecification as ambiguity at the discourse level (thus proliferating ambiguities). Disambiguation is itself a central-cognitive task that typically involves mindreading—a point that these authors grant (Lepore and Stone, 2015, ch.13). Drawing out the lessons for these views of my arguments in this paper would take considerable space, and so I will bracket the issue for now.

ous attempt at this line of thought is due to Emma Borg (2004; 2012), who uses it to defend her version of semantic minimalism. 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. On her view, then, the semantic module is subject to roughly the same constraints that I have outlined here, and yet it invariably outputs a proposition when given the LF of a declarative sentence as an input.

The challenge for a view like Borg's is to explain how the semantic module can identify the contents of semantically underspecified expressions. Borg has defended a controversial answer to this challenge. As a case study, she considers a situation in which a speaker utters (15), thereby thereby saying that A is red.

(15) That is red.

According to Borg, demonstratives are directly referential singular terms whose referents are determined by speakers' intentions. A hearer's semantic module therefore lacks the resources to identify A as the referent of 'that' in a "substantive, nonlinguistic sense" (Borg, 2012, xviii). Still, Borg thinks that the semantic module has a metalinguistically mediated way of identifying the propositional content of (15).

I claim that the semantic content a hearer is required to entertain when faced with an utterance of a sentence like 'That is red' in a given context *c*, will contain a singular concept in subject-position, the content of which is exhausted by the object to which the speaker refers. However I also allow that all that is required, from the perspective of linguistic understanding, is that the hearer be able to think about that object under the token-reflexive description (which gives the character of the concept) 'the actual object referred to by the speaker with this token of "that"'. (Borg, 2012, xviii)

So, the semantic module identifies the content of 'that' by grasping a coreferential singular concept *C*, and the hearer grasps *C* by formulating a rigidified metalinguistic description of *C*'s referent. Since this description will be the same for every occurrence of 'that', Borg thinks of it as the character of *C*—a piece of information that can be combined with information about extralinguistic context to determine *C*'s referent.

Here is a response to this view that I find tempting. Despite Borg's insistence that sentential semantic values are propositions, the theory she actually articulates sounds much more like constraint semantics than content semantics, as I sketched those views in §1. After all, on her view, what the semantic module delivers when

given (15) as an input is a piece of context-neutral information about what the speaker has said, rather than the context-specific content of what they said. Indeed, on Borg's view, the information delivered by a hearer's semantic module when given (15) as input will be the same on every occasion, and will be the same whether the hearer is the well-informed addressee or an eavesdropper who is completely uninformed about the speaker's intentions and the extralinguistic context. A speaker who has this information still needs to integrate it with information from the context in order to identify (15)'s content "in a substantive sense". As Borg (2012, 135) puts it, "that A is the referent of this utterance is settled by features beyond the reach of semantics and, furthermore, to put this semantic content to use (i.e. to use it to inform one's dealings with the world) the hearer normally needs to go on to non-linguistically identify A". Still, she maintains that "as far as linguistic meaning or semantic content is concerned such issues of substantive object-identification are irrelevant".

To put things a different way, grasping Borg's minimal content for (15) does not give one enough information to determine its context-specific intension. An agent who grasps (15)'s minimal content needn't thereby possess enough information to determine, for a world w , whether (15)'s intension maps w to truth or falsity. For example, suppose that two candidate referents for 'that' in (15) are the hearer's nose and the painting on the wall, and that w_1 is a world where the nose is red but the painting isn't, whereas w_2 is a world at which the reverse is true. Now, imagine a hearer who is omniscient with respect to w_1 and w_2 . Still, grasping (15)'s minimal content would not give the hearer enough information to decide which of w_1 or w_2 is mapped to truth by (15)'s intension, since grasping the minimal content does not allow the hearer to determine whether their nose, the painting, or neither is the referent of 'that' on this occasion.²³

I suspect that Borg would reject this interpretation of her view. In explaining minimal semantics, she repeatedly stresses that the semantic value of a sentence is its propositional content, and not merely a partial characterization of that content (e.g., Borg 2012, 4–5). It is difficult to decide this issue concretely, since Borg does not tell us how to implement her view in a compositional-semantic framework of the kind I introduced in §1. The only implementation suggestion that she does tentatively endorse is that a semantic theory should pair sentences with conditional truth conditions, such as the following:²⁴

²³Garía-Carpintero (2013) likewise argues that the motivations and some details of Borg's view are in tension with her insistence that sentential semantic values are propositions.

²⁴Borg (2012, 135) takes (16) from Higginbotham (1994, 92–3). Conditional truth-condition views have also been endorsed by Burge (1974); Gross (2005); Heck (2014); Higginbotham (1988, 2002); Larson and Segal (1995).

- (16) If the speaker of ‘this is red’ refers with the utterance of ‘this’ therein to x and to nothing else, then this sentence, as uttered in this context, is true if and only if x is red.

But again, it seems to me that a theory that pairs sentences with outputs like (16) would be better described as a constraint semantics than a content semantics. After all, (16) is something that I currently know about every utterance of ‘this is red’ that has ever been produced, without knowing anything about those utterances’ extralinguistic contexts, and (16) certainly does not itself determine an intension for any particular utterance of ‘this is red’.

There are also some problems with the view that the semantic module outputs conditional truth conditions like (16). One is that it is not clear how to give a fully compositional theory that delivers conditional truth conditions—an issue that Borg does not address. Another arises from the fact that conditional truth conditions like (16) are formulated so as to mention each of the semantically underspecified expressions that they contain. But most semanticists think that at least some sentences contain semantically underspecified expressions (or other contextual parameters) that are not visible in sentences’ surface structure, and that ordinary speakers therefore aren’t aware of. Many of these semantically underspecified expressions can occur bound or unbound, and in their unbound occurrences can be used deictically. But since ordinary speakers are unaware of the existence of such expressions, their central cognitive system(s) can’t be expected to deal in representations of them as in (16). Moreover, there is no evidence that denying the existence of unpronounced expressions impairs one’s ability to speak—a fact best explained by saying that our representations of unpronounced expressions all occur within an encapsulated module.

Borg responds to such worries in two ways. The first is to deny the existence of many unpronounced expressions. The second is to argue that the unpronounced elements that she does posit always occur bound by a silent existential quantifier. For example, she argues that the LF of ‘Rob is ready’ is, roughly, $(\exists x)(\text{Rob is ready to do } x)$, and that its semantic value is the proposition that Rob is ready for something (2012, 205–6). It would be an understatement to say that each of these views puts Borg at odds with mainstream syntax and semantics, which has found reams of evidence for unpronounced context-sensitive elements in myriad constructions. To take just one example, null-subject languages make it possible to leave the subjects of many clauses unpronounced, but there is excellent evidence that such unpronounced subjects enter into a wide range of syntactic relations and can have both bound and deictic occurrences, much like overt pronouns.²⁵ But if the semantic orthodoxy is correct that some unpronounced elements have deictic occurrences,

²⁵For a recent overview of the evidence for null subjects and their syntactic and semantic properties,

then Borg’s account implausibly requires speakers to represent these expressions in central cognition.

As I see it, Borg would be better off to give up her basic assumption that “every well-formed sentence, perhaps relativized to a context of utterance, is capable of expressing a proposition” (Borg, 2012, xiv). Although Borg defends propositionalism against various attacks (none of which I have discussed here), she offers few positive motivations for the view. One is that propositionalism has been presupposed by traditional work in formal semantics (Borg, 2012, 6–7).²⁶ My response to this argument will come in §3, where I show that it is relatively simple to modify the standard framework to output constraints rather than contents. Borg’s second motivation for propositionalism is that “it is natural” to think of sentences as conveying information, referring to contents when embedded in attitude constructions, standing in logical relations, and conveying reasons for belief (Borg, 2012, 7–8). But of course, it is equally natural to say that *speakers use sentences* to convey information, to refer to propositions when uttered in embedded contexts, to express contents that stand in logical relations, and to convey reasons for belief, and the latter description is compatible with my view here.²⁷

The failure to question content semantics has led others who are sensitive to

see Camacho (2013). For a helpful discussion of the evidence for other implicit contextual parameters that needn’t occur bound, see Glanzberg (2016).

²⁶Borg does not mention the well-established traditions in formal semantics, such as dynamic semantics, discourse representation theory, and variable-free semantics, that reject propositionalism for reasons unrelated to my arguments or position here.

²⁷More recently, Borg (2017) has argued that minimal content is needed to explain the truth-value judgments that we make when speaker’s strict commitments are at issue. She argues, for example, that lying entails failing to believe the minimal content of one’s utterance, and that semantic content “must be truth-evaluable (rather than a mere propositional radical or incomplete logical form)” because holding speakers accountable for lying involves judgments about “speakers holding minimal contents true or false” (Borg, 2017, 24n26). But it would be straightforward for a constraint semanticist to mimic Borg’s predictions. I could maintain that when we utter a sentence whose semantic value is a constraint, we are held strictly responsible for having said at least one proposition that meets this constraint. For example, Borg’s view predicts that I will be held culpable for lying if I utter ‘Simon is ready’ but fail to believe that there is something for which Simon is ready. My alternative proposal predicts that I will be held culpable for lying if there is no proposition p such that: for some eventuality φ , Simon is ready for φ , and I believe p . (This is only a slightly stronger prediction than Borg’s.)

So I *could* mimic Borg’s prediction. But I don’t think I *should*, because I think Borg gets our judgments about lying wrong. Imagine a case in which John believes that his brother, Simon is ready for something (say, playing video games), but knows that Simon is not ready for his final exams. The following exchange between John and their father takes place:

DAD: Is Simon ready for his final exams?

JOHN: Simon is ready.

John is clearly lying, but Borg’s view does not predict this. In at least some cases, the content of a lie is not one of Borg’s minimal propositions.

considerations about mental architecture either to conclude that compositional semantics, as usually conceived, is impossible (e.g. Fodor 1989; 1998; 2002), or to give up on the modularity of semantics and conclude that the composition of expressions' contents is actually a messy, *pragmatic* process that is performed by central cognition (Recanati, 2010). But again, each of these views is rooted in the erroneous assumption that semantic composition must deal in contents rather than constraints.

I wish to conclude, instead, that semantics should deliver constraints. One advantage of this view is that, unlike its competitors, it allows for a precise semantics–pragmatics interface that is grounded in underlying facts about human cognitive architecture. The semantics–pragmatics interface turns out to be an instance of the interface between central cognition and its peripheral input–output systems. If this is also where we locate the perception–cognition interface—an admittedly controversial view²⁸—then the semantics–pragmatics interface also turns out to be coincident with the interface between language perception and general-purpose cognition.

3 Constraint Semantics

My goal in this remaining section is to develop a constraint semantics minimally adjusting the standard framework, but without giving up on what makes this framework so successful.²⁹ I will proceed in stages, first showing how to give a compositional semantics of the overall kind I am after, then showing how to refine this theory to handle a wide range of semantically underspecified expressions.

Before I begin, a pair of methodological points:

First, my strategy is to show how to arrive at a constraint semantics by conservatively modifying the most well-known semantic framework available. This is almost certainly not the most elegant way to develop a compositional theory of constraints. It seems quite likely that heterodox frameworks could be made to deliver the semantic values I seek in more principled and less roundabout ways.³⁰ However, the point

²⁸For influential defenses of the claim that visual perception (and perhaps perception more broadly) is informationally encapsulated, see Firestone and Schöll (2015); Fodor (1983); Pylyshyn (1984, 1999). For criticisms of this view, see Churchland (1988); Collins and Olsen (2014); Dunning and Balçetis (2013); Hohwy (2013); Prinz (2006); Vetter and Newen (2014).

²⁹The theory presented here also inherits some limitations in its explanatory scope from the standard framework as regards, e.g., presupposition projection and discourse anaphora. There isn't space for adequate engagement with these phenomena here. I hope to address them in future work.

³⁰For example, alternative semantics provides us with tools that could be used to directly compose underspecified subsentential semantic values (Ciardelli et al., 2017; Kratzer and Shimoyama, 2002); variable-free semantics provides us with methods of delivering underspecified sentential semantic values without proliferating variables (Jacobson, 1999, 2014); and Pietroski (2005; 2018; 2008) has

of my approach is to demonstrate that radical departures from semantic orthodoxy aren't necessary in order to achieve a result that is consistent with the foundational lessons I have drawn.

Second, the standard way of arguing that we should revise our theory of the semantic values of a given kind of expression is to show that this would best explain our truth-value judgments about sentences that embed expressions of this kind. For example, I will be arguing that declarative sentences are type- $\langle st, t \rangle$ properties of propositions rather than the standard type- st propositions. The usual strategy would be to show that this view is needed to make sense of sentences containing embedded declaratives. This will not be my strategy here, although I agree that it is worth exploring in connection to my position. Some readers might therefore be tempted to complain that I have not provided adequate evidence for my position. To these readers I reply that they are missing the overarching point of this essay, which is that integrating semantics into broader cognitive science means finding new sources of evidence to constrain semantic theorizing. It should not be surprising, then, that my reasoning falls outside the usual practices of semanticists.

3.1 Semantic Values as Characters?

I have been arguing that the semantic value of a sentence is not its context-relativized content, but is something context-neutral that specifies the range of contents that can be said by using the sentence literally. How should a semantic value of this kind be represented?

One obvious possible answer to this question is that semantic values are Kaplanian characters—functions from contexts to contents. Instead of the usual context-parameterized content, (17), for example, the semantic value of ‘it stinks’ could be given as (18).

$$(17) \llbracket \text{it}_1 \text{ stinks} \rrbracket^c = \lambda w . g_c(1) \text{ stinks at } w$$

$$(18) \llbracket \text{it}_1 \text{ stinks} \rrbracket = \lambda c . \lambda w . g_c(1) \text{ stinks at } w$$

Here is a simple way of achieving this result.³¹ We leave the compositional machin-

developed a compositional-semantic framework that eschews contents as semantic values (along with many other standard orthodoxies) for cognitive-scientific reasons that complement those I have presented here. In future work I hope to explore how tools from these alternative approaches may better realize the foundational vision I have in mind.

³¹Another option would be to assign characters as lexical semantic values and reformulate our composition principles to act directly on characters themselves. We could reformulate Functional Application as follows, for example (cf. Charlow (2017)):

$$(FA^*) \llbracket \alpha\beta \rrbracket = \lambda c . \lambda w . \llbracket \alpha \rrbracket(c)(w)(\llbracket \beta \rrbracket(c)(w)) \quad \text{or} \quad \lambda c . \lambda w . \llbracket \beta \rrbracket(c)(w)(\llbracket \alpha \rrbracket(c)(w))$$

ery of content semantics in place, but posit a single extra step at the end of every compositional derivation. This extra step would be an operation that simply abstracts over the context parameter, thus transforming (17) into (18), for example.

On this view, 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, central cognition must identify the context in which it was uttered and apply the character to it in order to identify the utterance content. Content resolution is just the process of identifying the right context and plugging it into a sentence's character in this way—a central-cognitive task that happens after the semantic module has done its work. The semantic module still has to work with representations of contexts 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 a simple and attractive picture, and my positive proposal will resemble it in some ways. However, I don't think that it can be correct. My reason is that this proposal requires language users' central-cognitive systems to be capable of working with representations that are framed in terms of centrally inaccessible concepts, such as assignment functions and numerical indices. For example, if my central system is given (18), and if contexts of utterance are the sorts of things that can be mapped to propositions by (18), then identifying a context of utterance for 'it stinks' and deriving a proposition from it will require me to have beliefs (or other 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 this is a bad consequence for two reasons. First, it is implausible on its face that being a competent language user requires having beliefs or other attitudes about assignment functions and the indices in their domain. Second, and more importantly, 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 (18) 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 (18) 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.

I am not optimistic about this view. The precise mathematical structures of variable assignments are crucial to the semantic roles they play. It is implausible that our central-cognitive conceptions of speakers’ referential intentions are so well behaved as to be modeled by such precise mathematical objects. At the very least, articulating this view adequately would require saying considerably more than I have said here. But perhaps that could be done.

Still, I think this view would raise a new and closely related problem, because it would require language users to have beliefs and intentions 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. I mentioned null subjects in my criticism of Borg in §2.2, but semanticists have also posited many other unpronounced elements that are sensitive to context. If understanding speech does require representing these null elements, 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 represent sentences as containing them nonetheless. A plausible explanation of this fact is that null expressions are represented only at intermediate stages of processing within the semantic module, and do not feature in the module’s outputs. But in that case, our strategy for reinterpreting (18) in terms of speakers’ referential intentions leads us to yet another inadequate mixture of central and modular representations.

3.2 First Steps

My positive proposal is an attempt to capture the virtues of the foregoing account while positing sentential semantic values of a kind that both the semantic module and central cognition has the resources to handle. Specifically, I will assume that the outputs of semantics include representations of individuals, worlds, properties, relations, and propositions, but not of assignment functions, indices, or linguistic expressions 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:³²

³² Among the things that this semantic value abstracts away from are clause-type, tense, aspect, and

$$(19) \llbracket \text{it stinks} \rrbracket = \lambda p_{st} . (\exists x)(p = \lambda w . x \text{ stinks at } w)$$

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. This is a constraint on what can be said literally with ‘it stinks’. If I hear someone utter ‘it stinks’, and if I assume they’re speaking literally but don’t know anything about the context or their referential intentions on this occasion, all my semantic module tells me about what they’ve said is that it has this property. To work out what they’ve said, I need further, extralinguistic information.³³

How can semantic values like (19) be systematically derived? The following method is inelegant but simple. It is also similar to the method I described in §3.1 for computing sentences’ characters, in that it piggybacks on the computational mechanisms of the standard framework of content semantics, but then removes the undesirable features of that framework’s sentential semantic values at the last moment, via an abstraction operation. So, what’s needed is an abstraction operation that takes us from (17) to (19).³⁴

$$(17) \llbracket \text{it}_1 \text{ stinks} \rrbracket^g = \lambda w . g(1) \text{ stinks at } w$$

$$(19) \llbracket \text{it}_1 \text{ stinks} \rrbracket = \lambda p_{st} . (\exists x)p = \lambda w . x \text{ stinks at } w$$

This operation must remove the assignment dependency of its input while appending what I call a *prenex* to its front.

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

I’ll use the term ‘prenex quantifiers’ for the existential quantifiers contained within a prenex. Generalizing from this example, what’s needed is an abstraction operation that respects the following constraint: for any expression α , and for each numerical

the φ -features of ‘it’. I’ll show how to build the φ -features back in in §3.3.

Note that I leave the standard notation as it is, and write my preferred kind of semantic values using triple brackets. This notation is borrowed from Charlow (2017), who uses it for something else.

³³Some 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 (19) for example.

$$(19)' \llbracket \text{it stinks} \rrbracket = \{p : (\exists x)(p = \lambda w . x \text{ stinks at } w)\}$$

The intuitive idea 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. I have to pragmatically infer *which* of these propositions is the content of what they said.

³⁴For now, I ignore the context parameter and indexicals. I will show how to handle indexicals in §3.4.

index i to which at least one free variable in α is indexed, there will be a quantifier in the prenex of $\llbracket\alpha\rrbracket$ that binds all occurrences of a variable x , with one occurrence of x in $\llbracket\alpha\rrbracket$ substituted for each occurrence of $g(i)$ in $\llbracket\alpha\rrbracket^g$. A preliminary version of the sort of operation we need can be given as follows:³⁵

(21) *Constraint Abstraction* (preliminary version)

If α dominates unbound variables $v_i \dots v_n$ and $\llbracket\alpha\rrbracket^g \in D_\tau$

Then $\llbracket\alpha\rrbracket = \lambda p_\tau . (\exists x^i) \dots (\exists x^n) . p = \llbracket\alpha\rrbracket^{g^{i \rightarrow x^i \dots n \rightarrow x^n}}$

Constraint Abstraction is a general tool for converting assignment-relativized intensions into constraints of the kind that I have argued are the outputs of the semantic module. Assignment-relativized intensions of any semantic type can be converted in this way. However, it is sufficient for my purposes to assume that the semantic module applies this operation to the intensions of LFs' root nodes as the last step before sending the result as the module's output to central cognition.

So far I have offered only a sketch of an account, and some bugs remain to be fixed. Nonetheless, it can already be seen that the role of indices and assignment functions has been substantially reduced and demystified as compared to their role in standard content semantics. For one thing, the semantic values I have posited aren't relativized to assignment functions, or contexts, or anything else. Indices and assignments still play a role at intermediate levels of the compositional derivation. But since every variable is eventually bound, either in the usual way or by a prenex quantifier, indices and assignment functions play a mere bookkeeping role, coordinating binding relations until the final semantic value is constructed. In fact, it doesn't matter what is in the range of assignment functions, since they are always eventually discharged. It would be fine, for example, if all unmodified assignments were reflexive, simply mapping numerical indices to themselves. The values they assign are mere placeholders.

3.3 The Meanings of Variables

The semantics I've sketched so far is on the right track, but it is still problematic and underdeveloped. One problem is that my preliminary version of constraint abstraction is not compositional in the strict sense abided by the standard framework. In order to determine an output semantic value $\llbracket\alpha\rrbracket$, this principle needs *syntactic* information about which variables are contained within α , and not merely *semantic* information about its input semantic value, $\llbracket\alpha\rrbracket^g$. A second problem is that it

³⁵Notation: subscripts on variables are numerical indices, as usual. Superscripts on variables are merely for disambiguation.

remains unclear how to distinguish the semantic contributions of different semantically underspecified expressions. In this section, I'll address these problems.

I 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 to say that they trigger presuppositions that render their semantic values undefined relative to certain assignments. For any i and g , $\llbracket \text{he}_i \rrbracket^g$ is defined only if g maps i to a male, for example. Nodes inherit the presuppositions of unbound pronouns they dominate. And variable binders come with their own presuppositions that enforce feature agreement with the presuppositions of the variables they bind.³⁶ However, by binding variables that are free at LF, the prenex quantifiers that I have posited eliminate their assignment sensitivity, thereby nullifying presuppositions triggered by their gender features. So I need a different way of distinguishing the meanings of 'he' and 'she' in their unbound occurrences.

I can give a sense of what is needed by considering what the semantics should say about sentences containing free occurrences of 'he' and 'she'. If all I know about your utterance is that you used 'he smokes' literally, what can I know about what you've said? Roughly, that it is 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.

$$(22) \llbracket \text{he smokes} \rrbracket = \lambda p_{st} . (\exists x_e : x \text{ is male})(p = \lambda w . x \text{ smokes at } w)$$

$$(23) \llbracket \text{she smokes} \rrbracket = \lambda p_{st} . (\exists x_e : x \text{ is female})(p = \lambda w . x \text{ smokes at } w)$$

The trick, then, is to find a way to distinguish the meanings of 'he' and 'she' in a way that allows them to contribute different restrictions to the prenex quantifiers that bind them after constraint abstraction. 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 literally to refer to it. The constraint properties 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.

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:

$$(24) \mu(\text{he}) = \lambda x_e . x \text{ is male}$$

³⁶This account comes in several versions, some of which also marshal syntactic resources to explain some cases of agreement. See, e.g., Heim (2008); Sauerland (2008).

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

Because variables' constraint properties are part of the information that one's semantic module must have in order for one to qualify as a competent speaker, I adopt the following principle that expands on the role of the lexicon as it applies to variables.

(26) *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 each pronoun's semantic value in the form of a presupposition requirement on admissible assignment functions. I propose that this information is also stored in the variable's constraint property, which encodes part of the evidence that speakers offer to hearers when uttering an unbound occurrence of the variable, but that does not show up in the content of what speakers say with it. When I say 'he smokes', it is not part of the content of what I say that I am talking about a male; rather, this information is a clue I give my addressee in order to help them to recover my intended content.

What I need now is a compositional semantic theory that keeps track of the constraint properties of all of the variables in an expression, so that my new-and-improved Constraint Abstraction principle can install them as restrictors on the appropriate prenex quantifiers. I will accomplish this by assigning every expression a double-barreled semantic value, as follows:³⁷

(27) For every node α in every LF, $\llbracket \alpha \rrbracket^{w,g} = \langle \llbracket \alpha \rrbracket_1^{w,g}, \llbracket \alpha \rrbracket_2 \rangle$

The two coordinates within an expression's semantic value track different information about its meaning. The first coordinate, $\llbracket \alpha \rrbracket_1^g$, is α 's semantic value according to the standard framework—a world- and assignment-relativized content. The second coordinate, $\llbracket \alpha \rrbracket_2$, keeps track of the constraint properties of any variables in α . Constraint Abstraction, as I will define it below, puts these two values together into a single output semantic value, $\llbracket \alpha \rrbracket$.³⁸

Because only variables have constraint properties, and $\llbracket \cdot \rrbracket_2$ stores information about constraint properties, only variables are assigned substantive values by $\llbracket \cdot \rrbracket_2$.

³⁷My implementation is loosely inspired by Cooper's (1975; 1983) method of storing and retrieving the semantic values of object-position quantifiers, though the present account is entirely different in its theoretical aims.

³⁸That my two dimensions of meaning get reassembled in this way sets my theory apart from other recent multidimensional approaches to semantics (e.g. Potts 2005; McCready 2010; Gutzmann 2017).

(28) For any variable v and numerical index i , $\llbracket v_i \rrbracket_2 = \{\langle i, \mu(v) \rangle\}$

(29) For any non-variable lexical item α , $\llbracket \alpha \rrbracket_2 = \emptyset$.

On the other hand, we can stipulate that $\llbracket \cdot \rrbracket_1^{w,g}$ is identical to the interpretation function of the standard framework.³⁹

I still need an account of how my double-barreled semantic values compose. In practice, this means that I need an account of how the values of both $\llbracket \cdot \rrbracket_1^{w,g}$ and $\llbracket \cdot \rrbracket_2$ compose. Take the latter first. The value of $\llbracket \cdot \rrbracket_2$ for each non-variable lexical item is the empty set, and its value for each variable is a singleton set containing an index–constraint pair. At any given complex node in an LF, we need $\llbracket \cdot \rrbracket_2$ to store the constraint properties of whatever variables the node dominates. A first stab at this could be accomplished by taking $\llbracket \alpha\beta \rrbracket_2$ to be $\llbracket \alpha \rrbracket_2 \cup \llbracket \beta \rrbracket_2$. There is just one complication: if α contains multiple coindexed variables, we’ll want $\llbracket \alpha \rrbracket_2$ to contain just one corresponding index–constraint pair, with the coindexed variables’ constraint properties conjoined. The following operation collates and combines index–constraint pairs in this way.

(30) *Definition of \uplus*

For any numerical indices $x^1 \dots x^n, y^1 \dots y^m$ and functions $X^1 \dots X^n, Y^1 \dots Y^m$, if $\chi = \{\langle x^1, X^1 \rangle \dots \langle x^n, X^n \rangle\}$ and $\gamma = \{\langle y^1, Y^1 \rangle \dots \langle y^m, Y^m \rangle\}$, then $\chi \uplus \gamma$ is defined as follows:

For every $\langle x^i, X^i \rangle \in \chi$ and $\langle y^j, Y^j \rangle \in \gamma$:

- if $x^i = y^j$ and $\exists \tau : X^i, Y^j \in D_{\tau\tau}$ then $\langle x^i, \lambda y_{\tau} . X^i(y) \wedge Y^j(y) \rangle \in \chi \uplus \gamma$
- if $x^i \neq y^j$, then $\langle x^i, X^i \rangle \in \chi \uplus \gamma$ and $\langle y^j, Y^j \rangle \in \chi \uplus \gamma$

Nothing else is in $\chi \uplus \gamma$.

This operation takes two sets of index–constraint pairs, checks the indices against one another, conjoins constraints that are paired with matching indices, and otherwise unions the sets. This operation is how values of $\llbracket \cdot \rrbracket_2$ compose at branching nodes.

Values of $\llbracket \cdot \rrbracket_1^{w,g}$ compose by exactly the same principles posited by the standard framework. This is the sense in which my account here is conservative and non-destructive: the standard framework continues to be a true theory of $\llbracket \cdot \rrbracket_1^{w,g}$. I have merely added a theory of constraint properties to make sense of the meanings of unbound variables as well. Here, for example, is a version of functional application that will work for my purposes.

³⁹I stipulate that for any expression α , $\llbracket \alpha \rrbracket^g = \langle \llbracket \alpha \rrbracket_1^g, \llbracket \alpha \rrbracket_2 \rangle = \langle \lambda w. \llbracket \alpha \rrbracket_1^{w,g}, \llbracket \alpha \rrbracket_2 \rangle$.

(31) *Functional Application*

$\llbracket \alpha\beta \rrbracket^{w,g} = \langle \llbracket \alpha \rrbracket_1^{w,g}, \llbracket \beta \rrbracket_2 \rangle$, such that:

- $\llbracket \alpha \rrbracket_1^{w,g} = \llbracket \alpha \rrbracket_1^{w,g}(\llbracket \beta \rrbracket_1^{w,g})$ or $\llbracket \beta \rrbracket_1^{w,g}(\llbracket \alpha \rrbracket_1^{w,g})$ (whichever is defined); and
- $\llbracket \alpha\beta \rrbracket_2 = \llbracket \alpha \rrbracket_2 \uplus \llbracket \beta \rrbracket_2$

Other composition principles can be straightforwardly derived from the standard ones in analogous ways.

Lastly, here is my updated version of Constraint Abstraction, which works by combining the values of $\llbracket \cdot \rrbracket_1^{w,g}$ and $\llbracket \cdot \rrbracket_2$ at a given node, in order to create an appropriate output semantic value of the kind that I outlined at the beginning of this section.

(32) *Constraint Abstraction* (final version)

If $\llbracket \alpha \rrbracket_2^g = \{ \langle i, f^1 \rangle \dots \langle j, f^n \rangle \}$ and $\llbracket \alpha \rrbracket_1^g \in D_\tau$

Then $\llbracket \llbracket \alpha \rrbracket \rrbracket = \lambda p_\tau . (\exists x^1 : f^1(x^1)) \dots (\exists x^n : f^n(x^n))(p = \llbracket \alpha \rrbracket_1^{g^{i \rightarrow x^1 \dots j \rightarrow x^n}})$

It is straightforward to handle a range of 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, the constraint property will be nearly vacuous, consisting only of a specification of the variable's semantic type. If we think of DPs as including a hidden, type-*et* domain-restrictor variable, then its constraint property can be given as follows, for example.

$$(33) \mu(dom) = \lambda x_{et} . x$$

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. Glanzberg (2016) makes a similar point about the degree (or standard) variables usually posited to understand gradable adjectives.

On the other hand, Glanzberg argues that some other unpronounced, semantically underspecified expressions are richly semantically constrained. He posits an experiencer parameter in the semantic representations of taste predicates like 'fun' and 'tasty', for example, and argues that it can be saturated only by individuals or groups who occupy the thematic role of an experiencer in the event being described, and whose perspective is being adopted by the speaker (Glanzberg, 2007, 2016). Glanzberg does not show how he thinks this idea should be implemented in a compositional semantics, but I believe his ideas could be naturally fleshed out by thinking of his experiencer parameter as a variable with a relatively rich constraint property.⁴⁰

⁴⁰Glanzberg (2016) distinguishes between "functional" and "thematic" contextual parameters, ar-

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 in keeping with the spirit of constraint semantics as laid out in §1, I would like to keep it that way.

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 (1998, 244) 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 free variables’ semantic values are assignment-dependent (and, so, indirectly context-dependent) contents.⁴¹

guing that they differ in a range of syntactic, semantic, metasemantic, and pragmatic respects, and that these differences can be traced to the fact that functional parameters are projected (or otherwise invoked) by functional expressions whereas thematic parameters are projected by lexical expressions. This might be taken as a reason not to treat all semantically underspecified expressions uniformly, as variables, as I have done here. In particular, Glanzberg repeatedly suggests that while functional parameters are like variables, his thematic parameters are more like referring expressions. However, Glanzberg’s use of ‘variable’ is clearly narrower than mine here. As in the standard framework, I use ‘variable’ to mean any expression that may occur bound; my paradigmatic examples are personal pronouns. Glanzberg argues that thematic parameters, like his functional parameters, can be bound or free, and that thematic parameters work roughly like overt pronouns. Glanzberg also argues that functional parameters, unlike thematic parameters, have an “indirect metasemantics”, meaning that contextual factors other than the speaker’s intentions play a role in fixing their contents. I follow King (2013; 2014; 2017) in doubting that such a metasemantic distinction is warranted. In any case, the view I have advanced here places such considerations outside of semantics proper. What Glanzberg calls ‘metasemantics’ becomes the issue of how speakers and hearers bridge the partial and defeasible evidence provided by an expression’s semantic value and what the speaker used the expression to say—a central-cognitive process that may draw on all manner of information.

⁴¹One 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]_i believes that you_i 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; Partee, 1989; 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 like ‘we Canadians’ and ‘you people’, is to treat pronouns as the determiners of DPs whose nominals are often unpronounced, 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.

On my view, indexicals and demonstratives differ semantically from other variables 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.

(34) $\mu(I) = \lambda x_e . x$ is the speaker of the utterance being interpreted

(35) $\mu(\text{you}) = \lambda x_e . x$ is the addressee of the utterance being interpreted⁴²

(36) $\mu(\text{today}) = \lambda x_e . x$ is the day on which the utterance being interpreted was produced

(37) $\mu(\text{now}) = \lambda x_e . x$ is an interval of time during which the utterance being interpreted was produced

(38) $\mu(\text{here}) = \lambda x_e . x$ is a location in space within which the utterance being interpreted was produced⁴³

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 SPEAKER_c , which is just the speaker themselves. But, as I argued in §2, the semantic module lacks access to information about who the speaker is. Identifying the speaker of an utterance is a central-cognitive task that may require who-knows-what extralinguistic information. Since the semantic module can’t identify the speaker on its own, it can’t plug in the speaker as the semantic value of an expression.

Instead, the occurrences of ‘the utterance being interpreted’ in (34)–(38) 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

⁴²This entry deals only with singular ‘you’, and would need to be generalized in some way to handle plural senses of ‘you’.

⁴³There 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.

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 may turn out to be mistaken. Consider, for example, the well worn example 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. One property that the semantic module might thus misattribute 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. I take this phenomenon to be a genuine perceptual illusion wherein an agent's input systems persist in attributing properties to a stimulus that, the agent knows full well, the stimulus does not have.

What about the demonstratives, 'this' and 'that'? If 'this' and 'that', in their "bare" uses, are unstructured referring expressions, then their constraint properties can be given roughly as follows.⁴⁴

(39) $\mu(\text{this}) = \lambda x_e . e$ is "proximal" from the perspective of the speaker of the utterance being interpreted

(40) $\mu(\text{that}) = \lambda x_e . e$ 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.

Conclusions

There is much more to say about how the project of this essay could be expanded, refined, and used to understand further phenomena near the semantics–pragmatics

⁴⁴For arguments that demonstratives ought to be thought of as variables, see Nowak (2016, MS).

interface.⁴⁵ But I have run out of space.

For now, I hope I have made a plausible case for the following conclusions. Compositional semantics is the project of reverse-engineering the proprietary database of a centrally inaccessible and informationally encapsulated input–output system. This system’s job is to encode and decode partial and defeasible evidence of what speakers use sentences to say. Hearers must rely on central-cognitive resources to integrate this evidence with whatever else is available in order to identify what the speaker has said, as well as what they have otherwise meant. We can adopt this view without sacrificing a rigorous formal semantics. And, if we do, then the semantics–pragmatics interface turns out to be precise and grounded in underlying features of human cognitive architecture.

References

- Abney, S. P. (1987). *The English Noun Phrase in its Sentential Aspect*. PhD thesis, MIT, Cambridge, MA.
- Alonso-Ovalle, L. (2006). *Disjunction in Alternative Semantics*. PhD thesis, University of Massachusetts, Amherst.
- Anand, P. and Nevins, A. (2004). Shifty operators in changing contexts. *Proceedings of SALT*, 14:20–37.
- 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.
- Bloom, P. (2000). *How Children Learn the Meanings of Words*. MIT Press.
- Borg, E. (2004). *Minimal Semantics*. Oxford University Press.
- Borg, E. (2012). *Pursuing Meaning*. Oxford University Press.
- Borg, E. (2017). Explanatory roles for minimal content. *Noûs*, Early View.
- Burge, T. (1974). Demonstrative constructions, reference, and truth. *Journal of Philosophy*, 71(7):205–223.
- Camacho, J. A. (2013). *Null Subjects*. Number 137 in Cambridge Studies in Linguistics. Cambridge, Cambridge, UK.
- 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,

⁴⁵For example, some have argued that all or nearly all open-class vocabulary items, including common nouns and adjectives, are context-sensitive on grounds of polysemy and microlinguistic negotiability (e.g. Wilson and Carston 2007; Travis 2008; Ludlow 2014). Does this mean that they should be assigned thin semantic values that get fleshed out by central cognition? I am sympathetic to this idea, and in this respect I am sympathetic to proposals from Glanzberg (2014) and Gross (2005) to pass along the job of selecting contents for many open-class expressions to nonlinguistic systems. In the present system, this could be handled by treating, e.g., common nouns as type-*et* variables with thin constraint properties. I do not have space to articulate and defend this view here, but I plan to return to it in future work. Another topic for future work is the question of how the present system could make sense of context-shifting operators (what Kaplan (1989b) calls “monsters”), such as those posited by, e.g., Anand and Nevins (2004); Santorio (2012); Schlenker (2003). One option is to adopt a modified version of Rabern and Ball’s (2017) method for formalizing monstrous contents without using context-parameter-shifting operators. Another option would be to think of monsters as constraint shifters—operators that modify variables’ constraint properties.

- G., editors, *The Handbook of Pragmatics*, chapter 28, pages 633–656. Blackwell, Oxford.
- Carston, R. (2012). Word meaning and concept expressed. *Linguistic Review*, 29(4):607–623.
- Charlow, S. (2017). A modular theory of pronouns and binding. In *Proceedings of LENS 14*.
- Churchland, P. (1988). Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science*, 55:167–187.
- Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2013). Inquisitive semantics: a new notion of meaning. *Language and Linguistics Compass*, 7(9):459–476.
- Ciardelli, I., Roelofsen, F., and Theiler, N. (2017). Composing alternatives. *Linguistics and Philosophy*, 40:1–36.
- Clark, H. H. (1996). *Using Language*. Cambridge University Press, Cambridge, UK.
- Clark, H. H. and Brennan, S. E. (1991). Grounding in communication. In Resnick, L., Levine, B. M., and Teasley, S. D., editors, *Perspectives on Socially Shared Cognition*, pages 127–149. American Psychological Association.
- Clark, H. H. and Marshall, C. R. (1992). Definite reference and mutual knowledge. In Clark, H. H., editor, *Arenas of Language Use*, chapter 1, pages 9–59. University of Chicago Press and Center for the Study of Language and Information.
- Collins, J. A. and Olsen, I. R. (2014). Knowledge is power: How conceptual knowledge transforms visual cognition. *Psychonomic Bulletin and Review*, 21:843–860.
- Cooper, R. (1975). *Montague's Semantic Theory and Transformational Syntax*. PhD thesis, University of Massachusetts, Amherst, MA.
- Cooper, R. (1983). *Quantification and Syntactic Theory*. Reidel, Dordrecht.
- Cosmides, L. and Tooby, J. (1992). Cognitive adaptations for social exchange: Has natural selection shaped how humans reason. In Barkow, J., Cosmides, L., and Tooby, J., editors, *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford University Press.
- Cumming, S. (2008). Variabilism. *The Philosophical Review*, 117(4):525–554.
- Devitt, M. (2006). *Ignorance of Language*. Oxford University Press.
- Devitt, M. and Sterelny, K. (1999). *Language and Reality: An Introduction to the Philosophy of Language*. Blackwell, Oxford, UK, second edition.
- Dunning, D. and Balcetis, E. (2013). Wishful seeing: How preferences shape visual perception. *Current Directions in Psychological Science*, 22(1):33–37.
- Elbourne, P. (2005). *Situations and Individuals*. MIT Press, Cambridge, MA.
- Fara, D. G. (2015). Names are predicates. *Philosophical Review*, 124(1):59–117.
- 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*. Unpublished Lecture Notes, online at <http://web.mit.edu/fintel/fintel-heim-intensional.pdf>, spring 2011 edition.
- 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.
- Garía-Carpintero, M. (2006). Recanati on the semantics/pragmatics distinction. *CRÍTICA, Revista Hispanoamericana de Filosofía*, 38(112):35–68.
- Garía-Carpintero, M. (2013). review of 'pursuing meaning' by Emma Borg. *Mind*, 122(486):515–521.
- Glanzberg, M. (2007). Context, content, and relativism. *Philosophical Studies*, 136(1):1–29.
- 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.
- Glanzberg, M. (2016). Not all contextual parameters are alike. *Unpublished Manuscript*.
- Grice, H. P. (1975). Logic and conversation. In Cole, P., editor, *Syntax and Semantics 3: Pragmatics*. Academic Press.
- Gross, S. (2005). Context-sensitive truth-theoretic accounts of semantic competence. *Mind & Language*, 20(1):68–102.
- Guasti, M. T. (2017). *Language Acquisition: The Growth of Grammar*. MIT Press, Cambridge, MA, second edition.
- Gutzmann, D. (2017). *Use-Conditional Meaning: Studies in Multidimensional Semantics*. Oxford University Press, Oxford, UK.
- Hacquard, V. and Lidz, J. (2018). Children's attitude problems: Bootstrapping verb meaning from syntax and pragmatics. *Mind & Language*, online first:1–24.
- Heck, R. (2014). Semantics and context-dependence: Towards a Strawsonian account. In Sherman, B. and Burgess, A., editors, *Metasemantics: New Essays on the Foundations of Meaning*, pages 327–364. Oxford University Press, Oxford, UK.
- 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.
- Higginbotham, J. (1988). Context, models, and meaning: A note on the data of semantics. In Kempson, R., editor, *Mental Representations: The Interface Between Language and Reality*, pages 29–48. Cambridge University Press, Cambridge, UK.
- Higginbotham, J. (1994). Priorities in the philosophy of thought. *Proceedings of the Aristotelian Society, Supplementary Volume*, 68:85–106.
- Higginbotham, J. (2002). Competence with demonstratives. *Philosophical Perspectives*, 36(16):1–16.
- Hohwy, J. (2013). *The Predictive Mind*. Oxford University Press, Oxford, UK.
- Hunter, J. (2013). Presuppositional semantics. *Journal of Semantics*, 30(3):381–421.
- Jacobson, P. (1999). Towards a variable-free semantics. *Linguistics and Philosophy*, 22(2):117–184.
- Jacobson, P. (2014). *Compositional Semantics: An Introduction to the Syntax/Semantics Interface*. Oxford University Press.
- Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Kluwer, Dordrecht.
- 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.
- 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. (2017). Strong contextual felicity and felicitous underspecification. *Philosophy and Phenomenological Research*, online first.
- 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*.
- Kratzer, A. and Shimoyama, J. (2002). Indeterminate pronouns: The view from Japanese. In Otsu, Y., editor, *Proceedings of the Third Tokyo Conference on Psycholinguistics*, pages 1–25. Hituzi Syobo, Tokyo.
- 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.
- Lepore, E. and Stone, M. (2015). *Imagination and Convention*. Oxford University Press, Oxford.
- Lewis, D. K. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(3):339–359.
- Lewis, D. K. (1980). Index, context, and content. In Kanger, S. and Öhman, S., editors, *Philosophy and Grammar*, volume 143 of *Synthese Library*, pages 79–100. Springer.
- Ludlow, P. (2014). *Living Words*. Oxford University Press.
- May, R. (1985). *Logical Form: Its Structure and Derivation*. MIT Press, Cambridge, MA.
- McCready, E. S. (2010). Varieties of conventional implicature. *Semantics and Pragmatics*, 3(8):1–57.
- Michaelson, E. (2013). *This and That: On the Semantics and Pragmatics of Highly Context-Sensitive Terms*. PhD thesis, University of California, Los Angeles.
- Montague, R. (1970). Universal grammar. *Theoria*, 36:373–398.
- Montague, R. (1973). The proper treatment of quantification in ordinary english. In Hintikka, J., Moravcsik, J., and Suppes, P., editors, *Approaches to Natural Language*. Springer, Dordrecht.
- 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. (2007). Heavy hands, magic, and scene-reading traps. *European Journal of Analytic Philosophy*, 3(2):77–132.
- Neale, S. (2016). Silent reference. In Ostertag, G., editor, *Meanings and Other Things: Essays in Honor of Stephen Schiffer*, pages 229–342. Oxford University Press, Oxford.
- Newport, E. L. (1990). Maturation constraints on language learning. *Cognitive Science*, 14:11–28.
- 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. (1986). Noun phrase interpretation and type-shifting principles. In Groenendijk, J., de Jongh, D., and Stokhof, M., editors, *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*, Groningen–Amsterdam Studies in Semantics, pages 115–143. Floris, Dordrecht.
- 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.
- 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.
- Pettit, D. (2002). Why knowledge is unnecessary for understanding language. *Mind*, 111(443):519–550.
- Pettit, D. (2005). Belief and understanding: A rejoinder to gross. *Mind*, 114(453):67–74.
- Pietroski, P. (2005). *Events and Semantic Architecture*. Oxford University Press, Oxford, UK.
- Pietroski, P. (2006). Character before content. In Thomson, J. J. and Byrne, A., editors, *Content and Modality: Themes from the Philosophy of Robert Stalnaker*, chapter 3, pages 34–60. Oxford University Press, Oxford, UK.
- Pietroski, P. (2018). *Conjoining Meanings: Semantics without Truth Values*. Oxford University Press, Oxford, UK.
- Pietroski, P. M. (2008). Minimalist meaning, internalist interpretation. *Biolinguistics*, 2(4):317–341.
- Pinker, S. (1997). *How the Mind Works*. Norton, New York.
- Pinker, S. and Jackendoff, R. (2005). The faculty of language: what’s special about it? *Cognition*, 95:201–236.
- 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.
- Potts, C. (2005). *The Logic of Conventional Implicatures*. Oxford University Press.
- Prinz, J. (2006). Is the mind really modular? In Stainton, R. J., editor, *Contemporary Debates in Cognitive Science*, pages 22–36. Blackwell, Oxford, UK.
- Pylyshyn, Z. (1984). *Computation and Cognition*. MIT Press, Cambridge, MA.
- Pylyshyn, Z. (1999). Is vision continuous with cognition? the case for cognitive penetrability of vision. *Behavioral and Brain Sciences*, 22:341–423.
- Rabern, B. (2012). Against the identification of assertoric content with compositional value. *Synthese*, 189(1):75–96.
- Rabern, B. and Ball, D. (2017). Monsters and the theoretical role of context. *Philosophy and Phenomenological Research*, Online First.
- Recanati, F. (2004). *Literal Meaning*. Cambridge University Press, Cambridge, UK.
- Recanati, F. (2010). *Truth-Conditional Pragmatics*. Oxford University Press.
- Roberts, C. (2002). Demonstratives as definites. In van Deemter, 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.
- Roberts, C. (2005). Pronouns as definites. In Reimer, M. and Bezuidenhout, A., editors, *Descriptions and Beyond*, pages 503–543. Oxford University Press, Oxford, UK.
- Roberts, C. (2012). Information structure in discourse: Toward an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5:1–69.
- Rullmann, H. (2004). First and second person pronouns as bound variables. *Linguistic Inquiry*, 35(1):159–168.
- Santorio, P. (2012). Reference and monstrosity. *Philosophical Review*, 121(3):359–406.
- Sauerland, U. (2008). On the semantic markedness of phi-features. In Harbour, D., Adger, D., and Bejar, S., editors, *Phi-Theory: Phi-Features across Modules and Interfaces*, pages 57–82. Mouton de Gruyter, Oxford, UK.
- Schiffer, S. (1994). A paradox of meaning. *Noûs*, 28:279–324.
- Schiffer, S. (2003). *The Things We Mean*. Oxford University Press.
- Schlenker, P. (2003). A plea for monsters. *Linguistics and Philosophy*, 26(1):29–120.
- 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.
- 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.
- Stalnaker, R. (1973). Presuppositions. *Journal of Philosophical Logic*, 2:447–457.
- 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.
- Stojnić, U., Stone, M., and Lepore, E. (2017). Discourse and logical form: Pronouns, attention, and coherence. *Linguistics and Philosophy*, 40(5):519–547.
- Swanson, E. (2016). The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*, 45(2):121–146.
- Szabolcsi, A. (1987). Bound variables in syntax (are there any?). In Groenendijk, J., Veltman, F., and Stokhof, M., editors, *Proceedings of the Sixth Amsterdam Colloquium*, pages 331–353. Amsterdam. Institute for Logic, Language, and Information.
- Thomason, R. H. (1990). Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In Philip. R. Cohen, J. M. and Pollack, M. E., editors, *Intentions in Communication*, chapter 16, pages 325–364. MIT Press, Cambridge, Mass.
- Travis, C. (2008). *Occasion-Sensitivity: Selected Essays*. Oxford University Press.
- Unnsteinsson, E. (2016). A gricean theory of malaprops. *Mind and Language*, forthcoming.
- Vetter, P. and Newen, A. (2014). Varieties of cognitive penetration in visual perception. *Consciousness and Cognition*, 27:62–75.
- Westerståhl, D. (1984). Determiners and context sets. In van Benthem, J. and ter Meulen, A., editors, *Generalized Quantifiers in Natural Language*, pages 45–71. Foris, Dordrecht.
- Wilson, D. and Carston, R. (2007). 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.