

Communicating in Contextual Ignorance

Alex Davies (Department of Philosophy, University of Tartu, Estonia)

1 A conservative aim

The discernible context in which linguistic communication takes place typically underdetermines which proposition is literally expressed by a context-sensitive (declarative) sentence (in context) used to do the communication. A growing number of theorists believe this underdetermination to have exciting and unexpected implications for our understanding of linguistic communication. For instance, and roughly speaking, some think it shows that you cannot mean that a man is dying down at the local bar by saying “A man is dying down at the local bar”, while others think that in many circumstances where you wouldn’t expect it, if someone tells you that a man is dying down at the local bar by saying “A man is dying down at the local bar” then you cannot know that a man is dying down at the local bar because you don’t know that this is what that someone said. But it is not self-evident that underdetermination has any such special implications for the nature of linguistic communication: if it were self-evident, there would be no need to publish arguments defending the implications.

In this paper, I argue that four arguments used to defend some such implication are unsound. I think the supposed implications for linguistic communication of underdetermination are overblown and that we can see this by engaging in a careful examination of the arguments put forward for the implications. I’ll argue that each of the four arguments to be discussed elides from consideration the possibility that propositions have parts which are also propositions. This possibility (and I think actuality) makes space for partial understanding: i.e. parts of the unidentifiable proposition that a context-sensitive sentence expresses in context which can quite plausibly be identified on the basis of the available evidence in the context and whose identification typically suffices to achieve whatever one would hope to achieve through linguistic communication in said context. Once we recognize that there are such propositions, the unsoundness of the four aforementioned arguments will become apparent: or at least, this is the case I’ll make in what follows.¹

¹ This is a partner paper of Davies (forthcoming), where I defend an analysis of indirect speech reports according to which they can be used to report parts of the content of the sentence whose utterance is being reported. My hope is that what we get when we combine the arguments of the present paper with the arguments of its partner is a view on which verbs of indirect speech are built with a semantics that allows us to report what people say in contexts whose content-shaping details we don’t know exhaustively.

2 The standard story and underdetermination

Let's begin by getting in view a set of platitudes about linguistic communication. A speaker succeeds in communicating a proposition only if her audience understands that she means to communicate that proposition. Fortunately, the meaning of a declarative sentence is such that, when paired with a context, there's a unique proposition which the sentence literally expresses (or equivalently we may say there is a unique truth-condition which is the sentence's own). A speaker can use this feature of declarative sentences to communicate propositions to others who speak the language from which the sentences are taken. She can choose a sentence which, in the context she's in, literally expresses the proposition she wants to communicate. By uttering this sentence, she can make the proposition she wants to communicate plain to those who speak the same language. By uttering the sentence, a speaker may also mean (and perhaps succeed in communicating) other propositions which are not literally expressed by the sentence she uttered but which she nonetheless makes apparent by uttering the sentence (centrally: conversational implicatures). I set aside communication of these propositions in what follows: we're focusing on communication of the propositions which are literally expressed by sentences in context. The communication of these propositions is typically understood as their being *said* as opposed to implied, insinuated, suggested, etc. For example, consider:

- (1) Every beer is in the bucket.
- (2) Mbissine hasn't visited grandma yet.

The meaning of (1) doesn't fix which range of beers needs to be in which bucket for the sentence to be true. The meaning of (2) doesn't fix within which interval of time Mbissine has to have not visited grandma for the sentence to be true. However, when (1) and (2) are combined with a suitable context, what is left unsettled by the sentences in themselves becomes settled. What's said by uttering (1) and (2) in a suitable context is the proposition, in each case, that is literally expressed by the respective sentence when paired with the respective context. A speaker can get her audience to recognize that she means this proposition by uttering the sentence in context. Multiple speakers of the relevant language will agree that this proposition was expressed by the sentence in context. A hearer can come to know that this proposition is true by hearing the sentence in context.

I will call this set of platitudes *the standard story*.²

In recent papers, one or another aspect of the standard story has been attacked as false; and all on the basis of a common phenomenon, which we'll call *underdetermination*.

² One subset or other of these platitudes has been called 'an intuitively compelling and fairly standard picture of communication' (Abreu Zavaleta, 2019, p. 1), 'a simple account of communication' (Bowker, 2019, p. 1), 'a view implicit in much of twentieth-century philosophy of language' about linguistic communication (Buchanan, 2010, p. 340), and an assumption which is implicit in formulations of the conditions for testimonial knowledge or justification (Peet, 2016, p. 396).

Underdetermination is not just context-sensitivity: the fact that a sentence's meaning must be supplemented with context for the sentence to literally express a proposition. It is the thesis that the parts of the context which a typical interpreter of a sentence will be able to discern underdetermine which proposition the sentence expresses in that context: often there are multiple propositions that could be what the sentence expresses in a context for all that an interpreter can tell.

For example, suppose that (1) is used in the following context.

While preparing for their first party at their new off campus apartment, Chet and Tim go out to buy provisions for the night. After a long and heartfelt discussion, Chet convinces Tim that "sophisticated" partygoers, like the charming ladies next door, do not like to drink beer from a keg - 'especially if it is domestic, bro'. To cater to the sophisticates that they hope will show up later that night, they decide to go to a local corner store to pick up several cases of imported bottled beer which they will serve from a giant ice-filled plastic bucket, decorated in a pirate motif, which is to be located in their backyard. An hour before the party is to begin, Tim asks Chet 'Are we ready to rage?' 'So bro', Chet responds, 'We are totally ready. The living room totally looks like a pirate ship. The strobe lights are up. Every beer is in the bucket. I just need to find an eye patch to wear with this pirate hat.' (Buchanan, 2010, pp. 346-347)

It seems that, for all we can tell on the basis of the discernible context, the propositions expressed by each of the following sentences could be what was expressed by sentence (1) in the context just described:

- (1)
- a. Every beer we bought at the bodega is in the bucket in the backyard.
 - b. Every beer we will serve at the party is in the bucket decorated in pirate motif.
 - c. Every beer for our guests is in the bucket filled with ice.
 - d. Every beer at the apartment is in the bucket next to the hot tub.

But the discernible context doesn't indicate which of these propositions (amongst others) was literally expressed by (1) in that context. We have *no idea* which of these was expressed by the sentence in that context.

On the basis of underdetermination, the following conclusions have been defended. Abreu-Zavaleta ((2019), (2020)) argues that for a context-sensitive sentence in context, there is no proposition that two or more speakers agree is expressed by the sentence in context. Bowker (2019) argues that seemingly context-sensitive sentences, even when used in a context, don't express propositions (i.e. have truth-conditions) at all. If either Abreu-Zavaleta's or Bowker's argument is sound, then sentences in context cannot make plain any proposition the speaker means to express by uttering the sentence in that context. Buchanan (2010) argues that, for a wide range of context-sensitive sentences, speakers cannot mean any of the propositions that

the sentences could express by uttering those sentences. If that's right, then, for a wide range of context-sensitive sentences, linguistic communication understood in line with the standard story cannot happen. Peet ((2015), (2016)) argues that context-sensitivity undermines our capacity to gain testimonial knowledge in a wide range of circumstances that had hitherto not been considered dangerous to testimonial knowledge. If that were right, then although linguistic communication can happen, for a wide range of circumstances in which we thought we could get knowledge from testimony, we cannot.

When it proves useful to refer to these philosophers collectively, I'll refer to them as *the pessimists*. I'll address the arguments for these conclusions in sections 5-8. I first want to define content parthood and explain the role it can play when we communicate in contextual ignorance. I do things in this order so that readers can read the arguments and my explanation of why each is unsound, adjacently, rather than having to jump between separated parts of this paper.

3 States, state parthood, contents and content parthood

If we were to suppose that proposition P is a part of proposition Q if and only if Q entails P, then we would have to say, for instance, that part of the content of "The DPR is an unrecognized state" is the content of "*The Womanly Face of War* is an uncomfortable read or the DPR is an unrecognized state." Since it's not, an alternative supposition is preferable.

Classically the content of a sentence is understood as a function from worlds to truth-values: each world being a condition the world could be in, and each truth-value being the truth-value the sentence would have were the world in the corresponding condition. However, let's exchange these functions with functions from *states* to truth-values, where states are understood as parts of worlds. Let's conceive of states as follows. States are those fact-like things which are responsible for the truth or falsity of sentences *without remainder* (what Fine ((2017a), (2017)) calls *exact* verifiers and falsifiers). "Without remainder" means that they don't include any material that is irrelevant to the truth or falsity of the sentence in question. For example, "Kurakhove is a town in Ukraine" is a sentence that is made true by the sum of the state of Kurakhove being a town and the state of Kurakhove being in Ukraine. The obtaining of this sum suffices to make the sentence true. Other features of Kurakhove such as its population, its buildings, its distance from the front line, and so forth, are not constitutively relevant to whether the sentence "Kurakhove is a town in Ukraine" is true or false. The word "constitutively" here acknowledges the fact that if the front line engulfs Kurakhove and the Donetsk People's Republic consumes the town, this will cause the state of Kurakhove being in Ukraine to be extinguished. But what matters to the truth of the sentence is just whether that state obtains—never mind what causes it to obtain or to not obtain. Similarly, the state of Kurakhove not being a town and the state of Kurakhove not being in Ukraine each suffice to make the sentence "Kurakhove is a town in Ukraine" false. Further material is redundant.

State containment is not spatial containment. It is better to think of state parthood as follows. When you have two states, and you're wondering whether one is a part of the other, ask yourself whether what it is for one state to obtain is in part for the other to obtain. For instance, is what it is for Kurakhove to be a town in part for Kurakhove to be in Ukraine? No, it is not. But surely part of what it is for Kurakhove to be a town and in Ukraine is for Kurakhove to be in Ukraine. We'll use this heuristic to guide judgements about when one state includes another as a part.

With state parthood thus operationalized, let's turn to content parthood defined in terms of state parthood. Take the two sentences "Kurakhove is a town in Ukraine" and "Kurakhove is in Ukraine." What one says in uttering the latter seems intuitively to be a part of what one is saying in uttering the former. What can we observe about the relation between the verifiers and falsifiers of each sentence? Firstly, notice that "Kurakhove is a town in Ukraine" entails "Kurakhove is in Ukraine." So, every time we have a verifier of the former, we have a verifier of the latter. But we don't want to say that the verifiers of the former *are* verifiers of the latter—for the verifiers of the former include more material than is required to verify the latter. But, applying our heuristic for deciding when one state is a part of another, it seems that the verifiers of the former include verifiers of the latter. Secondly, notice that if "Kurakhove is in Ukraine" is true then "Kurakhove is a town in Ukraine" is at least partly true, even if not entirely true. This suggests that every verifier of the former is included as a part in a verifier of the latter. Thirdly, notice that if "Kurakhove is in Ukraine" is false then "Kurakhove is a town in Ukraine" is false and furthermore, that the falsity of "Kurakhove is a town in Ukraine" doesn't show that "Kurakhove is in Ukraine" is false. This suggests that every falsifier of the former contains a falsifier of the latter as a part.

From these observations, we arrive at a definition of content parthood modelled on Fine's ((2017a), (2017)).

Content Parthood

Propositional content P is part of propositional content Q iff:

- (a) Every verifier of Q contains a verifier of P.
- (b) Every verifier of P is contained in a verifier of Q.
- (c) Every falsifier of P contains a falsifier of Q.

To apply this definition to the examples we'll discuss in this paper, it'll help if we define a small formal language that we can use to represent the meanings of the sentences that appear in the examples. In order to apply this definition to a wider range of examples, one would need a more expressive formal language than I will provide. The absence of such a more expressive

formal language doesn't sink the aims of this paper. But I hereby acknowledge that they are contingent upon the statement of a more expressive language (for examples of such, see Fine's work).

Let a model M be a triple: $M = \langle D, S, I \rangle$. D is a set of individuals. We'll assume that the language is typed. e and t are types, and for all types σ and τ , $\langle \sigma, \tau \rangle$ and $\langle s, \sigma \rangle$ are types. For each type there corresponds a domain: D_e, D_t, \dots etc. S is a set of states closed under the sum operation. Worlds are largest states: for a world, every state is such that it is either a part of the world or incompatible with it. I is an interpretation function that assigns items to constants. Let's use the following notation. $\llbracket \alpha \rrbracket^{M,s,g}$ is shorthand for: the semantic value of α in model M , relative to situation s , under variable assignment g . We then adopt the following assumptions about the following formulas. The interpretation function assigns items from a domain of the right type to each constant. For instance, where α is an individual constant, $I(\alpha)$ is a function of type $\langle s, e \rangle$. Where α is a one place predicate, $I(\alpha)$ is a function of type $\langle s, \langle e, t \rangle \rangle$. Where α is a two-place predicate, $I(\alpha)$ is a function of type $\langle s, \langle e, \langle e, t \rangle \rangle \rangle$. Where α is a constant of whatever type, $\llbracket \alpha \rrbracket^{M,s,g} = I(\alpha)(s)$. Where α is a variable of whatever type, $\llbracket \alpha \rrbracket^{M,s,g} = g(\alpha)$. We'll consider only one non-atomic sentence viz. (1). For that it will be handy to have a specification of the verification and falsification conditions of a restricted universally quantified formula, which includes a definite description. Where α and β are of type t , and v is a variable of type e :

$$\llbracket \forall (v: \alpha) \beta \rrbracket^{M,g,s} =$$

1 iff s is a sum of states s''' such that for each k of D and state s' such that $\llbracket \alpha \rrbracket^{M,g[v \rightarrow k], s'} = 1$, there is a state s'' such that $\llbracket \beta \rrbracket^{M,g[v \rightarrow k], s''} = 1$ and s''' is the sum of s' and s'' .

0 iff s is a sum of two states s' and s'' such that for some k of D , $\llbracket \alpha \rrbracket^{M,g[v \rightarrow k], s'} = 1$ and $\llbracket \beta \rrbracket^{M,g[v \rightarrow k], s''} = 0$.

otherwise.

The universally quantified formula presupposes that something satisfies the restrictor, though this is not essential to our discussion. Where α is of type $\langle s, \langle e, t \rangle \rangle$:

$$\llbracket \iota v. \alpha(v) \rrbracket^{M,s,g} =$$

k iff $\{d: \llbracket \alpha(v) \rrbracket^{M,g[d \rightarrow v, s]} = \{k\}\}$

otherwise

We'll assume that an English sentence α is true relative to a world if and only if there exists a state which is a part of that world which α 's translation into the formal language maps to 1. We'll assume that an English sentence α is false relative to a world if and only if there exists a state which α 's translation into the formal language maps to 0. Otherwise α is neither true nor false.

4 How to communicate in contextual ignorance

For a sentence in a context, following Peet's (2016, p. 400) terminology, let an *epistemic candidate* be a proposition which the sentence in this context could express for all that can be discerned about the context (for instance: the propositions expressed by (1a)-(1d) are epistemic candidates for (1) in the context provided). If underdetermination is true of this sentence in this context, there'll be more than one epistemic candidate.

Epistemic candidates are not being identified using, what we might call, an *infallibilist* criterion: i.e. a proposition is an epistemic candidate for a sentence in context if and only if it is logically consistent with what can be discerned about the context that the proposition is the proposition the sentence in context expresses. As we'll see, this isn't a premise in any of the pessimists' arguments. What they rely upon is the observation that there is a range of propositions that the evidence provided by a context does absolutely nothing to distinguish between (not even on a fallibilist epistemology) *qua* candidates for the proposition expressed by the relevant sentence in context. It's the plausibility of *such* a premise that gives the pessimists' arguments their punch. Arguments whose weight is carried by excessively demanding epistemic standards will imply a general kind of scepticism that tells us nothing specifically about linguistic communication. And so, it's a fallibilist way of understanding an epistemic candidate on which we'll focus in what follows.

One thing one can do when met with a plurality of epistemic candidates (so understood) is identify a proposition which is a part of the proposition expressed, whichever epistemic candidate this proposition might have been. One does this by finding a proposition which is a part of every epistemic candidate. How does one go about identifying a proposition which is a part of every epistemic candidate? We can use the three conditions in the definition of content parthood as our guide. For ease of reference, let's call the proposition expressed the *total proposition* and let's call the proposition we're identifying the *proposition part*. Then, following condition (a), if the total proposition is one of the epistemic candidates, and if the proposition part is going to be a part of the total proposition, the proposition part should be such that every verifier of every epistemic candidate should have a verifier of this proposition part as a part. Moreover, we know that we want to limit what the verifiers and falsifiers of the proposition part can be: given condition (b), there can be no verifier of the proposition part which makes no appearance in a verifier of the total proposition. Given condition (c), there can be no falsifier of the proposition part which doesn't include a falsifier of the total proposition.

If we can identify a proposition that fits this brief, then we'll have a proposition which we can with high confidence believe to be a part of the total proposition even if we don't have any idea which of the epistemic candidates is the total proposition. In the following two subsections I'll illustrate how such a proposition can be identified with examples drawn from two of the pessimists: Buchanan and Abreu-Zavaleta.

4.1 Example: implicit domain restriction

Recall Buchanan's (1) and its context from section 2. I assume that we can translate (1) into our formal language as (3):

- (1) Every beer is in the bucket.
 (3) $\forall(x: \text{Be}(x))I(x, iy. \text{Bu}(y))$

Here, "Be" is a one-place predicate corresponding to "beer", "I" is a two-place predicate corresponding to "in", and "Bu" is a one place predicate corresponding to "bucket".

Given the context provided, are there any states we can reasonably believe to be verifiers of this formula? Given the semantics of restricted universal quantification, a verifier of this formula will be a state s such that for all x in D , s the sum of states s''' such that for all s' such that if $\llbracket \text{Be}(x) \rrbracket^{M, g, s'} = 1$ there is a state s'' such that $\llbracket I(x, iy. \text{Bu}(y)) \rrbracket^{M, g, s''} = 1$ and s''' is the sum of s' and s'' . Which objects can we be pretty confident satisfy the predicate Be? Those beers that are in the crates Chet and Tim bought from the local corner shop earlier on the night that (1) was uttered. Consequently, we can be pretty confident that any state in which there is an item that has these properties should be mapped to 1 by the denotation of Be; at least, for states in that world of evaluation. But as we consider different possible worlds in which different crates were bought, or the same crates were bought but which contain different bottles of beer, we can be pretty confident that the proposition expressed by (1) will track those beers that happen to be in the crates bought in that world. It is a claim about those beers that is being made. We can also be pretty confident that $iy. \text{Bu}(y)$ denotes the particular giant ice-filled plastic bucket, decorated in a pirate motif, which is located in their back yard. And it's reasonable to believe that had a different bucket been selected for the purpose, then (1) would be a claim about that other bucket. Thus, relative to alternative worlds in which a different bucket had been selected for this purpose, the states that verify (1) that are parts of such a world, would have to contain that different bucket. All this seems pretty certain given the discernible context. Given that we understand the denotations of "Be" and "Bu" in the ways just adumbrated, we can be pretty confident of the following: if each object x which, relative to a state s' satisfies $\text{Be}(x)$, also satisfies $I(x, iy. \text{Bu}(y))$ relative to another state s'' , then the sum of all the sums of such states s' and s'' is a part of a verifier of whatever proposition was expressed by (1).

Of course, there might be more inside a verifier of the total proposition expressed by the utterance of (1) in the given context than the states just described. In fact, if the epistemic

candidates for (1) in the described context are the propositions expressed by (1a)-(1d), then the verifiers of the proposition expressed by (1) in its context *must* contain more than what we just considered. For instance, suppose that the proposition expressed by (1) is the proposition which, in the context of Buchanan's paper, was expressed by (1a) "Every beer we bought at the bodega is in the bucket in the backyard." Any state which includes an object which is a beer bought from a bodega contains a state of that beer having been bought from a corner store because, to a first approximation, bodegas are corner stores that are in NYC. But since we have no idea whether the proposition expressed by (1a) (rather than any of the propositions expressed by (1b)-(1d)) is the total proposition expressed by (1) in its context, we have no idea whether the verifiers of the total proposition really do contain this extra material. But despite being clueless about *this*, given the description of the context, we can nonetheless be pretty confident that the verifiers of the total proposition contain as parts the states just described.

What about the *falsifiers* of (1) in its context? We know that a state that is a falsifier of the proposition part must, for each epistemic candidate, contain a state which falsifies that candidate. Can we identify such a state? We can be pretty confident that a state which includes an item that is one of the beers in the crates that were bought earlier in the evening but wherein that beer is not in the bucket selected for the purpose of the party will be sufficient to ensure the falsity of the total proposition expressed by (1) no matter which epistemic candidate it turns out to be. So, in identifying a proposition which is a part of the proposition expressed, we ensure that its falsifiers are such states.

What about (b), the condition requiring something of all verifiers of the part? To ensure that the proposition part satisfies this condition, we must ensure that there are no verifiers of the proposition part which are not parts of verifiers of the total proposition. To do this we must ensure that the states that verify the proposition part are always part of a state we can be pretty confident is part of what must be so if the total proposition is true. Since there are states that we can be pretty confident are parts of verifiers of the total proposition, we merely need to ensure that we never add in verifiers to the proposition part which are not parts of those states.

The result will be a proposition which we can be pretty confident is part of the proposition expressed, even if we have no idea which of the epistemic candidates was expressed by (1). And what's more, this proposition suffices to indicate Chet and Tim's state of readiness. Although there may be more to the proposition of Chet's sentence in context, this part of it is surely enough for their purposes — we haven't identified a useless part.

4.2 Example: gradable adjectives

Gradable adjectives are adjectives which admit intensifiers (e.g. "very", "really") and sometimes maximizers (e.g. "completely", "totally"). A now typical way to conceive of their (extensional) meanings is as functions from objects to degrees on a scale (cf. (Cresswell, 1977)

and (Kennedy, 1997)). For non-comparative gradable adjectives, it is assumed that context supplies a threshold point on this scale. An object has to be mapped to a degree which is higher than the threshold in order for the object to be mapped to true. For example, we could translate the sentence “Melania is tall” into our formal language as:

(4) $T(m)$

“ T ” is a unary predicate that corresponds to “is tall” and “ m ” denotes Melania. What is $\llbracket T(m) \rrbracket^{M, g, s}$? It’ll be a function from states to truth-values which makes implicit reference to a contextually supplied threshold: if s is a state in which Melania’s height is at least as high as the threshold, then the function maps s to 1; if s is a state in which Melania’s height is lower than this threshold, then the function maps s to 0.

As Abreu-Zavaleta ((2019), (2020)) has stressed, gradable adjectives easily lend themselves to demonstrations of underdetermination. For given how fine-grained the scales in question are presumed to be (real numbers), it seems that if the truth-condition of a sentence composed using a gradable adjective has a certain threshold, the discernible context in which the sentence is uttered will inevitably underdetermine exactly where this threshold falls on the scale.

However, ignorance of the details of the total proposition expressed in a context doesn’t mean ignorance of every part of that total proposition. Suppose some heterosexual men are discussing who Donnie could go on a date with. Someone suggests Melania. Aware of height-centred gender norms, one of them objects, “But Melania is tall.” Donnie’s crushed. They move on to other possibilities. Consider this utterance of “Melania is tall.” We cannot identify the exact threshold deployed by the total proposition expressed by this utterance. The discernible context doesn’t give us enough information to do so. However, for all pairs of propositions A and B which could be expressed by “Melania is tall”, and which differ only in where they locate the threshold on a scale of heights, if A sets the threshold higher than B then B is a part of A . Every verifier of A will contain a verifier of B : every state of Melania being at least as tall as A ’s threshold will contain a state of Melania being at least as tall as B ’s threshold. Every verifier of B will be contained in a verifier of A : every state of Melania being at least as tall as B ’s threshold will be a state that appears in some state of Melania being at least as tall as A ’s threshold. NB: not in the sense that whenever Melania is at least some smaller height she is at least some larger height. But rather in the sense that there will be states of being at least that larger height which contain the state of being at least that smaller height. Finally, every state that falsifies B will contain a state that falsifies A : every state of Melania not being at least as tall as B ’s threshold will contain as a part a state of Melania not being at least as tall as A ’s threshold.

A proposition identified in this way is a proposition everyone aware of the discernible context can be pretty confident was part of the total proposition expressed by “Melania is tall” even if none of us is in a position to know exactly which total proposition was expressed. We just need to identify a threshold we can be pretty confident is the same as or lower than whatever the threshold is for the total proposition expressed by “Melania is tall” in the given context.

Moreover, given the context, it'll be possible to identify a part of the total proposition expressed which is large enough to be informative. For it's clear from the context that Melania is being said to be at least taller than Donnie to an extent that would make him (a man who wishes to abide by certain gender norms) not want to go on a date with her. If Melania were not being said to be at least this tall, then it would be unclear why her height is being given as a reason to look for another date. For all we can tell, it's possible the speaker has in mind a proposition which requires Melania to be much taller still (a supererogatory proposition *vis a vis* the purposes of the exchange). But the part of that total proposition (whatever it is) that we can identify is all that's needed in the exchange.

Now, I haven't specified the threshold above which Melania must be at least being said to be taller than by using real numbers on an imperial or metric system of measurement. But this doesn't mean I haven't specified a degree on such a scale—just as I can point to a pencil mark on a wall to indicate a precise height in centimetres without being able to specify this height in centimetres *as such*: competence in using “tall” doesn't require acquaintance with the imperial or metric system. The description I have given of the threshold should be understood *de re*—since I am not proposing that the precise contents I am using to identify this threshold need be part of the proposition I am identifying.³ All I want to say here is that we can be confident that a proposition with this threshold (whatever it turns out to be when specified in centimetres or inches or whatever) is part of the total proposition the uttered sentence expresses in the specified context. The less informative the context, the lower the threshold of the largest proposition we can be confident is part of the total proposition expressed. But to the extent that it's implausible that the threshold is so low, the context provides sufficient information to conclude that a larger proposition is a part of the total proposition.

4.3 Summary

Even when we haven't any idea which of the epistemic candidates is the total proposition expressed by a sentence in context, the context can supply us with enough information to identify parts of that proposition. We've seen two examples. The content-parthood framework should make visible how, in principle, one can go beyond these examples. The trick is to identify the verifiers and falsifiers of a pair of propositions (in our case, expressed by the same

³ We could show Donnie a series of possible dates, with different heights, and identify, to a certain degree of precision, where the threshold is, specified as such in the imperial or metric system. The fact that we can only ever do this to a certain degree of precision means we cannot find the exact threshold—perhaps there isn't one. But regardless: we can always find the lower-bound of the threshold for the degree of precision attainable and be confident that a proposition with this threshold is a part of the total proposition expressed in the specified context. Given that the degree of precision is small enough, this proposition (the same as that described in the main text, but described in other terms) is very likely to be as informative as anyone in the context is going to need.

sentence) and then to use the Finean definition of content-parthood to figure out whether one is a part of the other. With other varieties of simple sentence (e.g. those involving multidimensional adjectives) this will be all that's required. For other varieties of sentence (e.g. those involving modal or attitude verbs, conditionals etc.) the rather inexpressive formal language provided here will need to be extended. The conclusions of this paper are contingent upon success in this endeavour.

I think the pessimists would even agree that such parts can be identified. None of them is suggesting that we're terrible at significantly narrowing down the range of epistemic candidates in which the total proposition might fall (see for instance (Bowker, 2019, p. 4250) and (Peet, 2016, p. 401)). But, so I'll now argue, in conceding what they do, they concede enough to cast doubt on their respective conclusions.

5 Abreu-Zavaleta – for most sentences, there's no proposition that two interpreters agree an utterance of that sentence to have expressed

Using the expression “(assertoric) utterance” to refer to sentences in context, Abreu-Zavaleta (2019, p. 3) aims to show that for most sentences in context, there is no proposition that ‘more than one language user believes to be that utterance’s truth-condition.’ Here’s the nub of his argument:

Nearly every assertoric utterance is such that, for any language user, there are enormously many extremely similar and equally eligible truth-conditional contents that language user could easily have believed the utterance to have. Given the large number of such extremely similar and equally eligible truth-conditional content candidates each of those utterances has, it would be extremely unlikely for there to be a proposition that more than one language user believes to be that utterance’s truth-conditional content. Thus, absent further explanation, there is no such proposition. (*ibid*)

Why should we think that for any assertoric utterance (i.e. declarative sentence in a context), ‘there are enormously many extremely similar and equally eligible truth-condition contents that language user could easily have believed the utterance to have’? If the discernible context underdetermines precisely which proposition a sentence expresses in a context, but nonetheless, interpreters take the sentence to have expressed some such proposition, then it seems very unlikely that two interpreters of the sentence will take it to express the same proposition. Abreu-Zavaleta illustrates the point with several examples. Here’s one of them (cf. (Abreu Zavaleta, 2019, p. 3):

Suppose Anna points to a certain box she has asked John to move and utters

(5) That box is heavy.

“Heavy” is a gradable adjective. There are many acceptable thresholds for being heavy compatible with the described context. Each corresponds to a different proposition (5) could be expressing in this context. Consequently, we don’t know exactly which total proposition is expressed by (5) in its context. There are many epistemic candidates.

However, although it might be unclear to John, given the discernible context, exactly which threshold Anna’s sentence in context requires the box’s weight to be above in order for it to be true, whatever it is, he can be reasonably sure that she’s warning him about its weight, given that he’s about to attempt to lift it. So she has to be saying that the box is at least heavy enough for her to think she should warn John about its weight: if the box were not detectably different in weight from the other boxes, then Anna’s utterance would be mysterious or obviously false. A proposition whose truth requires the box to be heavier than the threshold just identified (where the description just given is understood *de re*) is therefore part of whatever proposition her sentence in context expressed. So, although John cannot know what the *total* proposition is that Anna’s sentence in context expressed (there’s not enough information in the context for him to be reasonably confident about that—it might well be that Anna was saying the box is heavier than a much higher threshold—“No, I mean it’s really, really heavy.”), there’s no good reason to deny that he can be reasonably confident that the total proposition Anna’s sentence in context expressed at least had the part I’ve just described.

Abreu-Zavaleta discusses various ways of resisting his argument. Most of these don’t concern us: we’re not appealing to vagueness, naturalness or social externalism. But there is one that does. One might think that the *common ground* of a conversation is rich enough to enable different interpreters to identify the same proposition when trying to identify “the” proposition a sentence expresses in context. That might sound like the possibility I’m putting forward in this paper. But Abreu-Zavaleta rejects this possibility on the ground that each person’s understanding of the common ground will itself be slightly different. He tries to convince us of this by providing examples of inadequate common grounds and claiming these to be typical. However, by taking a look at one of Abreu-Zavaleta’s examples, we can see that, even if he’s right that the common ground doesn’t suffice for two persons to identify the same *total* proposition, there’s sufficient information to identify (informative) parts of whatever total proposition is expressed by the target sentence in context.

Suppose Anna and John are waiting for Carla, and see her walking towards them from afar. Looking at Carla, Anna starts a conversation with John by uttering:

(6) Carla is tall. I hadn’t noticed it before.

Abreu-Zavaleta’s contention is that Anna and John will diverge in the presuppositions they make about the height of various other objects. Abreu-Zavaleta (2019, p. 9) then reasons as follows:

...on the assumption that presuppositions about the standard of tallness [for judging whether Carla is tall] are arrived at on the basis of presuppositions about the heights of certain objects, it is also unlikely that Anna and John will make the same presuppositions about the standard of tallness relevant to their conversation.⁴

But none of this impacts the way of resisting Abreu-Zavaleta's reasoning presented above. Abreu-Zavaleta concludes that there are no propositions that two interpreters will take to be expressed by a sentence in context. For that to be false, it suffices to show that there are some propositions that two interpreters will, with great likelihood, take the sentence to express. Anna is expressing surprise at Carla's height. In all likelihood, Anna's sentence in context expressed a proposition which would be true if Carla were obviously, visibly taller than the average height for a human being of Carla's age, gender and locality. It may not be clear whether Anna's sentence in context expressed a proposition which is larger still. But if Carla were not at least this tall, Anna's sentence would be mysterious or clearly false. So this much must be part of whatever total proposition Anna's sentence in context expressed. If Anna's sentence expressed a proposition that contained this smaller proposition as a part, then it expressed this smaller proposition. But then, whatever other differences Anna and John may have, there *is* a proposition which Anna and John can agree Anna's sentence in context expressed.

One might say at this point: so what? So there are *some* propositions which two or more people are likely to agree are expressed by the relevant sentences in the relevant contexts. But the utility of shared belief about the propositions expressed by a sentence in context is that these allow information to be communicated by means of the utterance of such a sentence. Excessively small propositions are excessively uninformative. So Abreu-Zavaleta still shows this: that, for context-sensitive sentences, typically, the only propositions which can be identified as being expressed by the sentence in context are not informative enough for the purposes at hand.

But what does John need to glean from Anna's sentence in context in order for the sentence to be useful to him? Not that much. It's a warning that the box might be heavier than the others to a point where this might cause some trouble. If the box were not noticeably different from the others, then Anna's remark would be mysterious or false. So the proposition expressed by the sentence in this context must require for its truth that the box is at least that heavy. But John

⁴ Abreu-Zavaleta momentarily asks us to assume that neither Anna nor John know what each other think about the heights of objects outside their current fields of vision. But he abandons this assumption to make a more general statement, which I'm reporting above. If we don't abandon this assumption, the context discernible to Anna and John is reduced. This means they cannot identify as large a part of the total proposition Anna expressed as we can without the assumption. There still will be something. But it will be pretty uninformative. However, that seems to be a correct description of how people in such extreme circumstances would understand the utterance: as uninformative. The assumption is unusual and not reflective of the information ordinarily available to those communicating in a context.

really doesn't need to know more than that to make sure he doesn't, for instance, pull a muscle in his back by overzealous lifting of the box.

6 Bowker – seemingly context-sensitive sentences in context don't have truth-conditions

Bowker (2019) infers from underdetermination to the conclusion that those sentences which appear to be context-sensitive in fact don't literally express propositions at all (i.e. don't have truth-conditions).⁵ He reasons as follows. For a seemingly context-sensitive sentence, there will be multiple propositions such that, we don't know whether the sentence expresses one rather than another in a given context. But sentences in context cannot have unknowable semantic properties. From this he infers that there's no reason to accept that the sentence expresses any proposition.

However, for all that Bowker has said, it'll be possible to identify propositions which are parts of every epistemic candidate. So whatever proposition a sentence in context might express, it'll express these propositions. But then there are propositions which can (at least, given fallibilist criteria) be known to be expressed by the relevant sentence in context. Given this, if we apply the inference, that Bowker uses, from what's knowable to what's there, we arrive, not at the conclusion Bowker reaches (that seemingly context-sensitive sentences in context don't express propositions at all), but instead, that the only propositions they express are those propositions which are parts of every epistemic candidate.⁶ The largest such proposition will be the total proposition expressed by the sentence in context. In this respect, the implications of Bowker's considerations favour a metasemantics of context-sensitivity along the lines of King ((2013), (2014)): one that connects a context-sensitive sentence's content to what is recognizable in the context. Although this is not the only metasemantic theory of context-sensitivity in town, it is a serious contender. Its implications are not implausible.

There is a worry here. What if the proposition part we can identify is not itself an epistemic candidate? If, in such a case, we apply Bowker's reasoning, and conclude that this part is the total proposition expressed by the relevant sentence in context, then it seems we're committed to saying that the total proposition is a proposition which we have reason to believe is not the total proposition expressed by the sentence—because it's not an epistemic candidate. Are we committed to this result? Let's think for a moment. When *are* we confident that a proposition

⁵ I say "appear to be" or "seemingly" because Bowker aims to show that they are not in fact context-sensitive. Although that part of his discussion is not relevant to us here, I register this in order to acknowledge that Bowker is not trying to establish a conclusion that applies to all sentences.

⁶ Although I am very sympathetic to Bowker's incredulosity about unknowable semantic properties, I'm not certain whether to accept it. There do appear to be some semantic properties that are unknowable. No one will ever know what the referent of "The last words of Amelia Earhart" is. But that doesn't mean it doesn't have a referent. Perhaps a sentence in context can express a proposition that no one is ever in a position to completely grasp.

which was surely expressed by a sentence in context is itself definitely not the total proposition expressed thereby? These cases are not those in which, for all we can tell, the speaker's sentence in context expresses no more than the proposition we are capable of identifying: it won't be implausible to draw Bowker's inference in such cases and conclude that the part we could identify is the total proposition. Rather, the cases that come to mind are those in which you have reason to believe that there's evidence in the context, to which you're not privy, which will indicate as yet unidentified parts of the total proposition literally expressed by the relevant sentence in context. Perhaps you've had the experience of showing up late to a lecture. Although you understand the language of the lecturer (English, say), you are also aware that you're not grasping the full content of these sentences because you know you missed the beginning of the lecture by a lecturer who is held in high regard. You know you're only getting parts of the propositions expressed by the sentences the lecturer utters. You know this because you know there are stretches of the context of which you're uninformed yet which will probably indicate further parts of the total propositions the sentences you're hearing express. Compare the situation with another: you arrived early to the lecture and the lecturer has a reputation for being something of a charlatan. In that case, it'll be more reasonable for you to think you're getting everything there is to be got from what the lecturer is saying.

This is intriguing. Think back to the inference that Bowker draws from what is knowable to what is there. Such an inference isn't an inference from what is currently known to what is there. It works from what is in principle knowable. If however we know we're missing contextual information that we have reason to believe will show that the total proposition literally expressed by a sentence in context is bigger than we're capable of identifying, then it'll be knowable that the total proposition is bigger than the part we can identify. So the inference Bowker uses to restrict semantic content won't be licensed in this case.

It seems then that the cases in which it would be implausible to infer that the part of the total proposition we have identified is in fact the total proposition are those in which the inference isn't licensed. The inference Bowker uses to draw the conclusion that seemingly context-sensitive sentences in context don't express propositions can therefore be combined with the possibility identified in this paper (partial understanding) without fear that it'll force us to conclude that sentences in context express total propositions, when intuitively the sentences in context express total propositions which are larger than the part that can be identified.

7 Buchanan – speakers cannot mean a proposition literally expressed by a context-sensitive sentence in context

Buchanan (2010) notices the following. In order for a speaker to mean a given proposition by uttering a sentence, that proposition has to be such that, in order for a hearer to qualify as understanding the utterance, she must entertain that proposition. This much seems to follow from fairly standard Gricean assumptions about speaker meaning—none of which do I want to

dispute. But if we grant this general point, then underdetermination would seem to pose a challenge to the possibility of a speaker meaning a proposition by her utterance of a sentence. Let's restrict attention to propositions that a sentence in a context can literally express (i.e. let's continue to set aside things like conversational implicatures). Buchanan argues quite plausibly that typically when a context-sensitive sentence is uttered, there will be a set of propositions that a sentence can literally express with at least two members such that the entertaining of each would qualify the hearer as understanding the sentence in context. But if that's the case, then there's no proposition that the sentence can literally express in context such that entertaining that proposition is *necessary* for qualifying as understanding the sentence in context. Hence, given the general principle, the speaker could not have meant any of these propositions. From this it follows that, when you utter a sentence in context, you cannot mean any of the propositions that the sentence in context could literally express.

It should be clear by now how I think this line of reasoning goes awry. The propositions the entertaining of each of which suffices to qualify as understanding the relevant utterance will have a common part which is itself a proposition. But to entertain a proposition you have to entertain its parts. I don't mean this in the sense that in order to entertain a proposition you must go through a thought pattern of momentarily entertaining each of its parts. I mean that in successfully entertaining a proposition you entertain its parts, just as to pour the glass of water into the bowl is in part to pour a part of the glass of water into the bowl, even if you don't pour this part separately from pouring the entire glass. In failing to entertain a part of a proposition you fail to entertain the whole. But then a proposition which is a part of every proposition the entertaining of which is sufficient to understand the target sentence in context *is* such that, entertaining it is necessary for understanding the sentence in context. It doesn't suffice for entertaining those larger propositions. But it is such that, fail to entertain it, and you fail to entertain any of the larger propositions. Thus if there's a set of propositions each of which is such that entertainment of it is sufficient to understand a sentence in context, then the speaker of the sentence can mean a proposition by the sentence which is a part of all of the propositions in the set. For that proposition must be entertained in order for a person to qualify as understanding the sentence in context, no matter which total proposition the sentence expresses in context.

Think back to sentence (1). This is Buchanan's primary example for illustrating his line of reasoning. Sentences (1a)-(1d) are supposed to express propositions the entertaining of which would suffice for a hearer to qualify as understanding (1) in the specified context. But we have already seen that there is a content which is common to all of the propositions expressed by (1a)-(1d). Given that entertaining a proposition requires entertaining its parts, this content is a proposition such that if you don't entertain (at least) it then you don't qualify as understanding (1) in its context. This is a proposition which, despite Buchanan's principle, a speaker could mean by (1).

Buchanan doesn't consider anything like the possibility presented in this paper. I suspect this is both partly because definitions of content parthood have only become better known within philosophy well after the publication of Buchanan's paper, and partly because, in his 2010 paper, he happens to adopt a structured proposition approach to propositions: a framework which precludes application of the definition of content parthood.

8 Peet – underdetermination undermines the acquisition of knowledge from (context-sensitive) testimony in hitherto unexpected ways

We can acquire knowledge from someone having told us something. But seemingly, for that to happen, we must know what we have been told.⁷ Presumably, we've been told the content of the sentence uttered in its context. However, given underdetermination, it seems we commonly cannot identify the proposition that is the content of the sentence in context. Peet (2016) argues that because of this, the belief we end up forming on the basis of having been told something fails one or more plausible necessary conditions on knowledge: anti-luck, safety and sensitivity. Suppose there are several epistemic candidates: propositions the sentence uttered could have as its content, for all we can tell from the discernible context. Suppose furthermore that we interpret the sentence as having expressed one or other of these propositions and we go on to believe whichever proposition we interpret the sentence as having expressed. If one or more of the epistemic candidates is false, then, even if the proposition the speaker meant to express with her utterance were true, and we believed just that proposition, it would be a matter of luck that we believe a true proposition. Thus, an anti-luck condition on knowledge is violated. Now suppose that all of the epistemic candidates are true. Even so, the belief one forms is unlikely to be safe or sensitive. If one believes a proposition other than the one the speaker meant to express, in nearby possible worlds in which one formed the same belief in the same way it would be false (because the speaker wasn't checking to make sure that *that* proposition is true). Likewise, in nearby worlds in which the belief is false, one would still believe it, were one to form the belief in the same way. So, the belief formed would be likely to be neither safe nor sensitive.

But all of this is premised on the assumption that, when met with underdetermination, interpreters of a sentence will form a belief with the content of one or other of the epistemic candidates, despite having no evidence that it, rather than any other candidate, was the content of the sentence uttered. However, there is an alternative. No matter which (total) proposition was expressed by the utterance of a sentence in context, there'll be some parts of this proposition (whichever it is), which were expressed. The speaker is in a position to recognize these parts with high confidence, as is the hearer. But then, provided the speaker has ensured that these discernible parts are true, and provided the hearer forms beliefs with only these parts

⁷ Though see (Peet, 2018a).

as contents on the basis of the witnessed speech, underdetermination need pose no special problem for the acquisition of knowledge from testimony. Beliefs formed in this way won't be true by luck, and they will be safe and sensitive – for they will be beliefs whose truth is entailed by whatever proposition's truth the speaker has worked to ensure.

Peet (2016, pp. 408-413) does discuss a possibility which resembles the one considered here. He considers the possibility that interpreters of sentences only consider what Peet calls “coarse grained propositions” which are propositions that are true when any of the epistemic candidates are true and thus would include the constructed propositions (i.e. parts) we've been considering. Peet rejects this as a plausible explanation of what interpreters of a sentence do because he thinks that, by showcasing an example, he can show that interpreters commonly interpret sentences as having contents which are not entailed by each epistemic candidate. However, as discussed at length by Davies (2019), the example doesn't do its job. To make it more plausible that a hearer would believe a proposition that is finer grained than a proposition that is entailed by every epistemic candidate, Peet unwittingly updates the context with new information. But in doing so, he changes what the epistemic candidates are so that the proposition believed is entailed by each. So, the example doesn't show that a hearer plausibly believes a proposition not entailed by each epistemic candidate. Moreover, it seems likely that the same problem is going to arise for any attempt to provide a compelling example in which a rational interpreter will form a belief that is larger than the discernible context provides her reason to believe was expressed by the uttered sentence. For in order to make it clearly the case that a rational interpreter will form a belief with given content, we would describe a circumstance in which she has evidence to believe that the relevant sentence has at least this content. But insofar as we do that, we change the epistemic candidates in such a way that this content is plausibly a part of each epistemic candidate.

Quite generally, when studying doxastic behaviour, if we imagine a rational thinker presented with evidence that supports P but not Q, where Q is stronger than P, we typically don't suppose the thinker to, by default, believe that Q rather than just P. We don't suppose, for instance, that when it's a foggy day, and you see a figure in the fog, but you cannot make out who it is, that the person you see is Benedict Cumberbatch. We shouldn't assume belief formation in response to linguistic stimuli to be any different. Of course, we are all irrational some of the time. There will be circumstances in which irrational interpretation leads us astray (again, as Peet (2017), (2018)) himself explores). But there's no reason to believe that we become irrational simply because the total proposition expressed by a sentence in context is underdetermined by the available evidence.

The considerations Peet draws upon then don't imply that, when met with underdetermination, we cannot acquire knowledge from testimony. What they do imply is that often the knowledge we can acquire is limited to that part of the content expressed which can be discerned from the context. But surely, often enough, this suffices for us to do whatever we would want to do with

the content expressed (see (Pollock, 2020) for further discussion of the utility of partial linguistic understanding).

9 Clouds of contexts, underdetermination and proposition parts

I've now presented my reasons for thinking that four arguments, each with a different conclusion, but each sharing a premise (underdetermination), are unsound. In this closing section, I do some clarificatory work. I want to distinguish the phenomenon upon which we've been focusing from another, which some (e.g. (MacFarlane, 2020, pp. 56-57)) treat as interchangeable with the first. Both can be described in the language of "underdetermination." But, at least on the analysis offered here, they are different.

I have been assuming throughout that the utterance of a declarative sentence in context is being done to communicate the total proposition that the sentence, in that context, expresses (and thus, that it expresses such a thing in such a place). But instead of using a declarative sentence as a vehicle for communicating a proposition to someone, one can offer the sentence as, so to speak, a canvass for someone to draw their own proposition upon. Think for instance of Kieran Setiya asking philosophers on his podcast *Five Questions* five questions based on remarks by Iris Murdoch. When he asks these questions, he doesn't fully determine their content in advance. He allows his guests to define the content of the questions posed. He does this (presumably) because it's interesting: we learn something about the philosophers through how they nail down the content of the questions. The same thing can be done by uttering declarative rather than interrogative sentences: you let the recipient fill-out the content of the sentence rather than fixing the content in advance and using the sentence to communicate that content.

Von Fintel and Gillies (2011) describe this sort of use of language by appealing to the notion of a cloud of permissible contexts of interpretation: the speaker (and context) permit(s) her sentence to be interpreted relative to any of a range of different contexts.

Here then are two kinds of underdetermination: one in which there is one total proposition which a sentence in context literally expresses but an interpreter of the sentence doesn't have enough information to identify—the discernible context epistemically underdetermines the total proposition the sentence expresses; and another in which there is no single total proposition which the sentence in context expresses because the sentence isn't being used to communicate a particular (total) proposition but is rather being offered as a, so to speak, propositional canvass—the context itself (rather than the discernible context) metaphysically underdetermines which proposition the sentence expresses.

These are not the same thing: just because you cannot figure out what someone is saying doesn't mean they're letting you decide what it is they're saying. These two kinds of underdetermination can coexist though. For instance, we might imagine that Setiya only partially understands an answer given to one of his five questions: he would have performed

the kind of act described by von Stechow and Gillies in his question and then been able to identify but a proper part of the total proposition his guest's sentence expressed in response. This situation would be distinct from one in which Setiya's guest uses a declarative sentence to reciprocate the act Setiya performed with his question—a situation in which neither Setiya nor his guest is willing to nail down the content of their respective linguistic vehicles.

Acknowledgements

I'd like to thank the two referees for this journal who commented on this paper for helping me to improve the paper markedly beyond the initial submission. Their comments and objections were on point and have been a great help in improving the thing. The research that led to this paper was supported by the European Union's Regional Development Fund through the Centre of Excellence in Estonian Studies.

References

- Abreu Zavaleta, M. (2019). Communication and Variance. *Topoi*.
- Abreu-Zavaleta, M. (2020). Disagreement Lost. *Synthese*.
- Bowker, M. (2019). Ineliminable underdetermination and context-shifting arguments.
- Bowker, M. (2019). Saying a Bundle: Meaning, Intention and Underdetermination. *Synthese*, 196, 4229-4252.
- Buchanan, R. (2010). A puzzle about Meaning and Communication. *Nous*, 44(2), 340-371.
- Cresswell, M. (1977). The semantics of degree. In E. Partee (Ed.), *Montague Grammar* (pp. 261-292). New York: Academic Press.
- Davies, A. (2019). Testimonial Knowledge and Context-Sensitivity: a New Diagnosis of the Threat. *Acta Analytica*, 34(1), 53-69.
- Fine, K. (2017). Truthmaker Semantics. In B. Hale, C. Wright, & A. Miller (Eds.), *A Companion to the Philosophy of Language* (pp. 556-577). New York: Wiley Blackwell.
- Fine, K. (2017a). A Theory of Truth-Conditional Content I: Conjunction, Disjunctio and Negation. *Journal of Philosophical Logic*, 46, 625-674.
- Kennedy, C. (1997). *Projecting the adjective: The syntax and semantics of gradability and comparison*. New York: Garland.

- King, J. (2013). Supplementives, the coordination account and conflicting intentions. *Philosophical Perspectives*, 27, 288-311.
- King, J. (2014). The metasemantics of contextual sensitivity. In *Metasemantics: new essays on the foundations of meaning* (pp. 97-118). Oxford: OUP.
- MacFarlane, J. (2020). Lecture II: Seeing Through the Clouds. *Journal of Philosophy*, 117, 617-642.
- Peet, A. (2015). Testimony, pragmatics, and plausible deniability. *Episteme*, 12(1), 29-51.
- Peet, A. (2016). Testimony and the epistemic uncertainty of interpretation. *Philosophical Studies*, 173(2), 395-416.
- Peet, A. (2017). Epistemic Injustice in Utterance Interpretation. *Synthese*, 194(9), 3421-3443.
- Peet, A. (2018). Etiology, Understanding, and Testimonial Belief. *Synthese*, 195(4), 1547-1567.
- Peet, A. (2018a). Testimonial Knowledge Without Knowledge of what is Said. *Pacific Philosophical Quarterly*, 99(1), 65-81.
- Pollock, J. (2020). Context and Communicative Success. In T. Ciecierski, & P. Grabarczyk (Eds.), *The Architecture of Context and Context-Sensitivity*. Springer.
- von Fintel, K., & Gillies, A. (2011). 'Might' Made Right. In A. Egan, & B. Weatherson (Eds.), *Epistemic Modality* (pp. 108-130). Oxford: Oxford University Press.