

PRESUPPOSITIONS, ATTITUDES, AND WHY THEY MATTER

Caleb Perl

ABSTRACT: This paper introduces and defends a high-level generalization about the way that presupposition triggers interact with attitude verbs. This generalization tells us a great deal about what an adequate account of presupposition would have to look like. And it reveals one underappreciated way that presupposition is philosophically interesting.

Keywords: Presupposition, Attitude ascriptions, Epistemology

Utterances of (1a) and (1b) are standard examples of utterances that carry *presuppositions*.

(1a) John has stopped wearing plaid.

(1b) John hasn't stopped wearing plaid.

Make either utterance, and your hearer will take you to assume that John used to wear plaid. That is, uses of *stop* trigger a presupposition about the past. A wide range of philosophically interesting constructions trigger presuppositions. Ascriptions of knowledge presuppose that the complement is true; the use of a definite description presupposes that something satisfies the description; the normative or epistemic uses of some terms are plausibly associated with substantive presuppositions.¹

This paper argues that presupposition triggers behave very differently under attitude ascriptions than most other kinds of linguistic constructions, in a way that philosophers and linguists have failed to appreciate. So it attempts to establish a high-level empirical generalization about presupposition triggers. Theorists have noticed particular instances of this generalization. But they have missed its generality.

The high-level generalization is important because it tells us a great deal about what an adequate account of presupposition would have to look like. Some theorists offer *semantic* theories of presupposition, where presupposition triggers semantically encode something that distinguishes pre-

¹Several philosophers have argued that presupposition is important for understanding normativity (Alex Silk (2016, 2017) and Caleb Perl (2019)), epistemology (Daniel López de Sa (2008, 2015), Matthew Mandelkern (2017), and Caleb Perl (2017)), and as well as more linguistic topics, like complex demonstratives (Michael Glanzberg and Susanna Siegel (2006)).

supposition triggers from other kinds of constructions. These semantic accounts contrast with *pragmatic* theories of presupposition, where general conversational principles explain why we associate some utterances with presuppositions.

Standard kinds of pragmatic accounts do not predict my high-level generalization, though semantic theories do. So if my generalization is correct, we would need to rethink how to develop pragmatic theories of presupposition. This observation isn't a criticism of pragmatic accounts; I'm myself sympathetic to pragmatic accounts. The point is instead that the high-level generalization tells us a great deal about the form that pragmatic theories need to take. Getting clear on the high-level generalization helps us to develop an adequate pragmatic theory of presupposition.

Part I of this paper introduces the high-level generalization. It also shows that none of the standard tools available to pragmatic theorists explain this high-level generalization. Part I also argues, more strongly, there's no way to modify those standard tools to do what's needed. Then Part II will sketch some non-standard tools capable of capturing the relevant facts, while mentioning some philosophical payoffs.

Part I: Pragmatic theorists must expand their toolkit

1 Background about presupposition

This paper is about the class of presupposition triggers. I'll take this class to be a class of linguistic constructions that are associated with commitments that ordinarily *project*. A commitment of a sentence projects when (i) rational hearers will infer that the speaker accepts that commitment if the speaker is sincere in using the sentence, and (ii) rational hearers will also take the speaker to accept the same commitment if the speaker sincerely utters certain canonical embeddings of the sentence. In particular, a commitment p of a sentence S projects if (i) sincere uses of S ordinarily carry the commitment that p , and (ii) sincere utterances of \lceil not- S \rceil , \lceil perhaps S \rceil , and \lceil if S , then T \rceil also tend to carry that same commitment. For example, *John stopped wearing plaid* is associated with the commitment that John used to wear plaid. And that commitment is also associated with sentences like *John [didn't stop/perhaps stopped] wearing plaid* and *If John stopped wearing plaid, his dog is happier*.² The commitment projects.³

²I throughout italicize terms and sentences that I intend to mention.

³Gennaro Chierchia and Sally McConnell-Ginet (2000) helpfully discuss this sort of test for projection behavior in more detail.

Presupposition triggers aren't the only constructions that exhibit projective behavior. For example, non-restrictive relatives also exhibit projection behavior. *Mary's mother, who likes Mary, lives near her mother* is associated with the commitment that Mary's mother likes her. But this commitment also projects through these embeddings; it's also a commitment of *Mary's mother doesn't live near Mary, who likes her mother*.

The class of presupposition triggers differs from the broader class of constructions with projective behavior in part because the projection behavior of presuppositions is *defeasible*. Imagine someone saying: *John never wore plaid, so he didn't stop wearing plaid*. Even though *stop* triggers a commitment that usually projects through *not*, it doesn't in this case. The discourse is perfectly felicitous, and it wouldn't be if the commitment projected. Non-restrictive relatives pattern differently.⁴ The second hallmark of presupposition triggers is thus that their projection behavior is fairly defeasible – much more defeasible than other kinds of projective contents.

So understanding presupposition triggers requires understanding why the presuppositions project. Importantly, it's not just an idiosyncratic fact about some presuppositions that they project. Rather, *any* term that carries the same commitments as a presupposition trigger seem to *always* exhibit the same projection behavior. Take *stop*. Sentences containing *stop* are associated with two commitments: a commitment that something was happening in the past, and a commitment that that thing isn't happening in the present. It turns out that *any* term associated with those two commitments exhibits the same projection behavior. That is, when that term is embedded under a negation operator, the first commitment projects, and the second one is interpreted under the negation operator. This point holds across natural languages, even typographically distinct ones.⁵ This robust generalization calls out for explanation.

1.1 Pragmatic explanations of presupposition

Robert Stalnaker has developed a powerful and influential explanation of presupposition triggers that predicts that cross-linguistic uniformity of presupposition projection.⁶ He starts with a simple conjecture about the semantics of presupposition triggers like *stop*.

⁴Chris Potts (2005) helpfully works through some of those details (12 – 36)

⁵One of the central contributions of Tonhauser et al. (2013) is to show that this generalization does hold generally, even across typographically distinct languages.

⁶See Stalnaker (1973, 1974, 2002, 1999).

(1a) John has stopped wearing plaid.

For him, (1a) expresses a single proposition that entails that John used to wear plaid and isn't anymore.⁷ So what does (1b) express?

(1b) John hasn't stopped wearing plaid.

Stalnaker's answer is simple. (1b) expresses the proposition formed by composing the semantic contribution of *not* with the semantic contribution of (1a). For him, then, the semantics of presupposition triggers works in exactly the same way as the semantics for other linguistic constructions. Why does *not* compose with presupposition triggers in the same way across languages? For the same reason that *not* composes with an arbitrary sentence in the same way across languages. Nothing about the semantics of *stop* tells us anything about its projection behavior.

Stalnaker adds that presupposition is fundamentally something that speakers do, rather than a fundamental property of sentences or utterances. For him, a speaker presupposes a proposition if she takes it to be part of the common ground – that is, part of what parties to the conversation accept, and take others to accept, and so on. (Throughout the rest of the paper, I'll assume Stalnaker's common-ground-based conception of speaker presupposition. This conception of speaker presupposition is the most familiar one, so it's the best conception to assume as the baseline in describing the innovation that I want to make.⁸) The notion of utterances presupposing a proposition is less fundamental. An utterance of a sentence *S* presupposes *p* if and because speakers who use *S* tend to presuppose *p*.

The task of explaining the cross-linguistic uniformity of presupposition projection then reduces to a simpler task. We need to explain why we tend to interpret someone utters (1b) as presupposing something about the past.

Stalnaker executes this task by drawing on general facts about rational communication. In fact, he explores several options for doing so. One option focuses on *questions*. He asks about the conditions under which it's interesting to investigate whether (1a) is true or (1b) is true. Maybe it's

⁷In this paper, I'll simplify by assuming that sentences like (1a) express propositions, rather than something more complicated, like propositional matrices. I'm making this simplification just for terminological reasons. If you prefer, read occurrences of “the proposition that (1a) expresses” as “the propositional matrix that (1a) expresses”, or whatever your preferred ideology is.

⁸As it happens, I don't myself think that this conception of speaker-presupposition is the best one. I myself favor a view that builds from the framework that Craige Roberts (2012) has developed, as sketched in, for example, Simons et al. (2011).

normally only interesting to ask whether John used to wear plaid and isn't anymore if you already assume that John used to wear plaid, and want to figure out whether his clothing habits have changed. As Stalnaker puts the idea, “the propositions that P and that Q may be related to each other, and to common beliefs and intentions, in such a way that it is hard to think of a reason that anyone would raise the question whether P, or care about its answer, unless he already believed that Q” (Stalnaker 1974, 205).

Stalnaker's approach to presupposition is standardly called a *pragmatic* theory of presupposition, because the theory builds from Gricean assumptions about rational communication. Many others have followed in Stalnaker's footsteps in developing pragmatic theories of presupposition.⁹ Pragmatic theories are attractive because they promise to explain the cross-linguistic uniformity of projection. If presuppositions project because of general principles about rational communication, it's unsurprising that presupposition project the same way across all natural languages.

The upshot of Stalnaker's account of speaker-presupposition is that:

If you assertively utter (1b) (*John hasn't stopped wearing plaid*),

- you presuppose that John used to be wearing plaid.

This paper will focus on a further question that arises given a pragmatic account of presupposition. The question is: what does your use of (1b) ordinarily assert? There are two possible answers:

- you assert that \neg (John used to be wearing plaid and is currently wearing plaid).
- you assert that \neg (John is currently wearing plaid).

The first answer is that (1b)'s semantic value is asserted. And the second answer is that a proper part of its semantic value is asserted. I'll call the first answer the *Atom* answer, because it treats the semantic value of the embedded sentence as an atom – its contribution to what's asserted is not “broken up”. And I'll call the second answer the *Fission* answer, because it takes the semantic value of the embedded sentence to be “broken up” in making a contribution to what's asserted.

At first blush, only the Fission answer looks defensible. Competent hearers of a use of (1b) will infer that the speaker intends to communicate that John isn't currently wearing plaid. And in general, if someone asserts

⁹Examples include Philippe Schlenker (2010), Mandy Simons (2001), Robert Stalnaker (1973, 1974, 2002, 1999), and Mandy Simons et al. (2017).

something with the form $\neg(A \text{ and } B)$, it's a mistake to infer that they intend to communicate $\neg A$, and it's a mistake to infer that they intend to communicate $\neg B$. So it seems like the Atom answer would predict that it's a mistake to infer that the speaker intends to communicate that John isn't currently wearing plaid. This prediction would be a sufficient reason for rejecting the Atom answer.

However, the first-blush impression is mistaken: the Atom answer works just as well as the Fission answer. Given Stalnaker's account of speaker-presupposition, the speaker who assertively utters (1b) already accepts that John used to wear plaid. (The speaker presupposes that he used to, and accepting that p is part of presupposing that p .) And if the speaker asserts $\neg(A \text{ and } B)$ while accepting A , competent hearers will infer that she intends to communicate $\neg B$. In the first place, accepting $\neg B$ is the only way for what she asserts to be consistent with what she accepts. In the second place, the speaker can't intend to communicate $\neg A$; $\neg A$ is inconsistent with what she accepts. And she can't rationally intend to communicate merely $\neg(A \text{ and } B)$, because her communicative act itself reveals that she accepts something more informative: $\neg B$. As a result, the Atom answer *can* explain why hearers will infer that the speaker intends to communicate $\neg B$; that's the only way for the speaker's communicative act to be rational.¹⁰

And if the Atom and Fission answers are both viable, there are strong theoretical reasons for favoring the Atom answer. For one thing, the Atom answer is dramatically simpler. If it's right, Stalnaker has done everything he needs to do once he has explained why speakers who use (1b) tend to presuppose that John used to wear plaid. But given the Fission answer, Stalnaker *further* needs to explain why the speaker's presuppositions affect what the speaker asserts. So the Fission answer will involve complications that the Atom answer doesn't. Parsimony favors the Atom theory.¹¹

¹⁰Importantly, then, the disagreement between the Atom theory and the Fission theory is not a disagreement about presupposition projection; it's fully compatible with any theory of presupposition projection that's compatible with a pragmatic theory of presupposition. Both theories build from Stalnaker's account of speaker-presupposition, which is an account of why the presupposition that *stop* triggers projects. No pragmatic theorist could accept either the Atom theory or the Fission theory unless she already had an account of why that presupposition would project. So if a pragmatic theorist can make good on the predictions of the constraint-like conception from theorists like Irene Heim (1982, 1983), David Beaver (2001), or Daniel Rothschild (2011), both the Atom and the Fission theorist can too – and ditto for the anaphoric conception from theorists like Rob van der Sandt (1992) and Bart Geurts (1999).

¹¹Stalnaker himself accepted the Atom theory, as he claims, “the facts about presuppositions, I am suggesting, can be separated from a particular kind of

Moreover, the Fission answer involves something highly contentious: *pragmatic modulation* of the content that *stop* semantically expresses. Given the Fission answer, the contribution that *stop* makes to the proposition asserted by (1b) is narrower than what *stop* semantically expresses – that is, there is pragmatic modulation of what’s semantically expressed. Philosophers like Jason Stanley (2007) vigorously resist such pragmatic modulation. Compare the use of *drink* to assert something just about alcoholic drinks. Philosophers like Stanley posit further argument-places that are contextually saturated to explain these uses of *drink*, to avoid pragmatic narrowing of what’s semantically expressed. But pragmatic theories of presupposition cannot posit further argument-places in adopting the Fission answer. Positing further argument-places requires the semantics for presupposition triggers to contain something distinctive that’s absent in the semantics for other constructions. So giving a Fission answer must involve pragmatic modulation, which is highly contentious.

Extant pragmatic theorists have followed Stalnaker in adopting the Atom answer. Schlenker, for example, writes that “the semantics treats the presuppositional component d of an elementary expression dd' as if d were just part of its assertive component. But the pragmatics will give d a distinguished status through the requirement that the presupposition of an elementary expression be entailed by its local context” (Schlenker 2010, 384). The proposal in Simons et al. (2017) is similar, though less explicit. Abstracting from many complications, they take what’s asserted to be an answer to a salient question, while further holding that something is presupposed if it’s entailed by every answer to the salient question. In that case, what’s asserted must entail the presupposition; their theory is an instance of the Atom theory. It’s unsurprising that they adopt the Atom answer, since that answer is simpler and philosophically much less contentious.¹²

Pragmatic theories of presuppositions like Stalnaker’s contrast with other approaches where the semantics of presupposition triggers *does* encode something distinctive. A prominent example is the dynamic approach that Irene Heim pioneered. She takes words to express (something that contributes to) ‘context change potentials’: functions from bodies of information to updated bodies of information. For her, *stop* will contribute to

semantic explanation of those facts. This separation of the account of presupposition **from the account of the content of what is said** will allow for more diversity among presupposition phenomena than would be possible if they all had to be forced into the semantic mold” (Stalnaker 1974, 53, bold mine).

¹²Others who I take to accept the Atom theory include at least Paul Elbourne (2005), Bart Geurts (1998), Irene Heim (1992), and Judith Tonhauser, David Beaver, Craige Roberts, and Mandy Simons (2013), which is earlier joint work related to Simons et al. (2017).

a different context change potential than *used to and isn't now* does.¹³ For her, the lexicon does specify the presuppositional behavior of *stop*.

As suggested in the introduction, I'll introduce a high-level generalization, and argue that the high-level generalization tells us a great deal about what an adequate account of presupposition needs to like look. Semantic approaches like Heim's are immediately capable of explaining the high-level generalization I'll discuss in this paper. By contrast, it's an open question whether pragmatic accounts like Stalnaker's can capture that high-level generalization. As a result, my discussion of the high-level generalization will this focus on pragmatic views, since they are the ones for which the high-level generalization is pressing.

1.2 Presupposition triggers and attitude reports

The central question in this paper will be about presupposition triggers under attitude reports, like (2).

(2) I'm glad that John has stopped wearing plaid.

I'll explore what uses of (2) ordinarily assert.

Attitude verbs like *glad* are a kind of embedding, like embedding under *not*. And, as suggested in the previous section, there are two possible kinds of answers about what's asserted by the use of a sentence that embeds a presupposition trigger. An *Atom* answer is that the proposition asserted contains the semantic value of the embedded sentence, and a *Fission* answer is that the proposition asserted contains a proper part of the semantic value. So its possible to develop either an *Atom* theory or a *Fission* theory about what's asserted by (2).

- *ATOM* theory: I only use (2) to assert that I'm glad about both John's past and present activity.
- *FISSION* theory: I can use (2) to assert only that I'm glad just about John's *present* activity. In those cases, the only proposition communicated about my gladness is that I'm glad that John is not wearing plaid now.

Importantly, the *Atom* theory and the *Fission* theory are only theories about what can be *asserted* by the use of (2). They're entirely silent about the correct semantics for (2). The correct semantics for (2) could be the

¹³Irene Heim (1992) pioneered this idea, and it's broadly been taken up – for example, in Elbourne (2005) and in Geurts (1998).

straightforward semantics that Stalnaker accepts, where (2) semantically expresses the result of composing the semantic value of *I'm glad* with the semantic value of the complement.

Another option is for the correct semantics for sentences containing presupposition triggers to be *multi-dimensional*, expressing several distinct propositions. For example, (2)'s complement might semantically express:

- that John isn't wearing plaid, and
- that John used to be wearing plaid

Kent Bach (1999) has developed just such a suggestion. Importantly, though, this multi-dimensional semantics is compatible with *either* a Fission or an Atom theory. It's compatible with either because the Fission/Atom theories are simply about a different topic: about what's asserted, rather than what's semantically expressed. Now the Fission theory is a more natural fit for these multi-dimensional approach, but the multi-dimensional approach doesn't entail the Fission theory.

Over the next two sections, I'll introduce a high-level generalization about the behavior of presupposition triggers under attitude reports, and argue that this generalization requires pragmatic theorists to accept the Fission Answer about (2). So throughout the rest of the paper, I'll focus just on the Atom/ Fission theories as theories about presupposition triggers *under attitude verbs*.¹⁴ So one possible reaction to this paper would be to accept a Fission theory about attitude verbs but accept an Atom theory about all other embeddings. But there are plenty of reasons that even just a Fission theory about attitude verbs would be philosophically interesting, as we'll see later in the paper.

2 Why Atom theories are false

Atom theories are more ambitious than Fission theories. They attempt to explain the data by using a narrow toolkit – without supposing that the semantic value of the complement is “broken up” in interpreting some uses of (2). (Remember that one central difference between the Atom and Fission theories is that the Atom theories are dramatically simpler.) This section shows that this narrow toolkit can't work. It can't explain an important high-level generalization.

¹⁴I suspect that the two theories are empirically equivalent outside of hyper-intensional contents.

2.1 Basic data

Consider (2), again.

(2) I'm glad that John has stopped wearing plaid.

My initial observation is that (2) can be true even if I'm not glad that John used to wear plaid. There are several ways for Atom theorists to explain this initial observation. However, there is an additional observation that Atom theorists cannot explain, given their explanations of the initial observation. The pair of observations are the grist for my high-level generalization.

Atom theorists can explain my initial observation by taking it to follow from a more general fact about gladness. In general, you don't need to be glad about all the entailments of S to be glad that S. Consider (3a).

(3a) I'm sad that you're buying a car. But I'm still glad that you're buying a hybrid car. That minimizes the environmental damage of your choice.

If you're buying a hybrid car, you're buying a car. But I can be glad about the former even though I'm not glad about the latter. So there has to be some explanation of why you can be glad that S without being glad about everything that S entails.

My second important observation is that presupposition triggers are associated with a privileged commitment: a commitment that you *have* to be glad about in order for the gladness report to be appropriate. The privileged commitment associated with *John has stopped wearing plaid* is that John doesn't wear plaid. You can't be glad that John has stopped wearing plaid unless you're glad about that privileged commitment. As evidence for this fact, contrast (4a) and (4b):

(4a) James is sad that John used to wear plaid. He's still glad that John has stopped wearing plaid.

(4b) James is sad that John doesn't wear plaid. * He's still glad that John has stopped wearing plaid.

This fact is important, because other kinds of linguistic constructions behave differently. They do not display the same privileged commitment. Imagine that you were torn between buying a hybrid car and a hybrid

Vespa, and I'm relieved that you opted for the former. (I think that the former is more normal, and makes you less of a hippie.) Then (3b) is true.

(3b) I'm sad that you're buying something hybrid. But I'm still glad that you're buying a hybrid car. That's more normal.

In other words: (3a) is true if the speaker wants you to *minimize* damage to the environment, and (3b) is true if the speaker wants you to *maximize* damage. So *you're buying a hybrid car* doesn't carry any privileged commitment that I need to be glad about in order to be glad that you're buying a hybrid car.¹⁵

The high-level generalization that I'll describe is a generalization about the difference between (3a)/(3b), on the one hand, and (4a)/(4b), on the other. But to appreciate why this high-level generalization is important, we need to appreciate what Atom theorists have to say about the relationship between (3a)/(3b) and (4a)/(4b).

2.2 Attitude ascriptions without privileged commitments

Atom theorists explain the way that *stop* embeds under attitude verbs by appeal to general facts about attitude verbs, rather than by appeal to special facts about presupposition triggers. So to evaluate this claim, we have to understand how attitude verbs behave.

We've seen that you can be glad that S without being glad about everything that S entails. I'll describe this fact as the fact that *is glad* isn't closed under (logical) consequence. This section describes two representative accounts of this fact. It then identifies a baseline feature that these accounts share, and it explains why any plausible account of *is glad* has to share that baseline feature. The next subsection will show that that baseline feature prevents any of those accounts from explaining the way that presupposition triggers embed under attitude verbs. In other words: I'll be arguing that the two pairs in the previous section illustrate two distinct phenomena. You can't use the same tool to capture both.

It's not just ascriptions of *gladness* that fail to be closed under consequence.

¹⁵The point here is just that the sentence "I'm glad that you're buying a hybrid car" can be used appropriately both of these cases. The two uses may have different explanations; for example, one might express conditional wants, and the other unconditional wants.

(5) I want you to buy a hybrid car.

(5) can be true even if I don't want you to buy a car, but believe that you will anyway. And *wants* is in several ways simpler than *is glad*.¹⁶ So I'll use it to introduce two representative accounts of failure of consequence.

The first account starts from a Hintikka-style semantics for attitudes ascriptions. $\lceil A$ wants that $S \rceil$ is true iff all the worlds compatible with what A wants are worlds where the proposition that S is true. But this first account goes beyond the basic Hintikka-style semantics, by further restricting the worlds it quantifies over.

One version of this Hintikka-style approach further restricts to worlds compatible with what the matrix subject *believes*. As a first pass, take $\lceil A$ wants that $S \rceil$ to be true iff all the worlds compatible with what **A believes** and also what A wants are worlds where the proposition that S is true. To illustrate, go back to (5). Model my wants with a set of worlds that includes some worlds where you don't buy a car. An unmodified Hintikka-style semantics then predicts that (5) is false. There are worlds compatible with my wants where the complement is false. The modified Hintikka semantics, by contrast, correctly predicts that (5) is true. We no longer quantify over worlds where you don't buy a car, because we only quantify over worlds compatible with my beliefs. And my beliefs entail that you do buy a car.¹⁷ I'll call this first account a *modified-Hintikka* account.

The second account assumes that propositions have more internal structure than mere sets of worlds. Propositions are either built up out of objects, properties, or relations, or built up out of concepts for those things. So even if S entails T , there is no interesting and general relationship between the truth of $\lceil A$ wants that $S \rceil$ and the truth of $\lceil A$ wants that $T \rceil$. And there is a very general reason why there isn't any interesting relationship between the two: entailments between the complements are just irrelevant. For one thing, every sentence entails that arithmetic is incomplete, trivially. And proponents of structured propositions already deny that there is any interesting relationship between having an attitude towards an arbitrary proposition (say, that it's raining now), and having that attitude to the proposition that arithmetic is incomplete. I'll call this second explanation a *structured-proposition* account.

There are other possible accounts, too.¹⁸ I'll assume that one such

¹⁶For example, it's not factive.

¹⁷Irene Heim (1992) develops a sophisticated version of this account; the next page will dig into some further important features of her account.

¹⁸A third possible kind explanation builds from the sort of *alternative semantics* that Mats Rooth (1985) has developed; Elizabeth Villalta (2000) develops this kind of proposal; it's used to good effect in Pranav and Hacquard (2013),

account is correct, and also that the correct account extends to ascriptions of gladness. Now these explanations all share one baseline feature. They all predict *symmetries* among entailments. Wanting/ being glad that S doesn't require wanting/ being glad about any privileged entailment of S.

(SYMMETRIES AMONG ENTAILMENTS) If S asymmetrically entails T (if S entails T and T does not entail S), it's possible to [want/be glad] that S, without [wanting/being glad] that T.

It's easy to see that the structured-propositions account leads immediately to the Symmetry claim. Of course it's possible to V that S without V-ing that T – they're just different attitudes!

The modified-Hintikka account will also vindicate this Symmetry claim. Suppose that Jane already believes T but regrets that T. Because she regrets that T, she's not glad that T. Suppose further that she's glad that R, where R and T together entail S. Then the modified-Hintikka account predicts that she is glad that S: the proposition that she's glad about (that R) plus the proposition that she believes (that T) together entail that S.¹⁹

and it's related to some ideas in Simons et al. (2017).

¹⁹There are some complications. I've been focusing on a first pass version of the modified-Hintikka account, which vindicates the SYMMETRIES claim. But the reason why it vindicates that SYMMETRIES claim is the reason why it can only be a first-pass version of the account. If you already believe that T, then the first pass account takes you to want that T. After all, there just aren't any worlds compatible with what you believe, so all the worlds compatible with your beliefs and your wants are (trivially!) worlds where T is true. This result shows that the first-pass version can only be the first pass. Even if I believe you'll buy a car, I can want you to not buy it.

We thus need to refine the first-pass presentation of the modified-Hintikka account. To do that, we should hold that $\lceil A \text{ wants that } S \rceil$ is true iff A **prefers** each similar world where S is true to any **similar** world where $\neg S$ is true, and add that worlds compatible with what she believes are more similar than worlds that aren't compatible with what she believes. Then belief that you'll buy a car isn't enough for wanting you to buy a car. We compare worlds where you don't get a car with worlds where you do. And I can prefer a world where you don't get a car. Moreover, *I want you to get a hybrid car* can still be true, even if I don't want you to get a car. We compare worlds where you don't get a hybrid car with worlds where you do. Crucially, though, there are some worlds compatible with what I believe where you don't get a hybrid car, where you get a gas-guzzler instead. This sentence is true because I prefer worlds where you get a hybrid to those worlds. (Irene Heim (1992) develops this kind of proposal, inspired in part by some ideas from Stalnaker (1984).)

As a result, basic structural features of these two accounts force them to vindicate SYMMETRIES AMONG ENTAILMENTS.

SYMMETRIES AMONG ENTAILMENTS: If S asymmetrically entails T, it's possible to [want/be glad] that S, without [wanting/being glad] that T.

But I come to praise these explanations, not bury them. Any empirically plausible theory of attitude verbs will vindicate SYMMETRIES AMONG ENTAILMENTS. Go back to the earlier pair.

(3a) I'm sad that you're buying a car. But I'm still glad that you're buying a hybrid car. That minimizes environmental damage.

(3b) I'm sad that you're buying something hybrid. But I'm still glad that you're buying a hybrid car. That's more normal.

In the first sentence, I'm glad about the option that does the least environment damage. In the second one, I'm glad about the option that does the *most*. It's thus an example of the readings that SYMMETRIES AMONG ENTAILMENTS would predict.

In general, then, a theory of attitude verbs is empirically plausible only if it vindicates SYMMETRIES AMONG ENTAILMENTS. It's a virtue of all these accounts that they do.

2.3 Presuppositions are different

This section shows that there is no way for Atom theories about presupposition to be correct. The only way an Atom theory could correct is if it

The modified-Hintikka account does predict SYMMETRIES AMONG ENTAILMENTS. If S asymmetrically entails T, there's some further sentence U such that $\lceil U \wedge T \rceil$ is true iff S is true. Now suppose that my preferences are: $(U \wedge \neg T) \gg (\neg U \wedge \neg T) \gg (U \wedge T) \gg (\neg U \wedge T)$. Suppose further that I believe that T. The sentence $\lceil I \text{ want that } S \rceil$ is then true. There are worlds compatible what I believe where S is true (worlds where U is true), and worlds where it is false (where U is false). So we only compare worlds that are compatible with what I believe, and all those worlds are worlds where S is true. Crucially, though, the sentence $\lceil I \text{ want that } T \rceil$ is false. We have to look outside the worlds compatible with what I believe to find a world where $\neg T$ is true. And I prefer any of those worlds to a world where T is true.

incorporates one of the accounts from §2.2, and incorporating one of those accounts will force some mistaken predictions.

Think of this minimal pair:

(4a) James is sad that John used to wear plaid. He's glad that John has stopped wearing plaid.

(4b) James is sad that John doesn't wear plaid. * He's glad that John has stopped wearing plaid.

The minimal pair illustrates the crucial high-level generalization that I've been advertising since the beginning of the paper.

(HIGH-LEVEL GENERALIZATION) If a sentence 'S₁' contains a presupposition triggers, it's associated with a commitment T₁ that you have to V in order to V that S₁.

If S₁ is *John has stopped wearing plaid*, the commitment T₁ is that John doesn't currently wear plaid. So the HIGH-LEVEL GENERALIZATION predicts that being glad that John stopped wearing plaid requires being glad that he's not wearing plaid. (4b) is inappropriate because you can't be glad that S₁ unless you're glad that T₁. And the first conjunct in (4b) rules out your being glad that T₁.

This section shows that Atom theorists cannot explain this Generalization. Here's an overview of the argument. According to the Atom theorist, being glad that John stopped wearing plaid involves being glad about a conjunctive proposition, about John's past and present both. (4a)'s felicity shows that you can be glad that John stopped without being glad about John's past. The Atom theorist can capture (4a)'s felicity only by tracing it back to the facts about attitude verbs described in §2.2 – that you don't need to be glad about all of S's entailments to be glad that S.

But in offering the §2.2 explanation, the Atom theorist seems to predict that (4b) is also felicitous. After all, the theories in §2.2 all predicted symmetries among entailments of the complement. And the Atom theorist takes the proposition about John's *present* to be another entailment of the complement. So if James knows that John is doesn't wear plaid anymore, but is glad that John *used* to wear plaid, the §2.2 explanations predict that *James is glad that John has stopped wearing plaid* would be true. For the modified-Hintikka explanation, for example, the propositions that James believes (that John doesn't wear plaid anymore), plus the propositions that he's glad about (that he used to wear plaid) entail that John has stopped

wearing plaid. The Atom theory plus the §2.2 explanations thus predict that (4b) could be felicitous.

(4b) James is sad that John doesn't wear plaid. * He's glad that John has stopped wearing plaid.

And that prediction is mistaken.

Now Atom theorists do have a standard way to extend the tools from §2.2 to explain (4b) and the High-Level Generalization more generally. They start with the plausible observation that we interpret the matrix subject of an attitude ascription as *believing* presuppositions of the complement. That's why 'A V-es that S' doesn't require A to V S's presuppositions: we interpret A as already believing S's presuppositions. This strategy can capture some instances of the High-Level Generalization. An utterance of 'A hopes that John has stopped wearing plaid' is only appropriate if A already believes that John used to wear plaid. And you can't hope what you believe. So that utterance is appropriate only if A hopes that John doesn't wear plaid.²⁰

But this strategy simply cannot capture the High-Level Generalization in full generality. Stipulate that James knows that John used to wear plaid, and also knows that he doesn't wear plaid now. Even given that stipulation, the High-Level Generalization still holds – that is, (4a) is still felicitous and (4b) is infelicitous.

(4a) James is sad that John used to wear plaid. He's glad that John has stopped wearing plaid.

(4b) James is sad that John doesn't wear plaid. * He's glad that John has stopped wearing plaid.

Sadness that S and gladness that S are both compatible with knowledge that S. As a result, the Atom theorist simply has to fall back on the tools from §2.2. And those tools just can't explain the High-Level Generalization.

Philosophers and linguists have missed this point in part because their diet of examples is too narrow. Matrix belief in the presupposition plus the tools from §2.2 might capture the High-Level Generalization for attitudes that are incompatible with belief. But it just can't work in full generality. It can't capture emotive factives, because they're perfectly compatible with belief, and indeed perfectly compatible with knowledge.

²⁰Paul Elbourne (2005), Bart Geurts (1998), and Irene Heim (1992) all adopt this kind of idea; Heim is particularly explicit in doing so.

The examples in (4a) and (4b), with presupposition triggers, thus illustrate a different phenomenon than the examples in (3a) and (3b), without presupposition triggers.

(3a) I'm sad that you're buying a car. But I'm still glad that you're buying a hybrid car. That minimizes environmental damage.

(3b) I'm sad that you're buying something hybrid. But I'm still glad that you're buying a hybrid car. That's more normal.

And there are highly general reasons the two pairs illustrate distinct phenomena. Any viable account of (3a) and (3b) will predict SYMMETRIES AMONG ENTAILMENTS.

SYMMETRIES AMONG ENTAILMENTS If S asymmetrically entails T, it's possible to [want/be glad] that S, without [wanting/being glad] that T.

SYMMETRIES AMONG ENTAILMENTS is just a description of the data about the first minimal pair. But an account that predicts SYMMETRIES AMONG ENTAILMENTS can't explain the High-Level Generalization. If Generalization is true, there are *asymmetries* between the entailments associated with presupposition triggers.

I conclude that Atom theories can't capture the High-Level Generalization. I've introduced the problem for Atom theorists by focusing on (4a) and (4b). But I think the problem is highly general – that it's a problem about presupposition triggers, rather than a particular point about “stop”. Other examples of presupposition triggers exhibit the same pattern.

I'll suggest that the High-Level Generalization requires pragmatic theorist to adopt Fission theories, rather than Atom theories. Now it is possible to capture the Generalization without adopting a Fission theory. For example, dynamic views like Irene Heim's can capture it. But the key point is that the High-Level Generalization must be explained somehow. And pragmatic theorists who adopt an Atom theory can't explain it.

3 Defending the High-Level Generalization

But before moving on to Fission theories, I want to defend the High-Level Generalization from three important objections.

The first objection is that I've simply made a mistake about the empirical data. I've insisted that (4b) is always inappropriate.

(4b) James is sad that John doesn't wear plaid. * He's glad that John has stopped wearing plaid.

Imagine that James is sad that John doesn't wear plaid, because James knows that John's parents want him to and John neglects what they want. Imagine further that James is also glad that John is listening to his partner, who doesn't want John to wear plaid – so James is also glad that John has stopped wearing plaid. In that case, both sentences in (4b) are appropriate. So §2.2 was wrong to insist that (4b) can't be true.

This objection involves a subtle confusion. The imagined situation is a situation where James is sad and glad about the very same proposition. That combination of attitudes is possible when you're sad for one reason and glad for a different reason. James is sad that John doesn't wear plaid because of what John's parents want, and he's glad that John doesn't wear plaid because of what John's partner wants. Crucially, though, James can't be sad that John doesn't wear plaid and glad *for the very same reason* that he doesn't wear plaid. The important minimal pair from earlier just needs to be reformulated.

(4a) James is sad that John used to wear plaid – and he's glad **for the very same reason** that John has stopped wearing plaid.

(4b) James is sad that John doesn't wear plaid – * and he's glad **for the very same reason** that John has stopped wearing plaid.

This data is the data that Atom theorists can't explain.

The next objection is that the HIGH-LEVEL GENERALIZATION neglects an important possibility: the possibility that presuppositions are *locally accommodated*. Imagine that there's a prize given out to people who formerly wore plaid but don't anymore. In that case, I could be glad that John stopped wearing plaid because I'm glad about *both* his past and present activities. Isn't that incompatible with the HIGH-LEVEL GENERALIZATION?

(HIGH-LEVEL GENERALIZATION) If a sentence 'S₁' contains a presupposition triggers, it's associated with a commitment T₁ that you have to V in order to V that S₁.

Not at all. Cases where the presupposition is locally accommodated are cases where you have to be glad about the commitment T_1 plus another commitment. That is, they require more than the GENERALIZATION requires. The GENERALIZATION is itself silent about whether there are cases where more is required.

The third challenge is that the High-Level Generalization does not hold of the following case.

There's a demon who intends to kill everybody who has never worn plaid, and reward everybody who currently wears plaid. So James' preferences are:

- **First Best:** John is currently wearing plaid.
- **Second Best:** John used to wear plaid, but isn't currently.
- **Worst:** John hasn't ever worn plaid.

I have encountered people who think that, given these preferences, and James' knowledge that John used to but isn't currently wearing plaid, someone might say (6).

(6) James is glad that John has stopped wearing plaid.

(I'm not myself sure if this example is felicitous in this context. But it's worth discussing, since some people do find it felicitous.²¹)

If this utterance is felicitous, it seems to be true without James being glad about John's present activity. If so, it looks like evidence against the HIGH-LEVEL GENERALIZATION, because (6) seems to be true without James' being glad about the presupposition triggered from the complement. That is, it seems to be evidence that the HIGH-LEVEL GENERALIZATION is not a correct empirical generalization.

This context does raise interesting questions. But I think it raises the same sort of interesting questions that the first objection raises. The first objection was about a case where someone was sad for one reason but glad for another. And to control for that possibility, I switched to the following minimal pair:

(4a) James is sad that John used to wear plaid – and, for the same reason – he's glad that John has stopped wearing plaid.

²¹Stephen Finlay has insisted that it is; this example is originally due to him.

(4b) James is sad that John doesn't wear plaid – and, for the same reason – * he's glad that John has stopped wearing plaid.

Interestingly and importantly, (4b) isn't true even in the demon case. The key challenge to the Atom theorist is to explain this contrast – to explain why (4a) can be true even while (4b) cannot be true. And the HIGH-LEVEL GENERALIZATION is my label for this challenge. (So it should be refined to apply only when we hold fixed the agent's reasons for being glad.) But we can rest secure in the conviction that the empirical contrast I describe is robust, even across the cases that put the most pressure on it.²²

4 Fission explanations of the High Level Generalization

Let's recap. §1 contrasted two different accounts of what's asserted by a use of (2).

(2) I'm glad that John has stopped wearing plaid.

- the FISSION theory holds that I can use (2) without thereby communicating that I'm glad about the past and present both. Instead, I can use it to communicate that I'm glad just about the present. In those cases, the only proposition communicated about my gladness is that I'm glad that John is not wearing plaid now.
- the ATOM theory holds that the cases that the Fission theory is about are still cases where uses of (2) only communicate the proposition semantically expressed.

§2 argued that the Atom theory can't work, because it can't explain the High Level Generalization. In this section, we'll see how Fission theories can do better. In fact, it's easy for Fission theories to do better, because

²²Further reflection on the demon example bolsters this conclusion. Some people who think (6) normally sounds infelicitous think that enough setup can make it sound better: as a referee for the *Australasian Journal of Philosophy* put it, “assume we have a context in which James has already stated that he is glad that Tom has stopped wearing plaid, meaning to express that he is glad Tom no longer wears plaid, and then someone points out the demon facts to him and that his utterance therefore suggests he is glad Tom won't get a reward from the demon. James could then say ‘Hold on, what I said also presupposes that Tom once wore plaid and therefore that Tom won't be killed by the demon, so after all I can say that I am glad he has stopped wearing plaid!’ ” But even in this case, (4b) doesn't seem felicitous.

they hold that sentences with the form ‘A V-es that S’ usually only assert that A V-es the privileged commitment of S.

The central task for the Fission theorist is to motivate her theory in the first place. After all, the Atom theory has a simple and standard picture of what’s communicated by attitude ascriptions that embed a presupposition trigger. They communicate what a standard compositional theory takes the attitude ascription to semantically express. Since the Fission theorist doesn’t agree, she owes us her own systematic account of what is communicated. I will sketch one possible approach a Fission theorist might take in developing a systematic account. My goal here is to provide a proof of concept. There are other possible approaches a Fission theory might take; in fact, I think that the best Fission approach is dramatically more complicated than the one I’ll describe here.

My proof of concept starts with a fully traditional semantics for presupposition triggers, where sentences containing a presupposition trigger semantically express a conjunctive proposition.

(2) I’m glad that John has stopped wearing plaid.

- (CONJUNCTIVE) I’m glad that John used to wear plaid and isn’t anymore.

There are at least two ways for CONJUNCTIVE to be true. It can be true if PRESENT is true, and it can be true if PAST is true.

- (PRESENT) I’m glad that John isn’t wearing plaid.
- (PAST) I’m glad that John used to wear plaid.

Either proposition is sufficient for the truth of CONJUNCTIVE, because *is glad* isn’t closed under consequence.

Importantly, though, communicating PAST or PRESENT is more *informative* than communicating CONJUNCTIVE would be. For one thing, CONJUNCTIVE is true in a wider range of worlds than PAST is true in, and a wider range than PRESENT is true in.²³ Since it’s true in a wider range of worlds, it is less informative. On learning it, the learner hearer learns less about her location in logical space than on learning PAST or PRESENT.

So far, we just know that PAST and PRESENT are tied as interpretations. The proof-of-concept that I have in mind will break this tie by appeal to a pragmatic theory of presupposition. A pragmatic theory of presupposition explains why utterances of *John stopped wearing plaid* presuppose that John

²³It’s true at least in the union of the worlds where PAST with the worlds where PRESENT is true.

used to wear plaid. And part of what it is for an utterance to presuppose p is for competent speaker-hearers to understand that p is not part of the main point of the utterance. Consider a speaker who uses (2):

(2) I'm glad that John has stopped wearing plaid.

The speaker ordinarily presupposes that John used to wear plaid. That is, she's treating the proposition about John's past activities as not the main point of her utterance. It's then less likely that she intends to communicate about her gladness about PAST, about John's past activity. That's why competent speaker-hearers tend to understand utterances of (2) as communicating PRESENT.

This explanation also generalizes to third-person attitude ascriptions, given one further stipulation. Stipulate for a moment that attitude ascriptions somehow communicate that the matrix subject is disposed to presuppose presuppositions of the complement. With (6), for example, utterances of (6) somehow communicate that James is disposed to presuppose that John used to wear plaid.

(6) James is glad that John has stopped wearing plaid.

And in presupposing this proposition, James is disposed to treat it as not the main point of the utterance. As before, there are three interpretations of what uses of (6) communicate: CONJUNCTIVE, PRESENT, and PAST. CONJUNCTIVE is out, as less informative. Since James is disposed to treat propositions about John's past as not the main point, PAST is a less good interpretation of (6) than PRESENT.

My overarching goal in this paper is to shift the working hypothesis about presupposition triggers under attitude reports. The working hypothesis has traditionally been that the Atom theory is true. I think that the working hypothesis should rather be that the Fission theory is true. But it's hard for a single paper to shift the working hypothesis while also zeroing in on the very best version of the working hypothesis. One reason that it's hard to zero in on the best version of the Fission theory is that all versions of the Fission theory will face some important problems. For example, I just introduced an explanation that stipulates something important: that we normally interpret the matrix subject as being disposed to presuppose presuppositions of the complement. And this stipulation looks like a liability of this first explanation.

After all, the Atom theorist might claim to *explain* the facts that I've stipulated. For an Atom theorist, being glad that John stopped wearing plaid is being glad about a conjunctive proposition about the past and present both. And we should agree that being glad that p and q requires

believing p and q. And believing that John used to wear plaid is a great reason for being disposed to presuppose that he used to. Advantage Atom theory: it explains an important fact that my proof-of-concept stipulated.

This apparent advantage is illusory. The Atom explanation of this pattern doesn't work with the full generality that it needs. The pattern is highly general.

(GENERAL PATTERN) In order to V that S, you normally need to be disposed to presuppose S's presuppositions.

For example, *hoping* that John stopped wearing plaid also seems to require believing that John used to wear plaid. And the Atom account does not capture this general pattern. For one thing, it's not plausible that hoping that p and q requires believing p and q. (You can hope that it's cloudy and rainy, without believing that it's cloudy and rainy.) Extant work has tacitly stipulated that the General Pattern holds, without explaining why.²⁴ In other words: Atom theorists need an explanation of the General Pattern as much as Fission theorists do. So we can't use the General Pattern to discriminate between the Atom theory and the Fission theory.

There are other constraints on the correct way to develop Fission theories. For example, Fission theories also need to explain how presuppositions can be *locally accommodated* under attitude reports, as noted in §3. And this requirement does impose substantial constraints on the correct Fission theory, in a way that there isn't space to explore here.

5 Wrapping up

This paper has both an expository agenda and an argumentative agenda. The expository agenda is to cleanly articulate a fundamental divide between two kinds of pragmatic theories of presupposition: between Fission theories and Atom theories. Both are both broad tents that include lots of different concrete theories. But the choice between them is a deep and fundamental one. If you adopt any version of the Atom theory, you will automatically secure a particular range of advantages, and incur a particular range of costs. This same is true of the Fission theory – adopting any version of it automatically secures a particular range of advantages, and incurs a particular range of costs.

I've tried to describe the choice between the two theories at a high enough level to bring out the basic advantages and the basic costs. The

²⁴Paul Elbourne (2005), Bart Geurts (1998), and Irene Heim (1992) are some examples.

Atom theory is theoretically more ambitious. It tries to use a comparatively narrow toolkit to explain the full range of data. Its basic problem is empirical. You can't explain the full range of data just by appeal to that narrow toolkit. The Fission theory is theoretically less ambitious; it uses a comparatively broad toolkit. Its problem is theoretical. It needs to articulate the principles that determine what parts of the toolkit bear on which empirical phenomenon.

My argumentative agenda has been to eliminate the Atom theory. I've argued that it's *impossible* to make good on its theoretical ambitions. In particular, I've argued that it cannot explain a key generalization:

(HIGH-LEVEL GENERALIZATION) If a sentence 'S₁' contains a presupposition triggers, it's associated with a commitment T₁ that you have to V in order to V that S₁.

I take this generalization to force pragmatic theorists to be Fission theorists. (This GENERALIZATION also constrains the way we develop semantic theories of presupposition. But semantic theories of presupposition already tend to vindicate the HIGH-LEVEL GENERALIZATION, so accepting it has less dramatic upshots for semantic theories.)

Some theorists are already developing what I would take to be Fission theories; I think, for example, that it's helpful to think of Dorit Abusch (2010) and Márta Abrusán (2011, 2016) as developing Fission theories. I've in effect been arguing that they *underestimate* the importance of what they are doing. Such approaches are the only viable ways forward for pragmatic theorists. So the goal of this paper is to give us a more systematic perspective on presupposition triggers, rather than to develop a constructive explanation that competes with the sorts of theories that Abusch and Abrusán have developed.

My GENERALIZATION is also important because it has a wide range of important *philosophical* upshots, even outside of foundational questions about presupposition. For one thing, it requires pragmatic theorists to recognize just the sort of *pragmatic modulation* that Jason Stanley (2007) finds objectionable. More broadly, though, it has important philosophical upshots because a wide range of philosophically interesting constructions are presupposition triggers. One concrete example of this upshot is in metaethics. Moral realists take moral properties to be independent of and more fundamental than our individual evaluative attitudes. Wouter Kalf (2018), Alex Silk (2016, 2017), and Caleb Perl and Mark Schroeder (2018) have recently argued that commitment to those properties is a *presupposition* of moral discourse if it's present at all. Now "knows" is one kind

of attitude verb, and so falls within the scope of the HIGH-LEVEL GENERALIZATION.²⁵ If the GENERALIZATION is right, knowledge reports don't require knowledge of presuppositions triggered from the complement; they would only require accepting it.

So if the metaphysically realist part of moral discourse is part of a presupposition, moral knowledge could be much easier to acquire. It would only require *accepting* the metaphysically realist part of moral discourse, not knowing it. And the norms on acceptance are plausibly much less demanding than the norms on knowledge.²⁶ So if the HIGH-LEVEL GENERALIZATION is true, a presuppositional conception of moral discourse would have sweeping consequences for moral epistemology. I've introduced some of those sweeping consequences in Perl (2019). More generally, the epistemology of any subject-matter associated with substantive presuppositions would look very different given the HIGH-LEVEL GENERALIZATION. Knowledge in that domain would be much easier to acquire than we would initially assume. So even though I've been focusing very narrowly on a particular point about presupposition, the point under discussion promises to have significant philosophical upshots.

Acknowledgments: For extremely helpful comments on this material, I'm grateful to Steve Finlay, John Hawthorne, Barry Schein, Ralph Wedgwood, Jon Wright and the referees and editors for the *Australasian Journal of Philosophy*. I'm especially grateful to Mark Schroeder for comments and helpful guidance, and to Scott Soames for posing the question that this paper tries to illuminate and for vigorous discussions of the possible answers.

²⁵As evidence that it does, contrast these two examples:

(11a) I don't know whether John took out the trash. But I'll accept for now that he didn't. * I definitely know that it wasn't John who took it out. The trash is empty, so someone took out out.

(11b) I don't know whether someone took out the trash. But I'll accept for now that someone did. I definitely know that it wasn't John who took it out. I've been watching him all day, and he went nowhere near it.

²⁶I'm here assuming that acceptance and not mere belief is the norm on presupposition; for discussion, see Stalnaker (2002).

References

- Abrusán, Márta. 2011. “Predicting the presuppositions of soft triggers.” *Linguistics and Philosophy* 34 (6):491–535.
- . 2016. “Presupposition Cancellation: Explaining the Soft-Hard Distinction.” *Natural Language Semantics* 24 (2):165–202.
- Abusch, Dorit. 2010. “Presupposition Triggering from Alternatives.” *Journal of Semantics* 27 (1):37–80.
- Anand Pranav and Hacquard, Valentine. 2013. “Epistemics and Attitudes.” *Semantics & Pragmatics* 6 (8):1–59.
- Bach, Kent. 1999. “The Myth of Conventional Implicature.” *Linguistics and Philosophy* 22 (4):327–366.
- Beaver, David. 2001. *Presupposition and Assertion in Dynamic Semantics*. Stanford, CA: CSLI Publications.
- Chierchia, Gennaro and McConnell-Ginet, Sally. 2000. *Meaning and Grammar: An Introduction to Semantics*. Cambridge, MA: The MIT Press, second edition.
- Elbourne, Paul. 2005. *Situations and Individuals*. Cambridge, MA: The MIT Press.
- Geurts, Bart. 1998. “Presuppositions and Anaphors in Attitude Contexts.” *Linguistics and Philosophy* 21 (6):545.
- . 1999. *Presuppositions and Pronouns*. Oxford: Elsevier.
- Glanzberg, Michael and Siegel, Susanna. 2006. “Presupposition and Policing in Complex Demonstratives.” *Nous* 40 (1):1–42.
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts.
- . 1983. “On the Projection Problem for Presuppositions.” In Daniel Flickinger Michael Barlow and Michael Westcoat (eds.), *Second Annual West Coast Conference on Formal Linguistics*, 114–126. Stanford, CA: CSLI Publications.
- . 1992. “Presupposition Projection and the Semantics of Attitude Verbs.” *Journal of Semantics* 9 (3):183–221.

- Kalf, Wouter. 2018. *Moral Error Theory*. Basingstoke: Palgrave Macmillan.
- López de Sa, Dan. 2008. “Presuppositions of Commonality: An Indexical Relativist Account of Disagreement.” In Garca-Carpintero and Kolbel (eds.), *Relative Truth*, 297–310. Oxford University Press.
- . 2015. “Expressing Disagreement: A Presuppositional Indexical Contextualist Relativist Account.” *Erkenntnis* 80 (1):153–165.
- Mandelkern, Matthew. 2017. *Coordination in conversation*. Ph.D. thesis, MIT.
- Perl, Caleb. 2017. “Shifty Contextualism about Epistemics.” *Ergo* 4 (28).
- . 2019. “Might moral epistemologists be asking the wrong questions?” *Philosophy and Phenomenological Research* .
- Perl, Caleb and Schroeder, Mark. 2018. “Attributing error without taking a stand.” *Philosophical Studies* .
- Potts, Christopher. 2005. *The Logic of Conventional Implicatures*. Oxford: Oxford University Press.
- Roberts, Craige. 2012. “Information Structure: Towards an integrated formal theory of pragmatics.” *Semantics and Pragmatics* 5 (6):1–69.
- Rooth, Mats. 1985. *Association with focus*. Ph.D. thesis, University of Massachusetts.
- Rothschild, Daniel. 2011. “Explaining presupposition projection.” *Semantics & Pragmatics* 4 (3):1–42.
- Schlenker, Philippe. 2010. “Presuppositions and Local Contexts.” *Mind* 119 (474):377–391.
- Silk, Alex. 2016. *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics*. New York: Oxford University Press.
- . 2017. “Normative Language in Context.” In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 12*. Oxford University Press.
- Simons, Mandy. 2001. “On the Conversational Basis of Some Presuppositions.” In *Semantics and Linguistic Theory* 11, 431–448.
- Simons, Mandy, Beaver, David, Roberts, Craige, and Tonhauser, Judith. 2017. “The Best Question: Explaining the Projection Behavior of Factives.” *Discourse Processes* 54 (3):187–206.

- Simons, Mandy, Tonhauser, Judith, Beaver, David, and Roberts, Craige. 2011. “What projects and why.” *Proceedings of Semantics and Linguistic Theory (SALT)* 22:309 – 27.
- Stalnaker, Robert. 1973. “Presuppositions.” *Journal of Philosophical Logic* 2 (4):447–457.
- . 1974. “Pragmatic Presuppositions.” In Milton K Munitz and Peter K Unger (eds.), *Semantics and Philosophy*, 197–213. New York: New York University Press.
- . 1984. *Inquiry*. Cambridge, MA: The MIT Press.
- . 1999. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford: Oxford University Press.
- . 2002. “Common Ground.” *Linguistics and Philosophy* 25 (5-6):701–721.
- Stanley, Jason. 2007. *Language in Context: Selected Essays*. Oxford: Oxford University Press.
- Tonhauser, Judith, Beaver, David, Roberts, Craige, and Simons, Mandy. 2013. “Towards a Taxonomy of Projective Content.” *Language* 89 (1):66–109.
- van der Sandt, Rob. 1992. “Presupposition projection as anaphora resolution.” *Journal of Semantics* 9 (4):333–377.
- Villalta, Elisabeth. 2000. “Spanish subjunctive clauses require ordered alternatives.” In *Proceeding of SALT 10*, 239–256.

Affiliation: University of Colorado at Boulder