

Knowledge and Cancelability

Tammo Lossau¹

ABSTRACT: Keith DeRose and Stewart Cohen object to the fallibilist strand of pragmatic invariantism regarding knowledge ascriptions that it is committed to non-cancelable pragmatic implications. I show that this objection points us to an asymmetry about which aspects of the conveyed content of knowledge ascriptions can be canceled: we can cancel those aspects that ascribe a lesser epistemic standing to the subject but not those that ascribe a better or perfect epistemic standing. This situation supports the infallibilist strand of pragmatic invariantism according to which knowledge semantically requires absolute certainty but this claim is often pragmatically weakened: it turns out that exactly those aspects of the conveyed content are cancelable that this view claims are pragmatic. I also argue that attributor contextualism and relativism do not have an alternative explanation of this phenomenon.

KEYWORDS: Knowledge ascriptions; Pragmatic invariantism; Cancelability test

A major point of discussions in the recent debate about knowledge ascriptions has been to what extent our intuitions about their felicity may be due to pragmatic implications. I take the main positions in the debate to be the following:

Attributor Contextualism: The semantic truth conditions of knowledge ascriptions depend on the context in which the knowledge ascription is uttered (the speaker's context) and these account for the vast majority of our intuitions.

Relativism: The semantic truth conditions of knowledge ascriptions depend on the context of assessment (the "intuiter's" context) and this explains our intuitions.

¹ This is a post-peer-review, pre-copyedit version of an article published in Synthese. The final authenticated version is available online at: <https://doi.org/10.1007/s11229-020-02663-7>

Fallibilist Pragmatic Invariantism (FPI): The semantic truth conditions of knowledge ascriptions are invariant and require a fallibilist standard of justification that is often met by us. In contexts with high standards (including skeptical contexts) *pragmatic strengthening* leads to intuitions that many knowledge ascriptions are false, although semantically they still are true.

Infallibilist Pragmatic Invariantism (IPI): Knowledge ascriptions semantically require absolute certainty to count as semantically true. In most contexts this requirement is *pragmatically weakened* which explains why we often intuitively take many knowledge ascriptions to be true.²

In this paper, I argue that data on cancelability supports IPI. In particular, there is an asymmetry in which aspects of the conveyed content can be canceled: we cannot cancel the implication that the subject (the “knower”) is in a position of high or absolute certainty, but we can cancel the implication that she is in an imperfect epistemic state. The former is a well-known problem for FPI due to Keith DeRose (1998) and Stewart Cohen (1999), but the latter, and in particular the arising asymmetry have not been discussed at great length. I argue this asymmetry is also problematic for contextualism and relativism insofar as they lack a convincing *explanation* of it. Their best attempt of such an explanation is to appeal to some form of a Lewisian “rule of attention”, but I will argue that this attempt fails.

Let me begin by revisiting the problem for FPI which will help to make the asymmetry clearer. Fallibilist pragmatic invariantists such as Jessica Brown (2006), Patrick Rysiew (2007), Geoff Pynn (2015), and Alexander Dinges (2019) claim that knowledge ascriptions are semantically true under conditions that require only a reasonably good epistemic standing of the subject (plus, of course, the truth of the embedded proposition). While this is in accordance with most everyday uses of knowledge

² FPI is often referred to as “moderate pragmatic invariantism” (and Pynn (2015) calls it “pragmatic contextualism”) whereas IPI is often called “skeptical pragmatic invariantism” or “strict pragmatic invariantism”. I use this terminology because I do not believe that IPI is committed to a serious version of skepticism.

ascriptions, in some contexts we seem to require more than that of a subject to count as knowing. In particular it seems that we don't allow anyone to know anything about the external world in "skeptical contexts", i.e. contexts like a philosophical discussion about Descartes's Meditations. Champions of FPI argue that this is because when used in such a context knowledge ascriptions are pragmatically strengthened so that they pragmatically imply that the subject is in a perfect epistemic standing – a standing that she could in fact not be in. The idea often is that if the speaker was taken to assert simply the semantic content of the knowledge ascription, she would be claiming something irrelevant, and therefore participants of the conversation interpret the utterance as a stronger claim.³ Thus FPI ends up making very similar claims as contextualists about the conveyed content but maintain invariable and achievable truth conditions at the semantic level.

A problem for this account is that the pragmatic strengthening FPI claims lacks a feature typically associated with pragmatic implication: they are not *explicitly cancelable* (DeRose 1998, sec. 10-11; Cohen 1999, 60). To explicitly cancel a pragmatic implication P is to add something like "but I do not mean to say that P" to a statement that carries a pragmatic implication that P without making the statement infelicitous. For example, I can say "There is a gas station around the corner, but I do not mean to say that they sell petrol", whereas it is infelicitous to say "There is a gas station around the corner, but I do not mean to say that it is around the corner." However, whenever we try to cancel what FPI claims is a pragmatic implication of a knowledge ascription with respect to the subject's epistemic standing, we give rise to a so-called "concessive knowledge attribution" (CKAs). For example, imagine a skeptical context in which I try to make the following statement: "I know that I have hands, but I do not mean to say that I can rule out all possible errors." Rather than canceling a pragmatic implication,

3 More recently, Mikkel Gerken (2017) and Alexander Dinges (2018) have also at least partly appealed to biases that leads us to overrate the probability or relevance of skeptical scenarios.

this utterance appears to be a contradiction.⁴ Even worse, this appears to be a feature of concessive knowledge attributions in general (the best motivation of this is in Lewis 1996). But if CKAs appear contradictory, this means they fail to directly cancel the pragmatic implication – rather the apparent contradiction will even call into question the semantic content. Thus FPI must claim pragmatic implications that are not explicitly cancelable.

An initial response to this problem might be to point out that CKAs can sometimes be felicitous. For example, we may say:

CKA1: “I know they’re going to lose, but I’m going to carry on watching just in case.”

CKA2: “I know the opera starts at 8, but I don’t mean to say that I can rule out being deceived by an evil demon.”⁵

Dylan Dodd (2010, 390-3) and Trent Dougherty and Patrick Rysiew (2009; 2011, 400-2) discuss examples like CKA1 at some length. Dodd plausibly argues that CKA1 amounts to asserting something along the lines of “There is no chance they are not going to lose, but there is a small chance they are not going to lose.” Yet CKA1 is felicitous, because the speaker is not committed to the truth of this statement – the first half of CKA1 is uttered as a hyperbole, or an expression of frustration. Similarly, CKA2 can be felicitous in a low-stakes context, because the speaker appears to be clarifying what

4 Brown (2006, 428) denies that the supposed pragmatic implications are not cancelable. She argues that it is felicitous to say “S knows that p, but her belief wouldn’t match the facts in a really distant possible world.” This seems right, but I would respond that once we make explicit that we mean an epistemically possible world and explain what this means we involve ourselves in the same apparent contradiction as above (although the utterance may still be felicitous). On the other hand, if we do not explain what we mean, we will actually fail to cancel the pragmatic implication in question. I should also flag that I think this talk of “distance” of possible worlds is problematic, especially as skeptical scenarios resemble the actual world as far as our experience is concerned.

5 CKA1 is originally from Hawthorne (2004, 21, fn. 60) and is credited to Tamar Gendler and Brian Weatherson there. CKA2 was suggested by an anonymous reviewer.

information she is trying to convey – namely, that they are sufficiently well-informed, but may not be prepared to defend their information against a radical skeptic. However, once we ask whether the speaker *really knows* that the opera starts at 8, it appears that she is asserting a similar type of contradiction: even in the low-stakes context, the speaker appears to be expressing that it is both impossible and possible that the opera does not start at 8.

Another attempt of objecting to the problem for FPI raised above might be that there plausibly are pragmatic implications that are not explicitly cancelable. But is this true? Paul Grice (1989, 44) thought that at least all *conversational implicatures* are explicitly cancelable. Even more, it may seem that *all pragmatic implications* (not just implicatures) are in fact cancelable as they are “defeasible inferences” just like conversational implicatures (Dimmock and Huvenes 2014, 3249). If this is so, the pragmatic implications claimed by FPI could not exist. However, some (Weiner 2006, Rett 2015, Åkerman 2015) have tried to provide cases of conversational implicatures that are *not* explicitly cancelable. These cases rely on the fact that the cancellation will be subject to pragmatic alteration due to its obvious falsity (which will often make the speaker sound ironical). The challenge here is to find a way of making the speaker’s claim credible so that the pragmatic effects of the cancellation vanish. This may only be possible given the right setup of the conversation – in such a case, the pragmatic implications arguably remain *contextually cancelable*, i.e. they would be cancelable given the right context, but they would not be explicitly cancelable in their original context (see Blome-Tillmann 2008, Åkerman 2015). The following modification of one of Matthew Weiner’s (2006) original cases by Michael Blome-Tillmann (2008, 159) illustrates this:

Unbeknownst to Alice, who is sprawling over several seats on the recreation deck, Sarah and one of her engineering officers are testing a portable tractor beam. For the purposes at issue, the tractor beam has to be strong enough to make it impossible for Alice to make room for someone else to sit down next to her. After activating the beam Sarah asks Alice via the intercom:

Sarah: Alice, I'm curious as to whether it would be physically possible for you to make room for someone else to sit down.

Alice: (baffled) Why should I? There's nobody else here who wants to sit down. [...]

Sarah: Oh, I'm sorry. I didn't mean to imply that you should make room. We are testing a new tractor beam on you and we are curious as to whether you can do it. This would give us an important indication as to how strong the beam really is.

Sarah's initial statement gives rise to the implicature that Sarah thinks that Alice should make room (and Sarah would be conveying this in an ironical way). In most contexts, any attempt to cancel this implicature would only add to the irony. However, in this specific context, Alice is able to cancel the implicature.

Weiner's cases may be challenged further, as they seem to be relying on the *practical* difficulty of making the attempted cancellation credible to the hearer. However, for the sake of the argument, let us assume that there are cases of conversational implicatures that are not explicitly cancelable in certain contexts. Even then, it seems that FPI has trouble explaining the oddness of CKAs. First, the kind of exceptions that have been pointed out in the debate surrounding the cancelability test are very different from the ones FPI needs to rely on. The problem in these cases is to make the cancellation seem genuine and credible, and this explains why the cancellation fails – but it is easy for a speaker to convince us that they cannot rule out all possible errors without sounding ironical, so CKAs seem very different from those cases (cf. Dimmock and Huvenes 2014, 3250). Second, it seems that CKAs *always* appear contradictory. But given our discussion above, if FPI is correct, there should at least be some contexts in which they appear consistent. Third, advocates of FPI face the problem of explaining the contradictory appearance of CKAs in general, combined with the fact that strengthening the conveyed content of a knowledge ascription by way of claiming the stronger claim to be an implication *is* felicitous – I can say “S knows that P, *and by this I mean* that she can rule out all possible errors.” The problem of explaining this asymmetry will be discussed at greater length below.

But first, let us get a better view of that asymmetry itself. To do so, let us look at why this problem does not apply to IPI. IPI traces back to ideas of Peter Unger (1975) and Jonathan Schaffer (2004) and is championed by Herman Cappelen (2005), Earl Conee (2005), Wayne Davis (2007) and Igor Douven (2007). The claim is that knowledge ascriptions semantically express that the subject is in an ideal epistemic state with respect to the embedded proposition, i.e. she can rule out all possibilities of error. Of course, this is rarely the case, so outside of skeptical contexts this statement will be pragmatically weakened so that it expresses something fit to the purposes of the conversation. For example, if the question is whether we have to go to the bank today, my saying “I know that the bank is open tomorrow” expresses that I am certain enough about this to base our decision on it, but not that I can rule out deception by an evil demon (which I obviously cannot). Because of how rare skeptical contexts are, this pragmatic weakening occurs by default; these are what Grice would call *generalized conversational implicatures*. The implicatures in play here can be canceled, for I can say: “I know that the bank is open tomorrow, and by this I mean that I am absolutely certain – there is not even a remote possibility of error.” While such a statement is presumptuous, it does not seem to involve any kind of contradiction as CKAs do. We can also add less presumptuous statements of the form “... and by this I mean that I can even rule out that P” which do not itself make a claim of infallibility but do lead to a stronger statement being conveyed.⁶

6 One feature of most generalized conversational implicatures is that we can cancel them in an attempt to appear witty. For example, we can say humorously: “Jones has three children. In fact, she even has four of them.” This does not seem to be the case with the implicatures discussed here. The reason for this appears to be related to the fact that the pragmatic effects in question represent a weakening of the semantic meaning (whereas in the children example, the semantic meaning is strengthened, so that the original statement implicated that Jones has exactly three children). Given we are dealing with weakening, we cannot “fool” the listener into thinking something that we did not actually semantically entail. An apt comparison to the implicatures we are dealing with are domain restrictions. For example, we can say: “Smith always brushes his teeth. And by this I mean that he brushes his teeth all day and night.” This may be

The situation then is this: given any knowledge ascription, the speaker can add a statement that strengthens the conveyed content up until the point of ascribing infallibility with respect to the embedded proposition to the subject. However, the converse is not the case: it is infelicitous to add a statement to a knowledge ascription that is intended to weaken the conveyed content so that a lesser epistemic standing would be required of the subject to confirm the conveyed content. And this is why any pragmatic strengthening of knowledge ascriptions with respect to the subject's epistemic standing which we could claim would not be cancelable⁷ whereas any pragmatic weakening of them would be. So far we have seen that this is a problem for FPI that does not apply to IPI. But I suggest that we ask a further question: what *explains* this feature of knowledge ascriptions? *Why* can we add claims to knowledge ascriptions that strengthen the conveyed content but cannot do the same with claims that are intended to weaken the conveyed content? IPI has an obvious explanation of this: it is because the former cancel pragmatic implications (which are always, or almost always, cancelable) whereas the latter are attempts to cancel semantic implications while also drawing attention to the semantic meaning, which gives rise to a contradiction. It is also clear that other theories cannot make use of this explanation as they deny that this is the case. FPI is furthest from the explanation that IPI offers as it involves the claim that pragmatic and semantic implications are often divided up conversely to how IPI claims they are – the more so the weaker the claimed semantic meaning of knowledge ascription is.

funny, but only in the same way that it is funny to say “I know that the Queen is British. And by this I mean that I can rule out that my brain is being manipulated by aliens to think that the Queen is British.” (Thanks to an anonymous reviewer for raising this point.)

7 There are uncontroversial cases of pragmatic strengthening of knowledge ascriptions that do not apply to the subject's epistemic standing. For example, if I am asked whether I locked my door I could reply “I know that there is a lot of crime in my neighborhood”, giving rise to a conversational implicature that I did indeed lock my door.

But what about attributor contextualism and relativism? Can they offer alternative explanations of this asymmetry? Contextualists such as Keith DeRose (1998), Stewart Cohen (1999), David Lewis (1996) and more recently Michael Blome-Tillmann (2014), Zoltán Gendler Szabó and Jonathan Schaffer (2014) and Jonathan Ichikawa (2017) think that the semantic truth conditions vary with the *context of attribution* and that these determine the epistemic standing that is ascribed to the subject by the knowledge ascription. Relativists such as John MacFarlane (2005) make a similar claim except that they think that it is the *context of assessment* which has this effect. Pragmatic effects on both views do not, or only to a less significant degree, alter which epistemic standing is ascribed to the subject by the conveyed content. If the ascription of an epistemic standing is conveyed in virtue of the semantic meaning, this means that the semantic implications of knowledge ascriptions are cancelable, but only in one direction. That semantic implications can be cancelable is a controversial claim in itself, but I will here assume that this can be the case (see Davies 2017 for some discussion). Contextualists and relativists thus need to identify a feature of knowledge ascriptions that accounts for the cancelability of those implications that involve the claim that the subject is in a weaker or less-than-perfect epistemic state which does not lead to the prediction that implications of the subject being in a stronger or perfect epistemic state should also be cancelable.

The best candidate for such a feature appears to be some version of David Lewis's (1996, see Gerken 2017, 24-30 for a discussion of relevant evidence) "rule of attention". Lewis champions a relevant alternatives theory according to which a subject needs to be able to rule out all relevant possibilities that entail non-P in order to count as knowing that P. He spells out the notion of relevance by giving a tentative list of rules including the "rule of attention". According to this rule, a possibility is always relevant if speakers in the context of attribution are currently attending to it. Note that this rule has a sister rule applicable to relativism according to which attention by the subjects in the context of assessment leads to relevance. Lewis points out that this rule explains the infelicity of concessive

knowledge attributions: by mentioning the fact that there are uneliminated possibilities, or by even pointing out certain possibilities, we draw attention to possibilities and make these possibilities relevant. This, in turn, leads to the knowledge ascription actually being false, even though it might have been true had we not drawn attention to these possibilities.

Similarly the contextualist or relativist can claim that attempts to cancel a semantic implication of a subject being in a perfect or very good epistemic standing defy their own purpose by drawing attention (in the context of attribution or in the context of assessment) to possibilities the subject is (or may be) unable to eliminate. Such attempts will only make the statement more clearly false as it guarantees that these possibilities are being attended to. However, there is no such problem in canceling semantic implications of the subject being in a comparatively weak or less-than-perfect epistemic standing. These are guaranteed to work precisely because they point to potentially uneliminated possibilities of error making them immediately relevant.

The “rule of attention” thus would be precisely the kind of feature contextualists and relativists are looking for. However, it is implausible that knowledge ascriptions really have this feature. As, among others, Michael Williams (2004) and Michael Blome-Tillmann (2014, 19-20) observe, the rule of attention makes it too easy to raise the standards of knowledge. Blome-Tillmann considers an example where one sees one’s teenage son sneaking out of the house at night and tells him the next morning “I know you left the house yesterday.” The son then objects that one may also just have dreamed this. If the rule of attention is correct, that should then create a context in which it is false that one knows that the son sneaked out. But it is implausible that the mere mentioning of a remote possibility can serve as the son’s defense here – in Lewis’s terminology, one should intuitively still be allowed to “properly ignore” this possibility.⁸ Thus the rule of attention, as it stands, is not a plausible feature of knowledge

8 DeRose (2004) suggests that such cases may lead to “gappy” semantics: contextualists may say that semantic evaluation of knowledge ascriptions depends on an agreed-upon “scoreboard”, so if two speakers fail to agree on such a

ascription and therefore cannot explain the asymmetry of which aspects of their meaning are cancelable.

Can we fix the “rule of attention”? Maybe mere attention is not enough to make a possibility become relevant – perhaps we need a notion of “warranted attention”. But if this is the case, we no longer have a feature that explains why *any* attempt to weaken the conveyed content of a knowledge ascription is infelicitous. For example, it still seems wrong to tell one’s son “I know you left the house yesterday, but I cannot rule out I just dreamed this.” Maybe the speaker herself needs to embrace the possibility in some way. But then it seems that two speakers within the same conversation can have different truth conditions for their knowledge ascriptions which seems independently implausible. Maybe attending to an uneliminated possibility can make a knowledge ascription seem problematic in some way without making it false. But this seems to fly in the face of the contextualist or relativist strategy as it introduces an important non-semantic feature governing our judgments of felicity. Distinguishing semantic and non-semantic features governing those judgments in this way would also seem *ad hoc* as there is no indication that these judgments are caused in different ways.⁹

scoreboard their utterances will not be assigned a truth value. But this position has serious drawbacks as well. First, we need to accept that there can easily be situations which defy semantic evaluation because of speaker disagreement. Second, the parent in this example will likely not be persuaded that neither of them is right, so we still lack an explanation of our intuitions about knowledge ascriptions and their truth.

9 It is worth mentioning that IPI can give an unproblematic account of these examples precisely because it treats the contextual variations as pragmatic. According to IPI, the father’s assertion was semantically false even before the son raised an error possibility, and this semantic evaluation does not change. However, the statement’s pragmatic meaning was true because in the original context remote possibilities were irrelevant to its evaluation. IPI is able to use a modified rule of attention to explain how the raising of error possibilities can *sometimes* lead to a change in the pragmatic meaning of knowledge ascriptions. However, it has no problem saying that it depends on further features of the context whether this move is successful – and that the son’s reply is a case where the pragmatic meaning does not change. That is because this *pragmatic* rule of (warranted) attention is not part of IPI’s account of the asymmetry

Are there alternative strategies for an explanation of the asymmetry? I am not aware of any other resources, but one might well take up the task of finding one. In general, what contextualists and relativists need is a feature that privileges the infallibilist perspective so that walking away from it will be infelicitous whereas walking towards it is fine. However, it seems that the problems of the “rule of attention” generalize: if it is held that the attempted cancellation can change the truth conditions of the knowledge ascription itself (as it is without the cancellation), then it seems to become too easy to raise the standards of knowledge, at least insofar as this feature can either be exploited to betray the purpose of the conversation or is not applicable widely enough to do all the explanatory work. If this feature does not affect the truth conditions, we are moving away from contextualism or relativism respectively and owe an explanation why this exception of the strategy of accounting for judgments of felicity through semantic truth conditions is warranted. I do not take these to be decisive considerations why such a feature could not exist but rather reasons why we, in the absence of a convincing proposal, can legitimately assume that it does not. If that is so, contextualism and relativism lack an explanation of the asymmetry regarding which aspects of the conveyed content of a knowledge ascription can be canceled.

So to sum up, there is an asymmetry with respect to which aspects of the conveyed content of a knowledge ascription regarding the subject’s epistemic standing can be canceled. IPI has a straightforward explanation of this phenomenon: as is typically the case, we can cancel the pragmatic aspects but not the semantic ones. FPI is committed to the problematic claim that the semantic aspects can be canceled but at least some of the pragmatic aspects cannot. Contextualists and relativists are not committed to non-cancelable pragmatic aspects, but they lack an explanation why a certain group of the semantic aspects of meaning is cancelable but another one is not. It also seems that there are general

regarding cancelability.

obstacles to providing such an explanation. While there are of course many other important arguments on both sides that need to be addressed, the situation regarding cancelability then provides some support for infallibilist pragmatic invariantism.

Acknowledgements: I would like to thank Michael Williams, Richard Bett, Justin Bledin and two anonymous reviewers for very helpful comments and suggestions on the ideas in this paper. I would also like the audiences at the Tübingen Masterclass with Jonathan Schaffer, at the European Congress for Analytic Philosophy 8 in Munich, and at the Hammond Society Colloquium at Johns Hopkins.

References

- Åkerman, J. (2015). Infelicitous Cancellation: The Explicit Cancellability Test for Conversational Implicature Revisited. *Australasian Journal of Philosophy* 93(3), 465-474.
- Blome-Tillmann, M. (2008). Conversational implicature and the cancellability test. *Analysis* 68(2), 156-160
- Blome-Tillmann, M. (2014). *Knowledge and Presuppositions*. Oxford: Oxford University Press.
- Brown, J. (2006). Contextualism and Warranted Assertability Manoeuvres. *Philosophical Studies* 130(3), 407–435.
- Cappelen, H. (2005). Pluralistic Skepticism: Advertisement for Speech Act Pluralism. *Philosophical Perspectives* 19(1), 15-39.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives* 13(1), 57-89.

- Conee, E. (2005). Contextualism contested. In M. Steup & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (pp. 47–56). Blackwell.
- Davies, A. (2017). Entailments Are Cancellable. *Ratio* 30(3), 288-304.
- Davis, W. (2007). Knowledge claims and context: loose use. *Philosophical Studies* 132(3), 395–438.
- DeRose, K. (1998). Contextualism: An Explanation and Defense. In J. Greco and E. Sosa (Eds.), *The Blackwell Guide to Epistemology* (pp. 187-205). Blackwell.
- DeRose, K. (2004). Single Scoreboard Semantics. *Philosophical Studies* 119(1-2), 1-21.
- Dimmock, P. and T. Huvenes (2014). Knowledge, conservatism and pragmatics. *Synthese* 191(14), 3239-69.
- Dinges, A. (2018). Knowledge and Availability. *Philosophical Psychology* 31(4), 554-573.
- Dinges, A. (2019). *Knowledge, Stakes and Error: A Psychological Account*. Frankfurt a. M.: Klostermann.
- Dodd, D. (2010). Confusion about concessive knowledge attributions. *Synthese* 172(3), 381–396.
- Dougherty, T. and P. Rysiew (2009). Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions. *Philosophy and Phenomenological Research* 78(1): 123-132.
- Dougherty, T. and P. Rysiew (2011). Clarity about concessive knowledge attributions: reply to Dodd. *Synthese* 181(3), 395–403.
- Douven, I. (2007). A pragmatic dissolution of Harman’s paradox. *Philosophy and Phenomenological Research* 74(2), 326–345.
- Gerken, M. (2017). *On Folk Epistemology*. Oxford: Oxford University Press.
- Grice, P. (1989). Logic and Conversation. In: id., *Studies in the Way of Words* (pp. 1-143). Cambridge, MA: Harvard University Press.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.

- Ichikawa, J. (2017). *Contextualizing Knowledge: Epistemology and Semantics*. Oxford: Oxford University Press.
- Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy* 74(4), 549-567.
- MacFarlane, J. (2005). The Assessment Sensitivity of Knowledge Attributions. In T. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology*, vol. I (pp. 197-233). Oxford University Press.
- Pynn, G. (2015). Pragmatic Contextualism. *Metaphilosophy* 46(1), 26-51.
- Rett, J. (2015). *The Semantics of Evaluativity*. Oxford: Oxford University Press.
- Rysiew, P. (2007). Speaking of Knowing. *Noûs* 41(4), 627–662.
- Schaffer, J. (2004). Skepticism, Contextualism and Discrimination. *Philosophy and Phenomenological Research* 69(1), 138-155.
- Schaffer, J. and Z. Gendler Szabó (2014). Epistemic comparativism: a contextualist semantics for knowledge ascriptions. *Philosophical Studies* 168(2), 491–543.
- Unger, P. (1975). *Ignorance. A Case for Scepticism*. Oxford: Oxford University Press.
- Weiner, M. (2006). Are all conversational implicatures cancellable? *Analysis* 66(2), 127-30.
- Williams, M. (2004). Knowledge, Reflection and Sceptical Hypotheses. *Erkenntnis* 61(2/3), 315-343.