

Meaning without Information: Comments on Paul Pietroski's *Conjoining Meanings**

Paolo Santorio

Conjoining Meanings puts forward a bold conception of meaning in natural language, which is at odds with standard views in the field. And it does so with sophisticated and forceful arguments. Philosophers of language have much to learn from it, whether or not they accept Pietroski's positions.

The central claims of the book can be condensed into two theses. First, Pietroski defends a procedural conception of meaning: meanings are instructions for accessing and assembling mental representations. Second, composition of meanings in natural language happens via two conjunction-type operations. This is a major departure from all semantic theories inspired by Frege, on which linguistic items combine via saturation of functions. So, while this second thesis is technical, it touches on crucial parts of the architecture of a theory of meaning.

After giving an overview of the project, I will discuss three themes emerging from the book: the fineness of grain of meanings, polysemy, and composition rules.

1 Overview of the project

The theoretical background for the project is a picture of language highly inspired by Chomsky (1995, 2000). Language is a mental faculty with biological underpinnings, which involves the formation and manipulation of mental representations. Pietroski introduces the notion of a 'Slang' to denote "spoken or signed languages that human children can naturally acquire". Meanings are whatever Slangs connect with phonological strings that are realized in speech.

Pietroski starts from a negative claim: truth-conditional semantics does not yield a notion of meaning that is adequate to this task. By 'truth-conditional semantics' Pietroski has in the mind the tradition of formal theories of meaning that starts with Frege, and includes work by Church, Montague, Davidson, Lewis, Stalnaker, and Partee—just to mention a few. The central idea of

*For useful discussion of related topics, thanks to Fabrizio Cariani, John Collins, Josh Dever, Malte Willer, and the attendees of the 2019 Philosophy of Language conference in Dubrovnik.

truth-conditional (henceforth, TC) semantics so understood is that meanings are extensions, i.e. real-world entities (or set-theoretical constructions thereof) that Slang expressions denote. The notion of extension here is understood to encompass also what are usually called ‘intensional entities’, such as possible worlds. In particular, Pietroski takes the TC conception to be closely wedded to Davidson’s (1967a) claim that a theory of meaning should take the form of a Tarski-style theory of truth.

Pietroski’s objections to TC theories fall into four categories. (i) Extensions are not sufficiently fine-grained. (ii) The notion of truth leads to paradox. (iii) TC semantics cannot account for polysemy. (iv) TC semantics overgenerates semantic types, in the sense that it leads us to expect expressions with more complex semantic types than what we observe. I will discuss some of these points in detail below.

So what are meanings, if not truth conditions? Pietroski advocates a ‘procedural’ notion of meaning, on which meanings are literally instructions. Assume in the background a Fodorian picture of concepts: concepts are atomic mental representations that (i) apply to objects, (ii) are intentional (i.e. portray objects as being a certain way), and (iii) compose to form more complex representations. For Pietroski, the meaning of a lexical item is an instruction to fetch concepts. In particular, each lexical item is paired with a long-term mental address; this address is paired with one or more concepts. Meanings are instructions to fetch concepts from the relevant addresses.

A simple example: the meaning of *book* is an instruction *fetch@‘book’* to go fetch a concept from the address of *book*. Suppose, for now, that there is only one concept at this address—say, the concept *BOOK*. Then the meaning of *book* is an instruction to go fetch *BOOK*.

On this picture, meanings compose in the strong sense that meanings of composing expressions are literally parts of meanings of complex expressions. Since meanings are instructions, this entails that the instructions invoked by composing expressions are part of the instructions invoked by more complex expressions. For an illustration, take (1):

- (1) heavy book

heavy and *book* compose by means of an operation called ‘M-junction’. M-junction is an instruction to take two concepts Φ and Ψ and form their conjunction $\Phi \wedge \Psi$. As a result, the meaning of (1) consists in an instruction to fetch a concept from the address linked to ‘heavy’, fetch a concept from the address linked to ‘book’, and to form their conjunction:

- (2) **M-join**(*fetch@‘heavy’*, *fetch@‘book’*)

The output of this process is the conjunctive concept $\text{HEAVY}(_) \wedge \text{BOOK}(_)$.

In the rest of this note, I develop in detail three themes: the fineness of grain of meanings, polysemy, and composition rules. As a preview: I argue that some of Pietroski’s objections to TC semantics come from an overly restrictive view

of TC frameworks. At the same time, other challenges are very real and stand open. Overall, I agree with Pietroski that, in order to solve them, semantics should move in a psychologistic direction. But I don't see in-principle reasons why this should be done by entirely abandoning standard frameworks, rather than enriching them.

2 Fineness of grain

According to Pietroski, TC theories don't provide a fine-grained enough notion of meaning. Some of his problem cases are classics: they include sentences involving necessarily empty nouns, like *unicorn*, and mathematical sentences. I won't get into details about these. But I want to discuss another class of cases, which relate to Pietroski's work on event semantics (Pietroski 2004).

Suppose there are two chipmunks, Al and Theo, who are chasing each other. It appears that these two claims can be true together:

- (3) a. Al chased Theo gleefully and athletically, but not skillfully.
- b. Theo chased Al gleelessly and unathletically, but skillfully. (Pietroski 2018, p. 196)

On standard treatments of adverbial modification (Davidson 1967b), (3a) and (3b) involve existential quantification over events that satisfy the noun *chase*; the adverbs specify properties of the relevant events. So the LF of (3a) is:

- (4) $\exists e [Agent(Al, e) \wedge Chase(e) \wedge Patient(Theo, e) \wedge Gleeful(e) \wedge Athletic(e) \wedge \neg Skillful(e)]$.

Now, Al and Theo participated in one and the same chasing event, at least if this event is individuated in terms of its spatio-temporal coordinates and its participants. But there cannot be a single chase event that has all the properties described by (3a) and (3b), since a single chase cannot be both gleeful and gleeless, athletic and unathletic, skillful and unskillful. It follows that (3a) and (3b) cannot quantify over 'real world' events.

One obvious solution is to say that events should be individuated more finely. In particular, in the situation described there are two chases rather than one; Al is the agent of one, Theo the agent of the other. This despite the fact that the two events seem to carve out the very same chunk of space-time. According to Pietroski, this solution solves the technical problem, but runs against the spirit of TC theories.

The requisite values of the variable turn out to be individuated about as finely as their linguistic *descriptions*. And it is bad news for a *truth*-theoretic conception of meaning if defending it requires stipulations coupled with appeal to domain entities that are individuated in terms of meanings ... Davidson and others tried to defend a conception of events as relatively "coarse grained" entities

that are independent of how we think and talk about them. The events were supposed to *be there*, along with the chipmunks, as things speakers can quantify over and describe in endlessly many ways. (p. 200, emphasis in the original)

It's useful to distinguish two challenges here. The first starts from a metaphysical constraint on semantics. On this challenge, the project of TC semantics is to characterize meanings by using objects and properties whose existence is mind-independent, and which are individuated by purely extensional criteria. Introducing metaphysical distinctions to accommodate the semantics of adverbs gives up on the enterprise. The second challenge is more general. It grants that the entities we appeal to might not be mind-independent, but insists that we need to individuate them independently of our need to make semantic predictions.

I focus mainly on the first challenge, but I will say something about the second as well.

This challenge rests on an austere conception of TC semantics: semantics may only help itself to entities whose existence and individuation conditions are independently motivated on metaphysical grounds. This may indeed have been a constraint that motivated early TC theorists. But it's a constraint that many contemporary theorists reject. On an increasingly popular view, meaning is identified with *information*, where information is conceived in terms of alternative possibilities (see e.g. Dretske 1981).¹ Crucially, the relevant possibilities are not fixed antecedently by the metaphysics. Rather, possibilities are just determined on the basis of the needs of our semantic and psychological theories.

The informational conception has been extraordinarily productive over the last few decades, and has generated a number of theories that depart from classical notions of extension and intension. Here are a few examples. Expressivists about normative language, like Gibbard (1990, 2003) take possibilities to be pairs of a possible world and a normative standard or a plan. Expressivists about modal discourse like Yalcin (2007; 2011) and Charlow (2015) take possibilities to be information states of various sorts, or pairs of a world and an information state. In a different strand of literature, Chalmers (2006) and others endorse a notion of epistemic possibility to model differences in cognitive significance. More recently, inquisitive semanticists (see e.g. Ciardelli

¹The informational conception can be traced back at least to Stalnaker (1978). Here is a representative quote:

[O]ne thing that is common to [inquiries, deliberations, and conversations], and essential to them, is that the participants do seek to distinguish among alternative ways that things might be, or might have been ... [T]he best way to bring out the formal structure of such activities is to focus on what is done with a given set of alternative states of the world, *setting aside questions about the nature of the alternatives themselves*. (emphasis mine, p. 316)

et al. 2018) take possibilities to be sets of information states.

All these theories differ in various respects, but share two features. (i) They appeal to notions of possibilities and information that are more fine-grained than standard intensions and extensions, with the goal of gaining modeling power. Oftentimes, the relevant notion of a possibility involves reference to entities that figure in theories of attitudes (e.g. plans, information states). (ii) Even though they are more complex than standard intensional semantics, they are generalizations of it, rather than alternatives.

So the first challenge can be met simply by dropping the metaphysical constraint that gave rise to it. TC semantics does *not* need to map linguistic items to mind-independent, extensionally individuated entities. Going back to the specific problem raised by (3a) and (3b), we can help ourselves to an appropriately fine-grained notion of an event. This might involve a departure from some early versions of TC semantics. But many contemporary frameworks in the broad TC tradition resolutely make this departure anyway.

This leaves on the table the second challenge. Grant that semantics may help itself to a fine-grained notion of an event. One might still worry that, if the relevant notion is retro-engineered to get appropriate predictions, the semantics we obtain will be descriptively adequate, but unexplanatory. I agree that this is a worry. The response is to use a notion of event that we have independent grip on; only, this independent grip may be provided not by the metaphysics, but by a general theory of how events are represented in human cognition. Again, an analogy with informational-style semantics is helpful. Gibbard's plan-based semantics for normative language (2003) appeals to the notion of plan. This notion is supposed to play an important role in a general explanation of behavior and action, independently of subjects' linguistic practices. Similarly, we should expect that the notion of event that is useful for semantic theorizing will surface elsewhere in a theory of cognition. Of course, it is an open question whether this expectation is met. But I don't see in-principle reasons why this project may not be pursued within an informational framework descending from TC semantics.

3 Polysemy

On a first pass, polysemy is the phenomenon whereby the same lexical item is used to access concepts that are different, but related. A classical example is *book*, which can be used to access a concept that applies to physical objects (call this 'BOOK:PHYSICAL'), as in (5a), or a concept that applies to the information contained in the book (call this 'BOOK:INFO'), as in (5b).

- (5) a. The book was ruined by the rain.
- b. The book was so interesting that it kept me awake at night.

Polysemy affects a large amount of vocabulary from a variety of syntactic categories: some of Pietroski's examples include *rabbit*, *green*, *star*, *tofu*, *Venice*, and *water*.

The classical strategy to tackle polysemy in TC theories has been to simply assimilate it to homophony. In the same way as the string *bank* is linked to two meanings, *book* realizes two homophonous strings, which connect to different entries in the lexicon. But this idea is a non-starter, for three reasons. First, differently from homophony, polysemy is cross-linguistically robust. So it's not a simple coincidence that *book* is connected to a set of related concepts. Second, there are cases of *copredication*, i.e. cases where a single occurrence of a polysemous word is linked to more than one concept. For example:

- (6) The book was interesting but it was ruined by the rain and I couldn't finish it.

Third, there is evidence from language processing that polysemous items are processed differently from homophonous ones (see e.g. Eddington and Tokowicz 2015).

Here I fully agree with Pietroski: polysemy is a serious and urgent problem for standard semantic frameworks. Ultimately, the problem might require enriching our picture of lexical meanings, and perhaps redesigning composition rules.² Let me explain.

Also in this case, we can distinguish two challenges—an easy and a hard one. The easy challenge is that items like *book* have an extension that varies across contexts. This can be handled within standard TC systems via well-known strategies, such as covert variables or extra parameters. The hard challenge is that the compositional semantic values of polysemous items cannot be identified with their extension, understood as a set of actual or possible objects. This is the point of examples of copredication like (6). In (6), the same occurrence of *book* composes with both *interesting* and *was ruined*; but the two predicates ascribe properties to different kinds of objects (a body of information contained in the book, and the physical book). Hence the meaning of the predicates must combine with something different than an object, or a set thereof.

The challenge is very serious and cannot be brushed aside. To be sure, I am more optimistic than Pietroski about the prospects of solving it within standard frameworks. But this will require substantial work, and (most likely) more sophisticated views of lexical meaning than we currently have.

²I should note that some attempts at solving the problem in a fairly conservative way have been put forward: see e.g. Liebesman and Magidor 2017. At this stage, I am skeptical of this style of solution. In any case, it's interesting to consider whether the challenge can be met by pursuing a route other than by retrofitting the metaphysics to the needs of the semantics.

4 Composition rules

In standard TC theories, all linguistic items are assigned a semantic type. Some items have basic types, like e (for entities) or t (for truth values). Others are assigned higher types, recursively defined from the basic ones. The type-driven system exploits one basic mechanism of semantic composition, i.e. *functional application*: items of lower type saturate the argument places of items of higher type.

A system of this kind allows for lexical items with any number of argument places. Indeed, on classical TC frameworks, it is commonplace to treat some items as having three, four, or more argument places. Even setting aside arguments related to tense and modality, transitives verbs like *kick* are taken to have three arguments (one for an event, one for the subject and one for the object), and ditransitives like *give* four. Pietroski argues that, by contrast, all linguistic items are at most dyadic, and indeed are mostly monadic.

To be clear, Pietroski grants that we have supradyadic *concepts*. For example, the concept $\text{BETWEEN}(_, _, _)$ is triadic. But these concepts are not fetched directly by linguistic items. For illustration, he points out that natural language does not involve a triadic verb like *twixt*, such that (7) is synonymous with (8).

- (7) Plum twixted Scarlet White. (Pietroski 2018, p. 259)
(8) Plum was between Scarlet and White.

More in general, building on a neodavidsonian picture (see a.o. Williams 2021), Pietroski suggests that several items that are taken to be ordinary arguments of a verb on a Fregean picture should rather be treated as external arguments, which are linked to the verb via a thematic relation. For example, the verb *chase* is analyzed as having only arguments for an event and an object, and not a subject. The subject is analyzed as an external argument, related to *chase* via the thematic relation AGENT and the event argument. So e.g. the LF of (9a) is in (9b).

- (9) a. A dog chased a cat
b. $\exists e[\exists y \text{AGENT}(e, y) \wedge \text{DOG}(y) \wedge \exists x[\text{CHASE-OF}(e, x) \wedge \text{CAT}(x)]]]$

This constrained picture of adicity is paired with a nonstandard picture of composition rules. Pietroski rejects functional application and argues that composition happens via two conjunction-type operations, M-join and D-join. I already discussed M-join in §1: M-join is an instruction to take two monadic concepts and form their conjunction. As a result, *heavy book* expressed the composite concept $\text{HEAVY}(_) \wedge \text{BOOK}(_)$. D-join works in a similar way. Specifically, D-join takes as input a dyadic and a monadic concept $\Phi(_, _)$ and $\Psi(_)$ and forms the complex concept $\exists x \Phi(_, _x) \wedge \Psi(_x)$. Using an example of Pietroski's (p. 110), consider:

(10) above Bessie

Pietroski assumes that *above* fetches a dyadic concept $\text{ABOVE}(_,_)$ and *Bessie* a monadic one $\text{BESSIE}(_)$ (more about this below). The two concepts combine via D-join to form the following concept:

(11) $\exists_x \text{ABOVE}(_, -x) \wedge \text{BESSIE}(_)$

(11) applies to everything that is above Bessie—as is intuitive.

To recap: Pietroski defends two connected claims about adicity and composition rules, both in conflict with the classical Fregean picture of composition:

- (i) Most lexical items in natural language have lower adicity than on classical TC theories: in particular, they are all monadic or dyadic.
- (ii) The composition rules for natural language involve conjunction operations, and do not include functional application.

Crucially, (ii) is supposed to provide at least part of an explanation for (i). The reason why natural language only lexicalizes monadic and dyadic concepts is that compositional operations are simple, and cannot handle supradyadic concepts. By contrast, a Fregean picture, which countenances functional application, overgenerates, since it leads us to expect items of any adicity.

In the rest of this section, I will not question (i). Rather, I will focus on a specific challenge for (ii), coming from proper names. I have two goals. First, I want to put forward what I take to be a genuinely plausible argument in favor of some version of functional application. Second, I want to focus on a case study to show how the fate of Pietroski's framework is entangled with empirical debates about specific expressions in language.

Since Pietroski abandons functional application, he is also committed to claiming that there are no *e*-type items in language (p. 108). On the classical view, proper names like *Ralph* are *e*-type items. Names combine with predicates like *sleeps*, which are of type $\langle e, t \rangle$, via functional application. An alternative view, which Pietroski needs to endorse (p. 235), is predicativism. Predicativism holds that names are semantically analogous to predicates and hence are of type $\langle e, t \rangle$. Predicativism is motivated by examples like (12):

(12) There were three Ralphs at the party.

There are many versions of predicativism, but they share some core theses. (i) Proper names are of the same semantic type as common nouns. (ii) In bare singular occurrences (as in 13a), proper names are part of a determiner phrase which is formed via a covert determiner DET (as in 13b), which is akin in meaning to the definite article.

- (13) a. Ralph is a spy.
- b. [DET Ralph] is a spy:

As a result, (13a) is a phonological variant of (14):³

(14) The Ralph is a spy.

Predicativism has been at the center of much debate, both in the Seventies and recently.⁴ While the dust hasn't settled yet, it seems clear to me that the evidence speaks against it. I don't have the space to discuss the issue exhaustively, but I review two arguments that I find probative.

First, predicativism predicts that, since names are of the same semantic type as predicates, they should license *one*-anaphora (King 2006; Hawthorne and Manley 2012; Gray 2018⁵). This is what happens with predicates that appear in a definite phrase, as illustrated by (15a). But proper names do not license anaphora, as (15b) shows.

- (15) a. The dog_i barks whenever he sees another one_i.
b. ?? Ralph_i smiles whenever he meets another one_i

The only way for the predicativist to account for the contrast between (15a) and (15b) is to stipulate that DET somehow blocks anaphora, and that this blocking occurs robustly across languages. This seems implausible.

Second, predicativism predicts that names have scope interactions with other items, including other DPs, modals, and tenses. But these interactions are unattested. For an illustrative example, take the minimal pair in (16) (from Schoubye 2016, building on examples from Gray 2012):

- (16) a. In each of my classes, Louise asked the best question.
b. In each of my classes, the person called Louise asked the best question.

(16b) has a reading where the semantic value of the description *the person called Louise* is dependent on the phrase *in each of my classes* (each class includes a person called Louise, and that person asked the best question in that class). No such reading is available for (16a), contrary to the predicativist's predictions.⁶

Suppose we accept that predicativism is false, and that names are *e*-type. How can this be accommodated in Pietroski's picture? So far as I can see, the only way is to bring back a kind of functional application: names compose with verbs by saturating one of their arguments.

³Different versions of predicativism might assign different meanings to the covert determiner; for example, it might be taken to be a covert demonstrative.

⁴For some predicativist views, see a.o. Burge 1973, Matushansky 2008, Fara 2015.

⁵I should note that Gray 2018 is a predicativist in the sense that he thinks that *some* proper names have occurrences of type $\langle e, t \rangle$. But, crucially, Gray rejects the assumption that bare occurrences of names as in (13a) are part of a complex determiner phrase, as in (13b); according to him, bare occurrences of names are genuinely *e*-type.

⁶The classical observation that names and descriptions differ with respect to rigidity is, of course, a special case of the issue of scope interactions.

To be sure, this doesn't require bringing back the full Fregean picture. We could invoke a type-specific version of functional application, which only applies to items of type $\langle e, t \rangle$ and e . Here is a simple statement of this rule, in a framework in the style of Heim and Kratzer (1998):

E-type functional application (EFA)

If α is a branching node and $\{\beta, \gamma\}$ the set of its daughters, then, if β is of type e and γ is of type $\langle e, t \rangle$, $\llbracket \alpha \rrbracket = \llbracket \gamma \rrbracket(\llbracket \beta \rrbracket)$.

Accepting EFA is consistent with holding on to the bulk of Pietroski's program. In particular, claim (i)—natural language does not allow for items of adicity greater than 2—could survive.

At the same time, this move seems to undermine the explanatory advantage enjoyed by Pietroski's framework over the Fregean. As I mentioned, one strong argument for more constrained composition rules like M-join and D-join is that they lead to the prediction that there are no supradyadic items. But once Pietroski endorses EFA, the Fregean can make a parallel move. They can drop the full version of functional application, and endorse type-specific versions of it, like EFA. Of course, there is a stipulative element in this—if our language module allows for an operation of functional application, why is this operation restricted to items of low types? But this stipulative element is shared by both them and the revised version of Pietroski's view that I'm envisaging.

In sum: Pietroski's composition rules provide an interesting explanation of the generalization in (i). I have raised a challenge to this explanation from one particular case study, namely proper names. Whether this challenge is viable or not, it illustrates how the success of Pietroski's program is tied to empirical details concerning the semantics of particular expressions. High-level semantic theorizing and empirical details are closely intertwined in this case.

5 Conclusion

Conjoining Meanings is an extraordinarily rich book, covering an impressive amount of empirical material and locating it in the backdrop of a bold big picture. Let me close with an observation about this big picture. It is tempting to see the book as selling a package deal: it provides a sustained argument against the TC view, and in favor of the procedural view, which is built around conjunctive composition rules. But this take obscures the fact that the two main theses of the book are, so far as I can see, largely independent. The procedural conception could be combined with a traditional Fregean picture of composition. Conversely, Pietroski's conjunctive rules can be plugged into a more standard TC picture of meaning. So each of the two theses needs to be considered and evaluated on its own.

References

- Tyler Burge. Reference and proper names. *Journal of Philosophy*, 70(14):425–439, 1973. doi: 10.2307/2025107.
- David J. Chalmers. The foundations of two-dimensional semantics. In Manuel Garcia-Carpintero and Josep Macia, editors, *Two-Dimensional Semantics: Foundations and Applications*, pages 55–140. Oxford University Press, 2006.
- Nate Charlow. Prospects for an expressivist theory of meaning. *Philosophers' Imprint*, 15:1–43, 2015.
- Noam Chomsky. Language and nature. *Mind*, 104(413):1–61, 1995. doi: 10.1093/mind/104.413.1.
- Noam Chomsky. *New Horizons in the Study of Language and Mind*. Cambridge University Press, 2000.
- Ivano Ciardelli, Jeroen Groenendijk, and Floris Roelofsen. *Inquisitive Semantics*. Oxford University Press, 2018.
- Donald Davidson. Truth and meaning. *Synthese*, 17(1):304–323, 1967a. doi: 10.1007/bf00485035.
- Donald Davidson. The logical form of action sentences. In Nicholas Rescher, editor, *The Logic of Decision and Action*, pages 81–95. University of Pittsburgh Press, 1967b.
- Fred I. Dretske. *Knowledge and the Flow of Information*. MIT Press, 1981.
- Chelsea M. Eddington and Natasha Tokowicz. How meaning similarity influences ambiguous word processing: the current state of the literature. *Psychon Bull Rev*, 22:13–37, 2015. doi: <https://doi.org/10.3758/s13423-014-0665-7>.
- Delia Graff Fara. Names are predicates. *Philosophical Review*, 124(1):59–117, 2015. doi: 10.1215/00318108-2812660.
- Allan Gibbard. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press, 1990.
- Allan Gibbard. *Thinking How to Live*. Harvard University Press, 2003.
- Aidan Gray. *Names and name-bearing: An essay on the predicate view of names*. PhD thesis, University of Chicago, 2012.
- Aidan Gray. Lexical-rule predicativism about names. *Synthese*, 195(12):5549–5569, 2018. doi: 10.1007/s11229-017-1462-4.

- John Hawthorne and David Manley. *The Reference Book*. Oxford University Press, 2012.
- Irene Heim and Angelika Kratzer. *Semantics in Generative Grammar*. Blackwell, 1998.
- Jeffrey C. King. Singular terms, reference and methodology in semantics. *Philosophical Issues*, 16(1):141–161, 2006. doi: 10.1111/j.1533-6077.2006.00107.x.
- David Liebesman and Ofra Magidor. Copredication and property inheritance. *Philosophical Issues*, 27(1):131–166, 2017. doi: 10.1111/phis.12104.
- Ora Matushansky. On the linguistic complexity of proper names. *Linguistics and Philosophy*, 31(5):573–627, 2008. doi: 10.1007/s10988-008-9050-1.
- Paul M. Pietroski. *Events and Semantic Architecture*. Oxford University Press, 2004.
- Paul M. Pietroski. *Conjoining Meanings; Semantics Without Truth Values*. Oxford University Press, 2018.
- Anders J. Schoubye. The predicative predicament. *Philosophy and Phenomenological Research*, pages 571–595, 2016. doi: 10.1111/phpr.12336.
- Robert Stalnaker. Assertion. *Syntax and Semantics (New York Academic Press)*, 9:315–332, 1978.
- Alexander Williams. Events in semantics. forthcoming in *Cambridge Handbook for the Philosophy of Language*, 2021.
- Seth Yalcin. Epistemic modals. *Mind*, 116(464):983–1026, 2007. doi: 10.1093/mind/fzm983.
- Seth Yalcin. Nonfactualism about epistemic modality. In Andy Egan and B. Weatherson, editors, *Epistemic Modality*. Oxford University Press, 2011.