Received: 20 March 2018 | Revised: 26 August 2018 | Accepted: 26 September 2018

DOI: 10.1111/mila.12290

SUBMITTED ARTICLE

WILEY

Semantics without semantic content

Daniel W. Harris

Department of Philosophy, Hunter College, CUNY, New York City, New York, USA

Correspondence
Daniel W. Harris, Department of Philosophy Hunter College, CUNY, New York City, NY 10065.
Email: danielwharris@gmail.com

I argue that semantics is the study of the proprietary database of a centrally inaccessible and informationally encapsulated input–output system. This system's role is to encode and decode partial and defeasible evidence of what speakers are saying. Since information about non-linguistic context is therefore outside the purview of semantic processing, a sentence's semantic value is not its content but a partial and defeasible constraint on what it can be used to say. I show how to translate this thesis into a detailed compositional-semantic theory based on the influential framework of Heim and Kratzer. This 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.

KEYWORDS

cognitive architecture, compositional semantics, constraint semantics, modularity, semantic values

1 | INTRODUCTION

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 is a constraint on what speakers can say with the expression, and not the content of what they say. My answer to the second question is that semantics is the study of a modular component of the mind. I will defend each of these claims, in part by defending the latter view and showing that it entails the former (Section 3), and then by showing how to modify a standard compositional-semantic framework to make it fit with both ideas (Section 4).

 © 2020 John Wiley & Sons Ltd

2 | 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*.

2.1 | Semantic values as contents

Content semantics is the view that an expression's semantic value is its content and that this content may vary in a way that depends on extralinguistic context. This idea is implemented in different ways by different theoretical frameworks, in part because of disagreement about what contents are. The most influential framework identifies the semantic value of an expression ϕ, indexed to a world w, assignment g, and context c, with an extension $[\![\phi]\!]^{w,g,c}$ (or, by λ-abstracting over the world parameter, with an intension $[\![\phi]\!]^{g,c}$).[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) $[\![\text{Ann smokes}]\!]^{w,g,c} = 1$ iff Ann smokes at w

(2) $[\![\text{Ann smokes}]\!]^{g,c} = \lambda w$. 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—for example, 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) $[\![\text{Ann}]\!]^{w,g,c} = \text{Ann}$

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

(5) FUNCTIONAL APPLICATION

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

The overarching idea of clauses like (1)–(5) is that the semantic value of a declarative sentence is a propositional content that can be computed by positing contents for subsentential expressions and 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 cannot have any single proposition as its semantic value, for example, because it expresses (or is used to express) different propositions on different occasions.

[1] I will treat the framework of Heim and Kratzer (1998) and von Fintel and Heim (2011) as my foil, refer to it as "the standard framework," and use its claims and notation as a starting point. There are interesting questions about how the morals of this essay apply to various alternative frameworks, but I lack space to pursue them here.

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

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 (7) is understood to have an LF of the following form (e.g., Stanley, 2000; Stanley & Szabó, 2000; Westerståhl, 1984).

(10) $[_{\text{DP}}$ The man *dom*$]$ 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' 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 (11).

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

The assignment relative to which an expression's content is fixed is itself, according to the standard framework, given by context. So although context and assignment are often treated as

[2]For example, Kaplan (1989a), King (2013), and Heim (2008). 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, 2016), and Schiffer (2003).

distinct parameters, we can follow Heim and Kratzer in assuming that each context, c, determines an assignment, g_c, thereby collapsing the two parameters:

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

What about a context fixes the operative assignment? Heim and Kratzer (1998) 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" (p. 243). Most attempts to be more specific point either to the speaker's intentions (Heim, 2008; King, 2013) or to facts about the shared propositional attitudes of the conversation's participants (Clark & Marshall, 1981; Heim, 1982; Roberts, 2005; Stalnaker, 1978).

Why divide off a category 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 he$_1$ smokes.

In the standard framework, binding is understood in terms of compositional operations that abstract over the assignment-sensitive components of an expression, turning them into argument places in a complex predicate that can compose with the binding expression. As Heim and Kratzer (1998) put it, "roughly, what is meant by 'variable binding' is any semantic operation which removes (or reduces) assignment dependency" (p. 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 one crucial issue is whether they can be bound (Stanley & Szabó, 2000). I stress this distinction because it will become important again later: One quirk of the semantic framework I will build in Section 4 is that it classifies more expressions as variables than is sometimes assumed.

2.2 | 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 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. For example, 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. If a view like this is correct, there may be no need for compositional semantics to traffic in a theoretical notion of context at all.[3]

[3]Note that constraint semantics, as I conceive it, does *not* give up the idea that propositions are the things we assert, say, or mean.

I am not the first to advocate constraint semantics. Sperber and Wilson (1995) 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" (p.175; see also Carston, 2006, p. 633). According to Pietroski (2006), the meanings of declarative sentences "constrain without determining truth/reference/satisfaction conditions" (p. 34). Bach (1987) argues that "... the semantics of an expression gives the information that a competent speaker can glean from it independently of any context of utterance" (p. 5), and sometimes refers to sentential semantic values as "propositional radicals" (Bach, 2006, p. 437). Neale (2005) argues that a sentence's semantic value should be thought of as "a blueprint for (a template, a schematic or skeletal representation of) what someone will be taken to be saying when using [the sentence] to say something" (p. 189). Schiffer (2003, 112) argues that the semantic value of a sentence partially specifies the force and content of the illocutionary act that the sentence can be used to literally perform (cf. García-Carpintero, 2006).[4]

Many of the proposals I have just mentioned have been motivated roughly as follows. There is a deep distinction between the kinds of psychological capacities studied by semantics and pragmatics. Pragmatics studies an application of our 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 identifying 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 Section 3 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. 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.[5] I will therefore demonstrate one way of tackling this problem in Section 4.

3 | SEMANTICS AND COGNITIVE ARCHITECTURE

3.1 | The modularity of semantics

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

[4]Some have argued that semantics should pair expressions with both contents and context-neutral semantic values, referring to the latter, respectively, as "characters" (Kaplan, 1989b, p. 507), "reflexive contents" (Perry, 2001), or "propositional concepts" (Stalnaker, 1978). In Section 3, I will argue that semantics should deal *only* in context-neutral semantic values.

[5]An important exception is Pietroski (2018), who has fleshed out the formal-semantic details of his foundational vision, but who departs much further from the semantic orthodoxy than I will here.

My aim in this section is to argue that we have good reason to prefer constraint semantics over content 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.[6] 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 my view, 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.

This picture gives us reason to reject Lewis's view that semantics aims to model 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, semantic competence informs the operations of the semantic module. On the other hand, our beliefs and other intentional mental states are drawn upon by pragmatic reasoning but are off limits to the semantic module. An empirically adequate account of the psychological states underlying language use must not conflate these two bodies of information.[7,8]

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?[9]

First, note that semantic composition has many of the stereotypical features of processes carried out by Fodorian modules. It is fast. It is mandatory, in that we cannot voluntarily choose not to perceive linguistic utterances as meaningful. Although semantic composition is not 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 its proprietary database is amenable to computational modeling. Unlike many central-cognitive processes, semantic composition is not folk-psychologically scrutable.

Most importantly, semantic composition is performed by a system that is, in large part, informationally isolated from central cognition. The semantic module sends outputs to central cognition as one step in language comprehension, and takes instructions from central cognition

[6]I will not take a stand on whether what Fodor labels "central cognition" is itself divided into domain-specific subsystems with distinct evolutionary origins (Carruthers, 2006; Cosmides & Tooby, 1992; Sperber & Wilson, 2002).

[7]Lewis (1980) anticipated that "a grammar that assigns one sort of semantic value could fit better into future psycholinguistics than one that assigns another sort" (p. 83), but claimed that it would be speculative to draw any such conclusion at the time. I reply that this is not a speculative project at present, and that we now have good reasons to prefer constraints over contents.

[8]Some have denied that semantics is the study of any aspect of human psychology—for example, Soames (1985), Devitt (2006), and Katz (1981). I will bracket these views for lack of space and assume that semantics is the study of a body of information that plays a causal role in language perception and production.

[9]For previous defenses of modularity about semantics, see Borg (2004, Chapter 3) and Sperber and Wilson (1995, Chapter 4). I discuss Borg's view in Section 3.2. I agree, in essence, with Sperber and Wilson, but my arguments differ from theirs in appealing to contemporary semantic theory.

as one step in language production, but there are good reasons to think that it otherwise lacks access to the inner workings and informational resources of central cognition, and vice versa.

One half of this informational isolation is that the operations and database of the semantic module are *centrally inaccessible*. Language users are not aware of the intermediate steps of semantic composition or of the contents of the database on which it draws. Theorists have at various times disagreed about nearly all of the principles underlying semantic composition, and none of these debates has been resolved by direct cognitive access to the semantic module. Instead, the inner workings and informational resources of the semantic module must be painstakingly reverse engineered. 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 (May, 1985), in-situ type-shifting (Szabolcsi, 1987), 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 cognitively 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 we routinely represent linguistic inputs in terms of concepts that are sequestered from central cognition—concepts such as SEMANTIC TYPE and NUMERICAL INDEX. I will call these "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 or 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, 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, p. 544). The subject is then presented with example sentences such as (14).

(14) Let's have mud for lunch.

The subject has false beliefs about the meaning of "mud," both before and after encountering this sentence. But Pettit argues that these beliefs do not interfere with the subject's ability to understand utterances of sentences like (14). 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 (Fodor, 1983, p. 66).[10]

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. Does this not 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 is successful semantic composition. This is distinct from understanding a speaker, which entails forming a veridical representation (perhaps a true belief) about what they said or meant by their utterance. On my view, 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. 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. Imagine a hearer who 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 my view, this process will 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—perhaps unconsciously—the hearer infers that this evidence must not be accurate and, noting the similarity of "electrical" and "electoral," concludes that Berra meant (and perhaps *said*)[11] that Texas has a lot of electoral votes. 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 they may correct 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 intuitive semantic properties of open-class vocabulary, such as the belief that the word "mud" refers to mud. These are semantic features that are rough and obvious enough to be tracked by 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—that is, 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. Pauline vehemently denies the existence of raising and believes that object-position DPs are interpreted in situ with the help of type-shifting

[10]For a recent defense of the modularity of visual perception that emphasizes similar evidence for informational encapsulation, see Firestone and Scholl (2016). Although Pettit (2002) does not present his case as evidence for modularity, he briefly makes that implication explicit and mentions the analogy to visual illusions in a follow-up article (Pettit, 2005, p. 74).

[11]See Unnsteinsson (2017) for a defense of the claim that what a speaker says with a malapropism may be their intended meaning, and need not match the word's literal meaning.

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 and processes that allow them to use language. If semantic composition were not an informationally encapsulated process—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 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.

One complication for the idea that semantics is the study of a modular system arises from the fact that this system presumably plays 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. I see no reason to think that this fact undermines the informational encapsulation of the language module in its role as an input system. There is also a clear sense in which the semantic module *qua* output system is informationally encapsulated. 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 do not recognize themselves as designing particular semantic representations in the course of producing an utterance, and they cannot, in part because some of the concepts that would be needed in order to do so are centrally inaccessible. Speakers thus are not 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 cannot choose whether to represent proper names as type-*e* terms or as type-*et* predicates, for example. Likewise, our disagreeing semanticists, Robert and Pauline, cannot 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 production as well.

A second complication arises from the fact that the semantic module needs to acquire, maintain, and update some of the information in its proprietary database. We are not born representing the semantic value of "dog." How do we get it? I will not pretend to have an answer to this question; the study of semantic acquisition is nascent compared to the study of, for example, syntactic acquisition. However, we have reason to expect that a good answer will take the form of a computational theory of the kind offered by psycholinguists working within the framework of generative grammar.[12] Theories in this tradition typically posit domain-specific learning mechanisms that are, to a significant extent, dissociated from the operations of central cognition (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, and so forth (Pinker & Jackendoff, 2005, pp. 205–206). Since these are

[12]For a recent survey of the kind of work I have in mind, see Guasti (2017).

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.[13] However, from the fact that information flows from central cognition into the semantic module via an acquisition mechanism, it does not follow that this information can affect online processing by the semantic module in its role as an input system, and my arguments in this section suggest that it does not.

3.2 | Modularity and constraint semantics

What does the modularity of semantics have to do with the content–constraint debate that I outlined in Section 2? 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. However, if semantic composition 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. Identifying the contents of semantically underspecified expressions is a process that draws on all manner of central-cognitive information. If I am right that semantic composition is a modular process and content resolution is a central-cognitive process, then semantics can deliver only constraints, not contents.

My reason for thinking that content resolution is a central-cognitive process is that understanding what someone has 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—is a process that requires information about extralinguistic context, including information about other agents' mental states. This point is well rehearsed in the literature, and so I will not belabor it here.[14] But 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 that content resolution is a central-cognitive process does not depend on any particular view of what fixes the contents of semantically underspecified expressions. Some have argued that recognizing what a speaker says with semantically underspecified expressions is a matter of inferring the speaker's communicative or referential intentions.[15] Others have emphasized the importance of integrating grammatical information with the information that is commonly known or accepted, or otherwise shared by the interlocutors.[16] Either way, content resolution is a process that relies heavily on mindreading, which is a paradigmatic central-cognitive task.[17]

[13]Bloom (2000, Chapter 3) and Hacquard and Lidz (2018) both present evidence that semantic acquisition recruits mindreading, for example.

[14]See, for example, Carston (2002), Neale (2004, 2007), Recanati (2004); Sperber and Wilson (1995), Fodor (2002), Bach (1987), and Borg (2004, 2012).

[15]Kaplan (1989a), Bach (1992), Neale (2004, 2007), King (2013), and Heim (2008).

[16]Roberts (2002), King (2013), Stalnaker (1978), Lewis (1979), Heim (1982), and Clark and Marshall (1981).

[17]The alternative is to think that content resolution relies on information from a discourse context that is fully under grammatical control (Lepore & Stone, 2015; Stojnić, Stone, & Lepore, 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

I am not the first to defend a revisionary claim about the nature of semantic values with an argument from the modularity of semantics. One previous 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 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, p. xviii)

So the semantic module identifies the content of "that" by grasping a singular concept C, and they grasp C by formulating a 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. 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

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 & Stone, 2015, Chapter 13). Drawing out the lessons for these views of my arguments in this paper would take considerable space, and so I will bracket further discussion for now.

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 hearer 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) puts it,

> [T]hat 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. (p. 135)

Still, she maintains that "as far as linguistic meaning or semantic content is concerned such issues of substantive object-identification are irrelevant."[18]

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 (Borg, 2012, pp. 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 Section 2. 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:[19]

(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 a content for any particular utterance of "this is red."

As I see it, Borg would be better off without her assumption that "every well-formed sentence, perhaps relativized to a context of utterance, is capable of expressing a proposition" (Borg, 2012, p. 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, pp. 6–7).[20] My response to this argument will come in Section 4, 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, pp. 7–8). But of course, it is

[18]Garía-Carpintero (2013) likewise argues that Borg's claim that sentential semantic values are propositions is not obviously consistent with the rest of her view.

[19]Borg (2012, p. 135) takes (16) from Higginbotham (1994, pp. 92–93). Conditional truth-condition views have also been endorsed by Higginbotham (1988), Burge (1974), Heck (2014), Gross (2005), and Larson and Segal (1995).

[20]Borg does not discuss 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.

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.

The failure to question content semantics has led others who are sensitive to considerations about mental architecture either to conclude that compositional semantics, as usually conceived, is impossible (Fodor, 2002), or to give up on the modularity of semantics and conclude that the composition of expressions' contents is 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 semantics must deal in contents.

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 view[21]—then the semantics–pragmatics interface also turns out to coincide with the interface between language perception and general-purpose cognition.

4 | CONSTRAINT SEMANTICS

My goal in this remaining section is to develop a constraint semantics by minimally adjusting the standard framework.[22] 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.

I begin with 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 best way to develop an elegant 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.[23] However, the point of my approach is to demonstrate that radical departures from semantic orthodoxy are not necessary in order to achieve a result that is consistent with the foundational lessons I have drawn.

Second, the standard way of defending claims about 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. I will be arguing that the semantic values of declarative

[21]For defenses of the encapsulation of visual perception (and perhaps perception more broadly), see Fodor (1983), Pylyshyn (1999), and Firestone and Scholl (2016). For the opposing view, see Prinz (2006), Churchland (1988), and Vetter and Newen (2014).

[22]The theory presented here also inherits some limitations in its explanatory scope from the standard framework as regards, for example, presupposition projection and discourse anaphora. There is no space for adequate engagement with these phenomena here. I hope to address them in future work.

[23]For example, alternative semantics provides us with tools that could be used to directly compose underspecified subsentential semantic values (Alonso-Ovalle, 2006; Ciardelli, Roelofsen, & Theiler, 2017; Kratzer & Shimoyama, 2002); variable-free semantics provides us with methods of delivering underspecified sentential semantic values without proliferating variables (Jacobson, 2014); and Pietroski (2018) has developed a compositional-semantic framework that eschews contents (along with other orthodoxies) for reasons that complement those I have presented here. In future work I hope to explore how tools from these alternative approaches may better realize my foundational vision.

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 considering. 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 if the result falls outside semanticists' usual methodology.

4.1 | Semantic values as characters?

I have been arguing that the semantic value of a sentence is not its context-relativized content but 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 possible answer is that semantic values are 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) $[\![\text{it}_1 \text{ stinks}]\!]^c = \lambda w \,.\, g_c(1) \text{ stinks at } w$

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

Here is a simple way of achieving this result. We leave the compositional machinery 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 need not 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 do not think that it can be correct. My reason is that this proposal requires language users' central-cognitive system(s) to be capable of working with representations that are framed in terms of centrally inaccessible concepts, such as those of 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 (for example, 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 (2008, 35–36) 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 (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. 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.[24] Semanticists have also posited many other unpronounced, semantically underspecified expressions. 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 do not 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, the strategy of reinterpreting (18) in terms of speakers' referential intentions leads us to yet another inadequate mixture of central and modular representations.

4.2 | First steps

With my positive proposal I want to capture the virtues of the foregoing account while positing sentential semantic values of a kind that both the semantic module and central cognition can

[24]For a recent overview of the evidence for null subjects and their syntactic and semantic properties, see Camacho (2013).

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[25]:

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

In English: The semantic value of "it stinks" is a property possessed by any proposition p such that, for some x, p is the proposition that x stinks. This is a constraint on what can be said literally with "it stinks." If I hear someone utter "it stinks," and if I assume that they are speaking literally but do not know anything about the context or their referential intentions on this occasion, all that my semantic module tells me about what they have said is that it has this property. To work out what they have said, I need further, extralinguistic information.[26]

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 Section 4.1 for computing sentences' characters, in that it piggybacks on the standard framework of content semantics and then removes the undesirable features of that framework's sentential semantic values at the last moment, via an abstraction operation. What is needed is an abstraction operation that takes us from (17) to (19).[27]

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

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

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

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

I will use the term "prenex quantifiers" for the existential quantifiers contained within a prenex. Generalizing from this example, what is needed is an abstraction operation that respects the following constraint: For any expression α, and for each numerical index i to which at least one free variable in α is indexed, there will be a quantifier in the prenex of $[\![\![\alpha]\!]\!]$ that binds all

[25]This semantic value abstracts away from are clause-type, tense, aspect, and the ϕ-features of "it." I will show how to build the ϕ-features back in in Section 4.3. Note that I leave the standard notation as it is, and write my preferred kind of semantic values using triple brackets.

[26]Some readers may find it more intuitive to think of the semantic value of a sentence as the set of propositions that can be literally expressed with it. The following is equivalent to (19) for example.

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

When I hear someone utter "it stinks," my semantic module tells me that what they said is a proposition in this set (if they were speaking literally). I have to pragmatically infer *which* of these propositions is the content of what they said.

[27]For now, I ignore the context parameter and indexicals. I will show how to handle indexicals in Section 4.4.

occurrences of a variable x, with one occurrence of x in $[\![[\![\alpha]\!]]\!]$ substituted for each occurrence of $g(i)$ in $[\![\alpha]\!]^g$. Here is a preliminary statement of this operation[28]:

(21) *Constraint Abstraction* (preliminary version)
If α dominates unbound variables $v_i ... v_n$ and $[\![\alpha]\!]^g \in D_\tau$
Then $[\![[\![\alpha]\!]]\!] = \lambda p_\tau.(\exists x^i)...(\exists x^n)p = [\![\alpha]\!]^{g^{i \to x^i ... n \to x^n}}$

Constraint Abstraction is a general tool for converting assignment-relativized intensions into constraints of the kind that I take to be the outputs of the semantic module. 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 are not 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 does not 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.

4.3 | The meanings of variables

The semantics I have sketched so far is on the right track but still needs work. 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 $[\![[\![\alpha]\!]]\!]$, this principle needs *syntactic* information about the variables contained in α, and not merely *semantic* information about its input semantic value, $[\![\alpha]\!]^g$. A second problem is that I have not shown how to distinguish the semantic contributions of different semantically underspecified expressions. I will address these problems in this section.

I begin with the relatively straightforward case study of "he" and "she." As I illustrated in Section 2, 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, $[\![\text{he}_i]\!]^g$ 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.[29] However, by binding variables that are free at LF, the prenex quantifiers that I have posited eliminate their assignment sensitivity, thereby nullifying

[28]Notation: Subscripts on variables are numerical indices. Superscripts on variables are merely for disambiguation.
[29]This account comes in several versions, some of which also marshall syntactic resources to explain some cases of agreement (Heim, 2008; Sauerland, 2008).

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

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

What is needed is a way to distinguish the meanings of "he" and "she" so that they 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. 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.

I will write the constraint property of a variable v as $\mu(v)$. Here are the constraint properties of "he" and "she," for example:

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

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

Information about variables' constraint properties is included in semantic competence. I therefore 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 semantics that keeps track of the constraint properties of all of the variables in an expression and a new-and-improved Constraint Abstraction principle that installs them as restrictors on the appropriate prenex quantifiers. I will accomplish this by assigning every expression a double-barreled semantic value, as follows:[30]

(27) For every node α in every LF, $[\![\alpha]\!]^{w,g} = \langle [\![\alpha]\!]_1^{w,g}, [\![\alpha]\!]_2 \rangle$

The two coordinates within an expression's semantic value track different information about its meaning. The first coordinate, $[\![\alpha]\!]_1^{w,g}$, is α's semantic value according to the standard framework—an assignment-and world-relativized extension. The second coordinate, $[\![\alpha]\!]_2$, 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, $[\![\![\alpha]\!]\!]$.

Because only variables have constraint properties, and $[\![\cdot]\!]_2$ stores information about constraint properties, only variables are assigned substantive values by $[\![\cdot]\!]_2$.

(28) For any variable v and numerical index i, $[\![v_i]\!]_2 = \{\langle i, \mu(v)\rangle\}$

(29) For any non-variable lexical item α, $[\![\alpha]\!]_2 = \emptyset$.

On the other hand, we can stipulate that $[\![\cdot]\!]_1^{w,g}$ is identical to the interpretation function of the standard framework.[31]

How do my double-barreled semantic values compose? What is needed is an account of how the values of both $[\![\cdot]\!]_1^{w,g}$ and $[\![\cdot]\!]_2$ compose. Take the latter first. The value of $[\![\cdot]\!]_2$ 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 $[\![\cdot]\!]_2$ to store the constraint properties of whatever variables the node dominates. A first stab at this could be accomplished by taking $[\![\alpha\beta]\!]_2$ to be $[\![\alpha]\!]_2 \cup [\![\beta]\!]_2$. There is just one complication: If α contains multiple coindexed variables, we want $[\![\alpha]\!]_2$ 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 $x^1 ... x^n$, $y^1 ... y^m$ and functions $X^1 ... X^n$, $Y^1 ... Y^m$, if $\chi = \{\langle x^1, X^1\rangle ... \langle x^n, X^n\rangle\}$ and $\gamma = \{\langle y^1, Y^1\rangle ... \langle y^m, Y^m\rangle\}$, 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 $\exists\tau : X^i, Y^j \in D_{\tau t}$ 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 $[\![\cdot]\!]_2$ compose at branching nodes.

[30]My 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.

[31]For any expression α, $[\![\alpha]\!]^g = \langle [\![\alpha]\!]_1^g, [\![\alpha]\!]_2 \rangle = \langle \lambda w. [\![\alpha]\!]_1^{w,g}, [\![\alpha]\!]_2 \rangle$.

Values of $[\![\cdot]\!]_1^{w,g}$ compose by exactly the same principles posited by the standard framework. This is the sense in which my account is conservative: The standard framework continues to be a true theory of $[\![\cdot]\!]_1^{w,g}$. I have merely added to this theory.

Here is a version of functional application that will work for my purposes:

(31) *Functional Application*

$[\![\alpha\beta]\!]^{w,g} = \langle [\![\alpha\beta]\!]_1^{w,g}, [\![\alpha\beta]\!]_2 \rangle$, such that:

- $[\![\alpha\beta]\!]_1^{w,g} = [\![\alpha]\!]_1^{w,g}([\![\beta]\!]_1^{w,g})$ or $[\![\beta]\!]_1^{w,g}([\![\alpha]\!]_1^{w,g})$ (whichever is defined); and
- $[\![\alpha\beta]\!]_2 = [\![\alpha]\!]_2 \uplus [\![\beta]\!]_2$

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

Lastly, here is my updated version of Constraint Abstraction, which works by combining the values of $[\![\cdot]\!]_1^{w,g}$ and $[\![\cdot]\!]_2$ at a 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 $[\![\alpha]\!]_2 = \{\langle i, f^i \rangle ... \langle j, f^n \rangle\}$ and $[\![\alpha]\!]_1^g \in D_\tau$

Then $[\![[\![\alpha]\!]]\!] = \lambda p_\tau . (\exists x^1 : f^i(x^1)) ... (\exists x^n : f^n(x^n))(p = [\![\alpha]\!]_1^{g^{i \to x^1 ... j \to x^n}})$

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. Some variables' constraint properties may be nearly vacuous, specifying nothing more than the variable's semantic type. In other cases, a variable's constraint property may place complex constraints on how it may be literally used.

4.4 | 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 directly on the context parameter rather than the assignment parameter. The theory that I have outlined so far posits no context parameter, and in keeping with the spirit of constraint semantics as laid out in Section 2, I would like to keep it that way.

My strategy will be to treat indexicals as yet more variables, much like the pronouns that I have discussed so far. Although this strategy departs from the standard practice among most philosophers, Heim and Kratzer (1998, p. 244) treat first and second-person pronouns as variables that are distinguished from third-person pronouns only by their ϕ-features.[32]

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

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

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

(iii) Only I got a question that I understood (Partee, 1989).

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:

(33) $\mu(\text{I}) = \lambda x_e \,.\, x$ is the speaker of the utterance being interpreted

(34) $\mu(\text{you}) = \lambda x_e \,.\, x$ is the addressee of the utterance being interpreted[33]

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

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

(37) $\mu(\text{here}) = \lambda x_e \,.\, x$ is a location in space within which the utterance being interpreted was produced[34]

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 themself. But as I argued in Section 3, 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 cannot identify the speaker on its own, it cannot plug in the speaker as the semantic value of an expression.

Instead, the occurrences of "the utterance being interpreted" in (33)–(37) should be understood as attributive definite descriptions in the metalanguage in which the semantic module's outputs are framed. In language perception, the semantic module presumes that there is a unique utterance that it is interpreting, and it uses "the utterance being interpreted" to attributively pick out this utterance. It hands this description upstairs to the central system(s), one of whose jobs is to identify the description's denotation.

In some cases, the semantic module's presumption of a unique utterance that it is interpreting 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 language module 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.

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

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

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:[35]

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

(39) $\mu(\text{that}) = \lambda x_e \,.\, x$ is "distal" from the perspective of the speaker of the utterance being interpreted

I place "proximal" and "distal" in scare quotes 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" cannot be understood in purely spatial ways because 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.

5 | CONCLUSIONS

There is much more to say about how the project of this essay could be expanded and refined, but I have run out of space. For now I hope to 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. If we do, the semantics–pragmatics interface turns out to be precise and grounded in underlying features of human cognitive architecture.

ACKNOWLEDGEMENTS

For helpful feedback on earlier versions, thanks to David Braun, Simon Charlow, Jonathan Cohen, Kai von Fintel, 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*.

ORCID

Daniel W. Harris https://orcid.org/0000-0002-4764-3703

REFERENCES

Alonso-Ovalle, L. (2006). *Disjunction in alternative semantics*. (PhD thesis). University of Massachusetts, Amherst.

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

[35]For arguments that demonstratives are variables, see Nowak (2019).

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.

Bloom, P. (2000). *How children learn the meanings of words*. Cambridge, MA: MIT Press.

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

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

Burge, T. (1974). Demonstrative constructions, reference, and truth. *Journal of Philosophy*, *71*(7), 205–223.

Camacho, J. A. (2013). Null subjects. In *Cambridge Studies in Linguistics* (Vol. 137). Cambridge: Cambridge University Press.

Carruthers, P. (2006). *The architecture of mind: Massive modularity and the flexibility of thought*. Oxford: Oxford University Press.

Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication*. Oxford: Blackwell.

Carston, R. (2006). Relevance theory and the saying/implying distinction. In L. R. Horn & G. Ward (Eds.), *The handbook of pragmatics* (pp. 633–656). Oxford: Blackwell.

Churchland, P. (1988). Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science*, *55*, 167–187.

Ciardelli, I., Roelofsen, F. & Theiler, N. (2017). Composing alternatives. *Linguistics and Philosophy*, *40*, 1–36.

Clark, H. H. & Marshall, C. R. (1981). Definite reference and mutual knowledge. In A. K. Joshi, B. L. Webber & I. A. Sag (Eds.), *Elements of discourse understanding* (pp. 10–63). Cambridge: Cambridge University Press.

Cooper, R. (1975). *Montague's semantic theory and transformational syntax*. (PhD thesis). University of Massachusetts, Amherst.

Cooper, R. (1983). *Quantification and syntactic theory*. Dordrecht: Reidel.

Cosmides, L. & Tooby, J. (1992). Cognitive adaptations for social exchange: Has natural selection shaped how humans reason. In J. Barkow, L. Cosmides & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 163–228). Oxford: Oxford University Press.

Cumming, S. (2008). Variabilism. *The Philosophical Review*, *117*(4), 525–554.

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

Fara, D. G. (2015). Names are predicates. *Philosophical Review*, *124*(1), 59–117.

Firestone, C. & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for 'top-down' effects. *Behavioral and Brain Sciences*, *39*, (e229), 1–77.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

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

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

García-Carpintero, M. (2013). Review of *Pursuing meaning* by Emma Borg. *Mind*, *122*(486), 515–521.

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* (2nd ed.). Cambridge, MA: MIT Press.

Hacquard, V. & Lidz, J. (2018). Children's attitude problems: Bootstrapping verb meaning from syntax and pragmatics. *Mind & Language*, *34*(1), 1–24.

Heck, R. (2014). Semantics and context-dependence: Towards a Strawsonian account. In B. Sherman & A. Burgess (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 327–364). Oxford: Oxford University Press.

Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. (PhD dissertation). University of Massachusetts. Amherst.

Heim, I. (2008). Features on bound pronouns. In D. Harbour, D. Adger & S. Bejar (Eds.), *Phi-theory: Phi-features across modules and interfaces* (pp. 35–56). Oxford: Oxford University Press.

Heim, I. & Kratzer, A. (1998). *Semantics in generative grammar*. Oxford: Blackwell.

Higginbotham, J. (1988). Context, models, and meaning: A note on the data of semantics. In R. Kempson (Ed.), *Mental representations: The interface between language and reality* (pp. 29–48). Cambridge University Press: Cambridge.

Higginbotham, J. (1994). Priorities in the philosophy of thought. *Proceedings of the Aristotelian Society*, *68*(Suppl), 85–106.

Jacobson, P. (2014). *Compositional semantics: An introduction to the syntax/semantics interface*. Oxford: Oxford University Press.

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

Kaplan, D. (1989b). Demonstratives. In J. P. Joseph Almog & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford: 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*(1), 288–311.

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

Kratzer, A. & Shimoyama, J. (2002). Indeterminate pronouns: The view from Japanese. In Y. Otsu (Ed.), *Proceedings of the third Tokyo conference on psycholinguistics* (pp. 1–25). Hituzi Syobo: Tokyo.

Larson, R. & Segal, G. (1995). *Knowledge of meaning*. Cambridge, MA: MIT Press.

Lepore, E. & Stone, M. (2015). *Imagination and convention*. Oxford: Oxford University Press.

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 S. Kanger & S. Öhman (Eds.), *Philosophy and grammar* (pp. 79–100). Dordrecht: Reidel.

May, R. (1985). *Logical form: Its structure and derivation*. Cambridge, MA: MIT Press.

Montague, R. (1973). The proper treatment of quantification in ordinary English. In J. Hintikka, J. Moravcsik & P. Suppes (Eds.), *Approaches to natural language* (pp. 221–242). Dordrecht: Reidel.

Neale, S. (2004). This, that, and the other. In A. Bezuidenhout & M. Reimer (Eds.), *Descriptions and beyond* (pp. 68–182). Oxford: Oxford University Press.

Neale, S. (2005). Pragmatism and binding. In Z. G. Szabó (Ed.), *Semantics versus pragmatics* (pp. 165–285). Oxford: 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 G. Ostertag (Ed.), *Meanings and other things: Essays in honor of Stephen Schiffer* (pp. 229–342). Oxford University Press: Oxford.

Newport, E. L. (1990). Maturational constraints on language learning. *Cognitive Science*, *14*, 11–28.

Nowak, E. (2019). No context, no content, no problem. *Mind & Language*.

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

Partee, B. (1986). Noun phrase interpretation and type-shifting principles. In J. Groenendijk, D. de Jongh & M. Stokhof (Eds.), *Studies in discourse representation theory and the theory of generalized quantifiers, Groningen–Amsterdam studies in semantics* (pp. 115–143). Dordrecht: Floris.

Partee, B. (1989). Binding implicit variables in quantified contexts: Parasession on language in context. In C. Wiltshire, B. Music & R. Graczyk (Eds.), *Papers from the 25th regional meeting, Chicago linguistics society* (pp. 342–365). Chicago, IL: Chicago Linguistics Society.

Perry, J. (2001). *Reference and reflexivity*. Stanford: CSLI Publications.

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. (2006). Character before content. In J. J. Thomson & A. Byrne (Eds.), *Content and modality: Themes from the philosophy of Robert Stalnaker* (pp. 34–60). Oxford: Oxford University Press.

Pietroski, P. (2018). *Conjoining meanings: Semantics without truth values*. Oxford: Oxford University Press.

Pinker, S. & Jackendoff, R. (2005). The faculty of language: what's special about it? *Cognition*, *95*, 201–236.

Prinz, J. (2006). Is the mind really modular? In R. J. Stainton (Ed.), *Contemporary debates in cognitive science* (pp. 22–36). Oxford: Blackwell.

Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive penetrability of vision. *Behavioral and Brain Sciences*, *22*, 341–423.

Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.

Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.

Roberts, C. (2002). Demonstratives as definites. In K. van Deemyer & R. Kibble (Eds.), *Information sharing* (pp. 1–48). Stanford, CA: CSLI.

Roberts, C. (2005). Pronouns as definites. In M. Reimer & A. Bezuidenhout (Eds.), *Descriptions and beyond* (pp. 503–543). Oxford: Oxford University Press.

Sauerland, U. (2008). On the semantic markedness of phi-features. In D. Harbour, D. Adger & S. Bejar (Eds.), *Phi-theory: Phi-features across modules and interfaces* (pp. 57–82). Oxford: Mouton de Gruyter.

Schiffer, S. (2003). *The things we mean*. Oxford: Oxford University Press.

Schoubye, A. (2016). Type-ambiguous names. *Mind*, *126*(503), 715–767.

Soames, S. (1985). Semantics and psychology. In J. J. Katz (Ed.), *The philosophy of linguistics* (pp. 204–226). Oxford: Oxford University Press.

Sperber, D. & Wilson, D. (1995). *Relevance: Communication and cognition*. Oxford: Blackwell.

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

Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and semantics* (Vol. 9, pp. 315–332). Academic Press: New York.

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

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

Stojnić, U., Stone, M. & Lepore, E. (2017). Discourse and logical form: Pronouns, attention, and coherence. *Linguistics and Philosophy*, *40*(5), 519–547.

Szabolcsi, A. (1987). Bound variables in syntax (are there any?). In J. Groenendijk, F. Veltman & M. Stokhof (Eds.), *Proceedings of the sixth Amsterdam colloquium* (pp. 331–353). Amsterdam: Institute for Logic, Language, and Information.

Unnsteinsson, E. (2017). A Gricean theory of malaprops. *Mind and Language*, *32*(4), 446–462.

Vetter, P. & Newen, A. (2014). Varieties of cognitive penetration in visual perception. *Consciousness and Cognition*, *27*, 62–75.

von Fintel, K. & Heim, I. (2011). *Intensional semantics*. (Unpublished Lecture Notes). Retrieved from http://web.mit.edu/fintel/fintel-heim-intensional.pdf

Westerståhl, D. (1984). Determiners and context sets. In J. van Benthem & A. ter Meulen (Eds.), *Generalized quantifiers in natural language* (pp. 45–71). Dordrecht: Foris.

How to cite this article: DW Harris. Semantics without semantic content. *Mind & Language*. 2022;37:304–328. https://doi.org/10.1111/mila.12290