compos mentis

Semantics is not Just Semantics: On the Status of Formal Semantic Theory

Juan Ignacio Murillo Vargas
University of Toronto

ABSTRACT
What does it take to understand a sentence? To most linguists and many philosophers of language, understanding a sentence is a two-part process. First, we must understand a sentence's semantic content, as captured by a semantic algorithm that shows how the truth-conditions of a sentence are determined by the meaning of its parts. Second, we must understand a sentence's pragmatic content, as deduced from its semantic content and how that content would be cooperative in the context of the conversation. If we know both of these, we have understood the sentence. However, this picture says little about language understanding as a cognitive activity. Our phenomenal experience of our linguistic competence does not involve "truth-conditions" or the kinds of pragmatic derivations linguistics and philosophers of language have assumed. There is a gap, here, between what traditional theorizing says, and what linguistic competence is like.

In this paper, I argue that this gap is problematic: semantics should capture, at some level of abstraction and detail, the cognitive activity involved in understanding a sentence. I also argue that this is not accomplished by assuming semantic algorithms are part of our cognitive architecture, as scholars like Chomsky have proposed. Instead, I propose that we ought to take formal semantic theories to merely be reifications of the underlying cognitive activity that language understanding involves. I conclude by outlining how developing this proposal, both philosophically and empirically, leads to the conclusion that language understanding is possible because we, as a linguistic community, choose to keep our language sufficiently uniform so that our language processing system can be epistemically reliable.

KEYWORDS
Semantics, Semantic Competence, Language Understanding, Cognitive Activity

INTRODUCTION

What is involved in understanding the following sentences?

(1) Snow is white.

(2) The llama ate the kale.

(3) Juan should go outside more often.

One prominent answer in linguistics and some philosophy of language is as follows. First, you need understand the sentence's semantic content. For declarative sentences like (1) – (3), this involves knowing the sentence's truth-conditions: what the world would need to be like for the sentence to be true (cf. Heim and Kratzer 1998). These truth-conditions are determined by the meanings of sentence's parts, and how they are put together (cf. Frege 1997a, Frege 1997b, Frege 1997c, Frege 1997d). The task of a semantic theory, per this traditional view, is to capture and generate an algorithm that describes how the parts generate the truth-conditions of the whole sentence. (For examples of this approach, see Frege 1997a, Frege 1997b, Frege 1997c, Frege 1997d, Russell 2010, Lewis 1970, Kaplan 1989, Heim & Kratzer 1998, etc.)

The traditional story then tends to appeal to something like Grice 1993. He thought that what a sentence conveys in a conversation is determined by two parts: the sentence's semantic content, as given by a semantic theory; and the sentence's pragmatic content, wherein a competent hearer derives how the sentence's semantic content fits within the context of a conversation (cf. Grice 1993, chapter 2). The rules for this derivation are Grice's maxims: the rules all rational interlocutors must follow, and assume each other to follow, for their utterances to be cooperative in the context of the conversation (again, cf. Grice 1993, chapter 2). Take the conversation in (4), for instance:

(4) A: Hey, do you know where I could find some gas?

B: There's a gas station around the corner.

The semantic content of B's utterance in (4) is, roughly, "there is a gas station in our vicinity." However, that does not by itself answer A's question. Instead, A assumes that B is trying to be cooperative, and is thus genuinely trying to answer the question. So, A deduces that B must mean something of the form "there is a gas station in our vicinity, and that gas station has gas." This is a case of what Grice calls a "conversational implicature:" the additional pragmatic content that a speaker infers from a sentence's semantic content and how that content would be cooperative in the context of a given conversation.

The traditional view on how we understand a sentence, then, is that we i) know its semantic content, in virtue of knowing how its parts determine its truth-conditions, and ii) we know its pragmatic content, in virtue of knowing how the semantic content fits into the context of a conversation. If we know both things, we have understood the sentence.

Notice, though, that this view says little about language understanding as a cognitive activity—something we, as subjects, do.[1] Phenomenally, we do not experience ourselves using a complex algorithm or Gricean derivation to understand each other. We do not take ourselves to be interested in "truth-conditions," or any other formal machinery, when we speak to each other. There is a gap, here, between what traditional theorizing about language understanding says, and what the ordinary language understanding experience is like.

This paper has two goals. First, I argue that this gap is a problem for traditional theorizing: any successful account of language understanding should appropriately capture the fact that language understanding is a cognitive activity. Second, I try to satisfy this desideratum, by arguing that traditional semantic theories are reifications of semantic competence. In §1, I discuss the first point. In §2, I discuss the second.

1. For the purposes of this paper, I will be assuming that language understanding is a cognitive activity. While the kinds of semantics I am discussing come from a legacy of behaviorism—Quine was among those who popularized the Tarski-inspired semantics we use today—it would be problematic if being a semanticist required being a behaviorist.

1: SEMANTICS SHOULD CAPTURE COGNITION

This section argues for the following desideratum:

(5) Semantics-And-Cognition: a good semantic theory appropriately captures the fact that semantic competence is a cognitive activity, at some level of abstraction and detail.

Where "semantic competence" is defined as:

(6) Semantic competence $=_{df}$ an ordinary speaker's ability to understand a language. This involves being able to understand a sentence's semantic content, and being able to apprehend its pragmatic content.[2]

The argument for this desideratum proceeds as follows:

1. A semantic theory is an algorithm that calculates a sentence's truth-conditions based on the meanings of its parts [definition from the introduction].
2. A good semantic theory either appropriately captures that semantic competence is a cognitive activity, or it does not [tautology].
3. A good semantic theory does not appropriately capture that semantic competence is a cognitive activity [assumption for *reductio*].
4. A good theory captures what a phenomenon is like at some (discipline-specific) level of abstraction and detail.
5. Either semantic competence is not a cognitive activity, or a good semantic theory does not concern semantic competence at some level of abstraction and detail [From 3, 4].
6. Semantic competence is a cognitive activity [assumption].

2. There is some debate as to whether there is such a distinction between semantics and pragmatics, cf. Récanati 2010. Since this paper is partially focused on what linguists say language understanding involves, and since contextualism is not a live option for them, I will proceed as if there is a distinction between semantic content and pragmatic content.

7. So, a good semantic theory does not concern semantic competence at some level of abstraction [from 5, 6].
8. If a good semantic theory does not concern semantic competence at some level of abstraction and detail, it is unclear what it tells us about language.
9. It is unclear what a good semantic theory tells us about language [from 7, 8].
10. The result in 9 is absurd and must be rejected.
11. Both disjuncts in 5 are false [from 6, 10].
12. The assumption at 3 must be false [from 11].
13. THEREFORE: a good semantic theory appropriately captures that semantic competence is a cognitive activity, at some level of abstraction and detail [definition of Semantics-And-Cognition; from 2, 12].

The key premises at play here are 4, 8, and 10. Let us defend each in turn.

The claim that "a good theory captures what a phenomenon is like at some level of abstraction and detail" should be uncontroversial for most, but let us cover our bases. For one thing, this is clearly our intuition. I suspect most people would rebuke a physics that could not handle how ordinary, marcroscopic objects move. Furthermore, the whole purpose of developing a theory, we might think, is to understand what something is like and how it works to some capacity. If so, a theory that does not capture a phenomenon at a certain standard of abstraction and detail fails to be a theory at all—it runs against the very point of theory-crafting.

What about "if a good semantic theory does not concern semantic competence at some level of abstraction and detail, it is unclear what it tells us about language?" Well, if a good semantic theory does not concern semantic competence, then it either concerns some other dimension of language (say, phonological or syntactic competence) or it does not concern language at all. The former disjunct is clearly absurd: if semantics captures phonological or syntactic competence, it is entirely redundant. There are other fields, with other formal models, that study those things. So, lest we want to say that semantics is redundant, we would have to say that it is unclear what it tells us about language.

Finally, why should we think that semantics not concerning language is an absurd result? For one thing, this would require semanticists to revise what we take

ourselves to be doing, since we take ourselves to be studying natural language (e.g., Heim and Kratzer 1998, 2). Furthermore, and again, if semantics did not concern natural language, the discipline would be at high risk of being redundant, since it is simply unclear what else semantics could concern. Unless someone has an account for what else semantics could be about, we should think it absurd that it is not about natural language.

Since we have defended premises 4, 8, and 10, we can safely conclude—at least for now—that Semantics-And-Cognition is a desideratum for any good semantic theory. However, notice that it leaves indeterminate the level of abstraction and detail that a semantic theory can or should operate in to be consistent with Semantics-And-Cognition. I take this issue up in the next section.

2: SEMANTICS AS A REIFICATION

2.1 Why Not Hard Internalism?

If we are coming from a linguistics background, we might think that there is already a view of semantics consistent with Semantics-And-Cognition: Chomsky 2000 proposes a view that I will be calling "hard internalism."[3] He thinks that, insofar as we want to study language through naturalistic inquiry (e.g., the kind of inquiry in physics, chemistry, neuroscience, etc.), we should only care about "I-languages;" an individual speaker's internal, intensional associations between an expression in the lexicon, with phonological and semantic features, and how it fits into the syntax (Chomsky 2000, 26). Why does he think this? Because most of the other terms we might use to describe a public language—e.g., "communicative intent," "mutually held beliefs," etc.—are folk concepts, not natural kinds, and so they simply cannot be part of a naturalistic inquiry (cf. Chomsky 2000, chapter 3).

Of course, Chomsky owes us a full account of what "naturalistic inquiry" is, or why we might want to study natural language under its lens. Nonetheless, we might be tempted to accept his view because it easily satisfies Semantics-And-Cognition. If what semantics captures are an individual brain's intensional associations, then semantics just spells out the cognitive architecture associated with semantic competence.

Such a radically internalist view of semantics, however, infamously leads Chomsky to conclude that semantics does not concern "reference" or "extension:"

3. Pietroski 2018 is a more recent example of this type of view.

> Contemporary philosophy of language follows a different course. It asks to what a word refers, giving various answers. But the question has no clear meaning. The example of "book" is typical. It makes little sense to ask to what thing the expression "Tolstoy's War and Peace" refers, when Peter and John take identical copies out of the library. The answer depends on how the semantic features are used when we think and talk, one way or another. In general, a word, even of the simplest kind, does not pick out an entity of the world, or of our "belief space." Conventional assumptions about these matters seem to me very dubious. (Chomsky 2000, 17)

The argument looks weak, but it follows from the rest of Chomsky's framework. If we are to be studying natural language in general—and semantics in particular—under a naturalistic lens, then "reference," which is not a natural kind, and is not solely part of an individual speaker's intensional associations, simply cannot be the object of study. "Extension," for the same reason, falls by the wayside too.

I find Chomsky's view implausible for three reasons. For one, there might be a layer of normativity fundamental to language (e.g., rational patterns of use, norms of communication, etc.), and a naturalistic lens—at least as Chomsky seems to understand it—would miss out on that entirely. Moreover, there is something prima facie implausible about adding highly complex semantic algorithms to our cognitive architecture, especially considering they are not part of our phenomenal experience as language users. Finally, most semantic theory appeals to "extension," if not "reference" (cf. Heim and Kratzer 1998); so losing out on those notions, just to satisfy Semantics-And-Cognition, is simply not desirable.

What we should want, I think, is a view of semantics that keeps "extension" and "reference," but satisfies Semantics-And-Cognition without adding semantic algorithms to our cognitive architecture. (We should also leave some room for normativity in language use while we are at it.) In the next section, I will be proposing and developing my preferred way of accomplishing this: by treating formal semantic theories as a reification of the cognitive aspect of semantic competence.

2.2 Soft Internalism

Since I have lost all pretense of humility by this point, let us call my view "soft internalism:"

> (7) Soft internalism $=_{df}$ the view that semantic theories are just *reifications* of the cognitive activity involved in semantic competence.

This proposal, again, straightforwardly satisfies Semantics-And-Cognition. If semantic theories are just reifications of the cognitive activity involved in semantic competence, then, by definition, they capture what semantic competence involves cognitively soft internalism and hard internalism, then, is two-fold.One, soft internalism does not propose nor entail that the algorithms semanticists study are literally part of human cognitive architecture. Two, and consequently, soft internalism takes semantics to be operating at a higher level of abstraction than what hard internalists, like Chomsky, typically suppose.

Soft internalism comes with a few upsides. For one thing, it is more parsimonious: unlike Chomsky, I think we can and should keep the explanatory power of semantic theory without adding it to our metaphysics of the mind. Moreover, it also paves the way for normativity, extension, and reference to play a fundamental role in semantic competence, precisely because soft internalism views formal semantic algorithms as mere reifications for whatever is metaphysically at bottom in semantic competence.

For this proposal to work, however, we need to give a derivation from the cognitive activity involved in semantic competence, to what semantic theories actually look like. Otherwise, we would have no account of how it is that semantics, as a reification, can capture the cognititive aspect of semantic competence. If so, we might be tempted to accept hard internalism instead, since the story would be much more straightforward: semantic algorithms capture the cognition involved in semantic competence because they are what is at bottom in semantic competence.

One derivation to solve this is challenge is as follows:

> 1. Semanticists present native speakers with written utterances of a language, and develop formal models that i) capture our intuitions on what that utterance means, and ii) predict the intuitive meaning of similarly-structured utterances.

2. Native speakers process written and natural language incrementally: as an utterance occurs, their brains make forecasts about what the rest of the structure will look like.
3. Our intuitions about what an utterance means are generated by what our forecasting system delivers [from 1, 2].
4. Linguistic structure is relatively uniform.
5. If linguistic structure is relatively uniform, then the forecasting process will work the same way in both the written language case and the natural language case.
6. The forecasting process will work the same way in both the written language case and the spoken language case [from 4, 5].
7. The intuitions that semanticists model match the spoken language case [from 3, 6].
8. The spoken language case involves, among other things, an ordinary speaker's semantic competence.
9. THEREFORE: The intuitions that semanticists model match, among other things, an ordinary speaker's semantic competence [from 7,8].

Let us defend premise 1, 5, and 8 first, since they are easier. 1 is a basic fact about how semantics research is conducted (e.g., Heim and Kratzer 1998, Copley 2006, Yalcin 2007, von Fintel and Gillies 2010, Silk 2018, etc.). Premise 5 is a relatively plausible empirical claim: if linguistic structure is uniform, there would be no prima facie reason for the human processing system to differentiate between written and spoken language—the underlying structure would be the same. (There is some empirical backing for this idea; see Kuperberg and Jaeger 2016 for an overview.) Finally, premise 8 is straightforward; when we communicate with each other using natural language, that involves us producing sentences others can understand and understanding the sentences others produce. The second clause of that sentence is simply the definition for "semantic competence" I gave above.

This leaves us with premise 2 and premise 4 to defend. Premise 2 is an important empirical claim, which is often called "incrementality" in the psycholinguistics literature:

(8) INCREMENTALITY $=_{df}$ the way the brain processes language is not by "waiting" for the utterance to finish. Instead, as we hear parts of the utterance, the brain generates and processes forecasts[4] as to what the rest of the structure will look like.

Incrementality has been found across many different experiments, for both written and spoken language, and across different experimental paradigms (Rubio-Fernandez and Jara-Ettinger 2020, 1; again, for an instructive overview, see Kuperberg and Jaeger 2016). An intuitive motivation for this claim, however, comes from so-called "garden path" sentences. Take (9):

(9) The old man the boat.

At first glance, (9) might seem ungrammatical. It is not though: "the old" is the subject, "man" is the verb, and "the boat" is the object. The reason it appears ungrammatical is because "the old man" tends to occur as a single noun phrase. So, our brain forecasts the structure in (9) to treat it as such. When this forecast turns out to be wrong, we initially process the sentence to be ungrammatical.

Now, only premise 4 remains to be justified. At first, we might be tempted to say "it is an empirical claim" and leave it at that. However, a lot of research (especially in the Generativist tradition) implicitly presupposes that language is uniform. And relying on the empirical work, to support the presupposition underlying the empirical work, would create a nasty vicious circle. Instead, we need to give a philosophical argument here. While it will take us on an odyssey of sorts, it will also let us say something interesting about the epistemology of language understanding in the process. So, off we go.

2.3 Why Language is Relatively Uniform

To give the argument for the claim "linguistic structure is relatively uniform," let us start with the following first principles:

(10) ARBITRARINESS $=_{df}$ there is no intrinsic relationship between a natural language expression and the content it conveys.

4. The word "forecast" is used intentionally: there is an interesting question as to whether these cognitive processes are representational, perceptual, or neither. I will not take that question up here.

Instead, the relationship is *arbitrary*: members of a linguistic community choose what the association is.

(11) Structural-Rationality $=_{df}$ humans are, or at least aim to be, structurally rational. We do not want our beliefs to be inconsistent, and we do not want our actions to be inconsistent with our beliefs, needs, or goals.

(12) Communicative-Need $=_{df}$ the mind has a basic need to communicate.

Arbitrariness dates back to Ferdinand de Saussure (Saussure 2003, 79. Holdcroft 1991, 52 explicitly attributes the idea to Saussure; and Chomsky 2000, 9 also mentions it in passing, though he does not cite Saussure). In some sense, though, it just has to be right. If it were false—that is, if there were some intrinsic relationship between expressions in natural language and the content they convey—we would be unable to explain how different languages could have different words to convey the same content. Considering this is obviously the case (e.g., "dog" and "perro" refer to the same kind of cute canine), Arbitrariness must be true.

Of course, these arbitrary associations are not the result of some communal deliberative process. There is no "original position" where English speakers decided "dog" would refer to a type of cute canine. Instead, our collective choices on the relationship between expressions and their content are much more subtle.[5] For instance, we often use expressions deferentially or parasitically: we simply use them to mean whatever other speakers, usually experts, take the expression to mean (e.g., non-physicists like myself use "quantum physics" to refer to whatever it is that physicists understand by "quantum physics.")

Structural-Rationality, meanwhile, is standardly accepted amongst philosophers of language (e.g., Williams 2020, 14). Moreover, insofar as semantics is also concerned with rational patterns of inference (as Kaplan 1989 seems to think, cf. chapter 17 and 18), something like Structural-Rationality must be right. For argument's sake, I am also supposing that, since this kind rationality is minimal

5. I will not try to develop a full list or account of these choices here. In general, though, I take the choices to be more akin to the "invisible hand" economics sometimes appeals to, rather than any kind of deliberate, simultaneous choice-making activity.

and merely formal, most individuals and collectives are structurally rational in the way Structural-Rationality prescribes. While most readers might (rightfully) be skeptical, remember: Structural-Rationality makes no prescriptions on the content of the relevant beliefs and wants. So, most subjects' beliefs or wants—no matter how prima facie absurd they are—can easily meet it.

Communicative-Need is decidedly trickier to defend, and I simply do not have the space to do so here. So, let us interpret it as an empirical claim: as part of the evolutionary process, humans developed a need to communicate with each other to increase their chances of survival. This claim still needs empirical substantiation, but it is plausible enough for now.[6]

With these three principles at hand, all we are missing is Incrementality, which we already saw earlier in §2.2. Now, we can build the argument for the claim that linguistic structure is relatively uniform:

1. Linguistic structure is either relatively uniform or not [tautology].
2. Linguistic structure is not relatively uniform [assumption for *reductio*].
3. If linguistic structure is not relatively uniform, we would not know when or whether we understood someone's utterance, since our forecasting system would often be wrong [presupposes Incrementality].
4. We do not know when or whether we understood someone's utterance, since our forecasting system is often wrong [from 2, 3].
5. Our decisions surrounding linguistic structure will be consistent with our need to communicate [from Arbitrariness, Structural-Rationality, and Communicative-Need].
6. The need to communicate entails that we need to know when or whether we understand each other.
7. Our decisions surrounding linguistic structure will be consistent with our need to know when or whether we understand each other [from 5, 6].

6. This is not to say there is no philosophical justification to go along with it; I think there is and should be.

8. 4 contradicts 7, since we would not make a choice that would prevent us from knowing when or whether we understand each other [from 4, 7].
9. Since 4 follows from 2 and 3, and 3 is true, the assumption at 2 is false [from 8].
10. THEREFORE: Linguistic structure is relatively uniform [from 1, 9].

For this argument to work, we need to justify premises 3 and 6. The latter is straightforward: if we did fail to understand each other, it is unclear how we would count as "communicating" in the first place. The more philosophically interesting premise, however, is 3. Notice that it is a linguistic version of what David Hume said when he first posed the problem of induction. ("Induction," for Hume, is an inference about the future based on present evidence; see Hume 2011, 78.)) For induction to be justified, nature needs to be uniform (or, at least, relatively uniform), since only then could experiences in the present reliably tell us about the future (Hume 2011, 79). Except, the same constraint applies to our brain's forecasting system: if language were not relatively uniform, then the forecasting system would constantly misfire, as it did when it first encountered the garden-path sentence in (9).

However, language, unlike nature, is entirely up to us. Arbitrariness entails that we, as a linguistic community, choose what linguistic structure looks like. And we choose to keep that structure relatively uniform because we have a need to communicate, and that need requires us to keep things such that our semantic competence is epistemically reliable. This is a deeply interesting result for the epistemology of language understanding: we reliably understand each other because we choose to make linguistic structure relatively uniform, such that reliable understanding is possible.

Moreover, we have also given an argument for the claim that "language is relatively uniform," which was the remaining premise to justify in §2.2. So, we have also completed our derivation for how semantics, seen as a reification, can capture the cognitive activity involved in semantic competence.

3: SUMMARY

This paper asked two questions: 1) should semantic theories relate to the cognitive activity involved in semantic competence, and 2) if so, how can and should they do so?

§1 answered 1) in the affirmative: if semantics fails to capture what semantic competence involves cognitively, its explanatory role in a broader account of language is unclear.

§2 answered 2) in three parts. §2.1 argued against a Chomskyan solution, since his view is implausible for a variety of reasons. §2.2. put forth my own proposal: semantics is best seen as a reification of the cognitive activity involved in semantic competence. I also gave and defended a derivation for how semantics captures this cognitive activity: if language is relatively uniform, and if written and spoken language are subject to the same forecasting system, then the intuitions semanticists model will match what ordinary speakers understand in virtue of their semantic competence. Finally, §2.3. defended the claim that language is relatively uniform. As language users, we choose to keep linguistic structure uniform so we can meet our cognitive need to communicate.

I think what I have said here raises at least three further questions. One, we need to fully defend the claim that the mind has a need to communicate, either empirically or philosophically (ideally both). Two, we can ask further questions about the status of the other theories in linguistics, i.e., what phonology, syntax, and morphology are. Finally, if soft internalism is correct, we should also find out what cognitive realities underlie our essential semantic tools (e.g., "propositions" and "possible worlds," in the way Stalnaker 1970 and Stalnaker 1999 use them). If they are merely reifications for underlying cognitive activity, we should be able to find what these tools serve as reifications for.

REFERENCES

Almog, Joseph, John Perry, Howard K. Wettstein, and David Kaplan, eds. 1989. *Themes from Kaplan*. New York: Oxford University Press.

Chomsky, Noam, and Neil Smith. 2000. *New Horizons in the Study of Language and Mind*. 1st ed. Cambridge University Press. https://doi.org/10.1017/CBO9780511811937.

Copley, Bridget. 2006. "What Should '"should"' Mean?" *Language Under Uncertainty Workshop, Kyoto University*, 22.

de Saussure, Ferdinand. 2003. "The Linguistic Sign." In *Classic Readings in Semiotics: For Introductory Courses*, 69-87. Ottawa: LEGAS.

von Fintel, Kai, and Anthony S. Gillies. 2010. "Must... Stay... Strong!" *Natural Language Semantics* 18 (4): 351–83.

Frege, Gottlob. 1997a. "Begriffsschift: A Formula Language of Pure Thought Modelled on That of Arithmetic." In *The Frege Reader*, 47–78. Blackwell Publishing.

Frege, Gottlob. 1997b. "Function and Concept." In *The Frege Reader*, 130–48. Blackwell Publishing.

Frege, Gottlob. 1997c. "On Concept and Object." In *The Frege Reader*, edited by Michael Beaney, 181–93. Blackwell Publishing.

Frege, Gottlob. 1997d. "On Sinn and Bedeutung." In *The Frege Reader*, edited by Michael Beaney, 151–71. Blackwell Publishing.

Grice, Paul. 1993. *Studies in the Way of Words*. 3. print. Cambridge, Mass.: Harvard Univ. Press.

Heim, Irene, and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Blackwell Textbooks in Linguistics 13. Malden, MA: Blackwell.

Holdcroft, David. 1991. *Saussure: Signs, System and Arbitrariness*. 1st ed. Cambridge University Press. https://doi.org/10.1017/CBO9780511624599.

Hume, David, and Lorne Falkenstein. 2011. *An Enquiry Concerning Human Understanding*. Broadview Editions. Peterborough, Ont. ; Buffalo, NY: Broadview Press.

Kuperberg, Gina R., and T. Florian Jaeger. 2016. "What Do We Mean by Prediction in Language Comprehension?" *Language, Cognition and Neuroscience* 31 (1): 32–59. https://doi.org/10.1080/23273798.2015.1102299.

Lewis, David. 1970. "General Semantics." *Synthese* 22 (1–2): 18–67. https://doi.org/10.1007/BF00413598.

Pietroski, Paul M. 2018. *Conjoining Meanings: Semantics without Truth Values*. First edition. Context and Content. Oxford, United Kingdom: Oxford University Press.

Récanati, François. 2010. *Truth-Conditional Pragmatics*. Oxford ; New York: Clarendon Press.

Rubio-Fernandez, Paula, and Julian Jara-Ettinger. 2020. "Incrementality and Efficiency Shape Pragmatics across Languages." *Proceedings of the National Academy of Sciences* 117 (24): 13399–404. https://doi.org/10.1073/pnas.1922067117.

Russell, Bertrand, and David F Pears. 2010. *The Philosophy of Logical Atomism*. London: Routledge.

Stalnaker, Robert. 1970. "Pragmatics." *Synthese* 22 (1/2): 272–89.

Stalnaker, Robert. 1999. "Assertion." In *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford Cognitive Science Series. Oxford ; New York: Oxford University Press.

Williams, J. Robert G. 2020. *The Metaphysics of Representation*. 1st ed. Oxford University Press. https://doi.org/10.1093/oso/9780198850205.001.0001.

Yalcin, Seth. 2007. "Epistemic Modals." *Mind* 116 (464): 983–1026. https://doi.org/10.1093/mind/fzm983.