A relativist semantics for *know*: *wh*-complements and intermediate exhaustivity

Ahmad Jabbar*

University of Connecticut

1 A Puzzle

(1) John knows who called.

In (1), *know* embeds an interrogative complement. Such knowledge-*wh* ascriptions are thought to have at least two readings: Strongly Exhaustive (SE) and Intermediately Exhaustive (IE), truth conditions of which are mentioned below.[1]

(2) SE conditions: (i) For all true answers to *who called?*, John believes that they are true; (ii) for all false answers to *who called?*, John believes that they are false.

(3) IE conditions: (i) For all true answers to *who called?*, John believes that they are true. (ii) of no false answers to *who called?*, does John believe that they are true.

*Many thanks to audiences at UConn Meaning Group; and to Derek Ball, Mitch Green, Ali Hamed, Drew Johnson, Magdalena Kaufmann, Zhiyu Luo, William Lycan, Floris Roelofsen, Lionel Shapiro, and Nadine Theiler. Thoughts and time of three anonymous reviewers are also greatly appreciated. All mistakes are my own.

[1]Note that these are instances of mention-all readings of knowledge-*wh* ascriptions, where the subject is required to believe of all the true answers to the question that they are true. Mention-some readings, where the subject is said to know w.r.t a question when she knows at least one answer to the question, have been explored in recent work by George (2011, 2013) and Phillips and George (2018). Moreover, earlier literature on questions features a third mention-all reading, Weakly Exhaustive (WE) reading, but recent experimental evidence suggests strongly against its existence. See Karttunen (1977) and Heim (1994) on weak exhaustivity, and Cremers and Chemla (2016) for evidence against WE readings. Also see Cremers, Tieu, and Chemla (2017) for recent acquisition data on exhaustive strength.

Groenendijk and Stokhof (1982, 180) present a strong argument against the existence of IE readings for first person knowledge-*wh* ascriptions like (4).[2]

(4) I know who called.

> "Suppose that John knows of everyone who walks that he/she does; that of no one who doesn't walk, he believes that he/she does; but that of some individual that actually doesn't walk, he doubts whether he/she walks or not. In such a situation, John would not say of himself that he knows who walks. We see no reason to override his judgement and to claim that in this situation, John does know who walk."

The observation is reasonable. If John believes of all those who called that they did, but doubts whether one of the non-callers called or not, then although John fulfills the IE conditions for *who called?*, he would not know that he fulfills (3i): of all true answers to *who called?*, he believes that they are true. Therefore, he wouldn't ascribe himself knowledge using (4), when he fulfills its IE conditions. Moreover, note that it is not just that John cannot assert (4), as upon being asked he would deny knowing who called. In other words, upon being asked, John asserts the negation of (4). Nor is it that John wouldn't assert (4) so as to not give rise to misleading expectations. As MacFarlane notes, this may be good reason to not assert (4), but not for asserting (4)'s negation.[3]

That IE readings for first person ascriptions like (4) don't exist may be taken as a reason to conclude that they don't even exist for third person ascriptions like (1) either. However, recent data by Cremers and Chemla strongly suggest that IE readings for third person ascriptions do exist. The following discussion relates to Experiment 2 in Cremers and Chemla (2016). The aim of this experiment was to determine the status of IE readings for third person ascriptions. Participants were given the following prompt.

> "John is playing a card game to train his memory. Each turn, he picks a card, quickly looks at it and flips it. John tries to recall what he saw on the card. Then he looks at the card again, and checks if he was right. You will see the card, and how John remembered the card before checking it. Using this information you will have to judge whether some sentences about the card and John's knowledge are true or false. All cards are made up of 4 squares, each of which can be red, green or blue." Cremers and Chemla (2016, 63)

The participants were then asked to look at the actual card that John saw (Figure 1(a)), and then consider John's beliefs about it, as shown in Figure 1(b) and Figure 1(c). The '?' in the bottom right square in Figure 1(c) represents John's agnosticism about the color of the bottom right square in Scenario 2. For each of the two scenarios, (b) and (c), the participants were asked to judge the truth value of (5).

(5) John knew which squares were blue.

[2]Strictly speaking, this argument is presented against weakly exhaustive readings, which is what Karttunen (1977) derives, but the argument works equally against IE readings.

[3]See Rysiew (2001) and MacFarlane (2014).

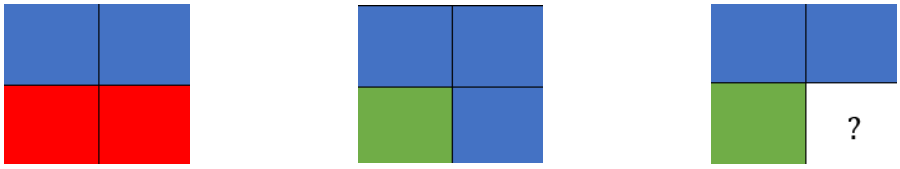

(a) *The actual card*

(b) *John's beliefs in Scenario 1*

(c) *John's beliefs in Scenario 2*

Fig. 1

For Scenario 2, the participants saliently judged (5) to be true. The rate of the response 'True' for (5) in Scenario 2 was over 95%. Let's break down Scenario 2 to understand the response.

In Scenario 2, John does not believe any of the false answers to *which squares are blue?* to be true answers. Moreover, of all the true answers to *which squares are blue?*, John believes that they are true. However, John is agnostic about one of the false answers to *which squares are blue?*. This means that (5) cannot be judged true under the SE reading, as that would require John to believe of all the false answers that they are false. This is not the case in Scenario 2, which displays the IE conditions for *which squares are blue?*. Thus, an IE reading for (5) is available to the participants of the experiment. The more general conclusion we can draw from this experiment is that external observers do ascribe knowledge to subjects who fulfill IE conditions for a given question.

It is crucial to note that, in Scenario 2, John wouldn't ascribe himself knowledge w.r.t. *which squares are blue?* for reasons mentioned above in Groenendijk and Stokhof's observation. If John is agnostic about the bottom right square's color, for all he knows, that bottom right square can be blue. Then, in Scenario 2, even though he believes of all the blue squares that they are blue, he doesn't know that this is the case. Combining this consideration with Cremers and Chemla's findings, we arrive at an interesting puzzle.

(6) ***Puzzle***: Under IE conditions, (i) while external observers ascribe knowledge to John, (ii) John does not ascribe himself the same knowledge.

All semantic theories of knowledge-*wh* ascriptions combine a semantics for *know* with a semantics for interrogatives. (6i) requires a theory to be able to derive IE readings.[4] In addition, (6ii) requires the IE reading to be somehow blocked in first person ascriptions. If we build such a blocking constraint in the semantics for *know*, we should be wary that it doesn't block the SE reading for first person ascriptions too.

(7) If the blocking constraint required for (6ii) is built in the semantics for *know*, it should not block the SE reading for first person ascriptions.

[4]See for instance Klinedinst and Rothschild (2011), Theiler et al. (2018), and Uegaki (2015).

Plan for the rest of the paper: (6) and (7) together give rise to an interesting problem, for which there is only one attempted solution in the literature: Theiler, Roelofsen, and Aloni, who, in their brilliant 2018 paper, propose to solve (6) by invoking an ambiguity for *know*. One entry for *know* requires the subject to fulfill the introspection condition; the other doesn't. Although this explains (6), in the next section, we argue that an ambiguity approach, especially Theiler et al.'s, is inadequate. Theiler et al.'s proposal and our arguments against it are explained in §2 and §3. In §4, we present our own solution to the puzzle. We propose a novel semantics for *know*. *Know* is interpreted relative to an information state as a parameter of the index. This allows the requirement to fulfill the introspection condition be contextual, and thus we don't need to invoke an ambiguity to solve the puzzle above.

2 The Ambiguity Proposal

Theiler et al. build their proposal in the Inquisitive Semantics framework (Ciardelli et al. 2018). Here, abstracting away from the framework, we explain their proposal as generally as possible. The reader is directed to pages 414-26 in Theiler et al. (2018) for a substantial account of how IE readings are derived. To give an overview, they take the interrogative complement in knowledge-*wh* ascriptions to be composed of an exhaustivity operator E and an interrogative nucleus. There are two types of E: E_{+cmp} and E_{-cmp}; and two types of nuclei: exhaustive and non-exhaustive. The combination of E_{+cmp} and the non-exhaustive nucleus yields the denotation of the complement under its IE reading. Note that this derivation is important, as IE readings for third person ascriptions exist. However, once we derive the IE complement, the following still requires explanation: when *know* embeds the IE complement, it blocks the IE reading for first person knowledge-*wh* ascriptions.

Theiler et al. have a fix. They propose that *know* is ambiguous between (8) and (9).

(8) $[\![know]\!]_{ext} = \lambda f.\lambda x.\lambda w.\ Dox_x^w \in f(w)$

(9) $[\![know]\!]_{int} = \lambda f.\lambda x.\lambda w.\ (Dox_x^w \in f(w) \wedge \forall v \in Dox_x^w : Dox_x^v \in f(v))$

Let's break these down. f, the denotation of the interrogative complement, is a function that maps w to the downward closed set of truthful resolutions to the interrogative in w. The set of truthful resolutions yielded by $f(w)$ varies with the reading (SE, IE, or MS) that the interrogative complement receives. The relevant one for now is f of the interrogative under its IE reading, so, to reduce clutter, we don't subscript f below. To see what truthful resolutions are, consider the toy model involving A and B with the four possibilities w.r.t. *who called?*: A and B both called (w_1); only A called (w_2); only B called (w_3); no one called (w_4). Suppose that the actual world is w_2. f takes w_2 to the set containing the

propositions depicted in Figure 2. These propositions are truthful resolutions of *who called?*, under its IE reading.[5]

(a) *Proposition 1* (b) *Proposition 2*

Fig. 2

Moreover, Dox_x^w in (8) and (9) is the set of worlds compatible with what x, the subject, believes in w. Now, it's easy to understand (8). If the subject's doxastic state is one of the members of the set in Figure 2, then the subject knows who called in w_2, in our toy model. However, if *know* is used under its interpretation in (9), then in addition to fulfilling the membership condition in (8), the doxastic state, Dox_x^w, must be such that for all worlds v in Dox_x^w, $Dox_x^v \in f(v)$. This amounts to the introspection condition, as all it says is that the subject knows that her doxastic state is a member of the set of truthful resolutions. As the subject's doxastic state being a member of the set of truthful resolutions amounts to knowing according to (8), $\forall v \in Dox_x^w : Dox_x^v \in f(v)$ amounts to knowing that one knows.

Introspection: $\forall v \in Dox_x^w : Dox_x^v \in f(v)$

Given the ambiguity, we can offer the following solution to the puzzle: the speaker denies knowing under the IE conditions due to their use of (9) with the IE complement—which is derived using the Inquisitive apparatus; moreover, as *know* can also be used, as interpreted as (8), the external observers can ascribe knowledge to John when he fulfills the IE conditions, as is the case in Cremers and Chemla's experiment. This explains Groenendijk and Stokhof's observation, while also accounting for Cremers and Chemla's data. Therefore, the ambiguity proposal seems to offer a neat solution to the puzzle. In the next section, by presenting a variety of data, we argue that an ambiguity approach for *know* is mistaken.

3 The case against ambiguity

For this section, we take Mary to be an agent whose beliefs are identical to those of participants in C&C's experiment above who consider (5) to be true under IE conditions. Then, Mary takes (5) to be true, while our subject, John, doesn't

[5]To compare, if we interpret f to be the denotation of the complement under its SE reading, then $f(w_2)$ is the singleton containing Proposition 2.

self-ascribe knowledge under IE conditions. As Mary and John believe contrary contents, a semantics for *know* should not preclude disagreement as in (10).

(10) John: I don't know which squares are blue.
Mary: No John, you do know.

Recall that Theiler et al. explain the puzzle above by proposing two different entries for *know*. John's assertion of ignorance is explained by his interpretation of *knowing which squares are blue* with $[\![know]\!]_{int}$. Mary's assertion above is explained by her interpretation with $[\![know]\!]_{ext}$. This explanation entails that the content of John's assertion is in no way incompatible with Mary's assertion, as they are asserting different contents. Why Mary disagrees with John above remains unexplained. As this is a serious objection, which potentially generalizes to any account that proposes ambiguity based on introspection requirement for *know*, it is worthwhile to consider other ways of explaining disagreement above.[6]

It is customary to motivate relativism by considering disagreement and retraction data (cf. MacFarlane (2014)). Recently, there has been pushback against this strategy – especially by Plunkett and Sundell (2013) and Khoo (2015), who explore ways of explaining disagreement, other than incompatibility of content expressed by the disagreeing parties. Here, we sketch what seems to us one such promising non-canonical explanation for (10).[7] Building on Barker (2002), P&S propose that some disagreements can occur when two parties have competing opinions on how to use a relevant word. For instance, when Ali thinks that the curry is spicy and Beth thinks that it isn't, their disagreement can be explained by their competing proposals for the threshold for spiciness. Co-opting this strategy for disagreement in (10) would amount to: John and Mary holding competing opinions about the standards of knowledge; while John believes that knowing requires introspection, Mary believes that it doesn't. This dovetails well with the ambiguity approach as this difference in opinion is then encoded in their use of *know* under different senses. This explanation has an unwelcome consequence in that it purports that some speakers, if they fulfill the IE conditions for *which squares are blue?*, would assert that they know which squares are blue. This is problematic due to G&S's argument sketched in §1, which shows that IE readings for first person ascriptions don't exist. Another consequence of this pragmatic explanation for (10) is that it entails that John would not retract his assertion, once Mary disagrees and John believes Mary; or, alternatively, that retraction here would amount to changing one's mind about how *know* is supposed to be used. We are agnostic about whether John would retract his assertion or not, but espousing the above sketched pragmatic explanation does commit a proposal to a view about retraction following the exchange in (10).

Moreover, the proposal that $[\![know]\!]_{ext}$ features in the semantic composition of Mary's ascription to John, and $[\![know]\!]_{int}$ in that of John's self-ascription, allows Mary to agree with John, as in (11).

[6]Thanks to Mitch Green and an anonymous reviewer for suggesting this line of exploration.

[7]Our consideration here doesn't exhaust the alternative explanations available. However, this illustrates the work to be put in for an ambiguity theory to work.

(11) John: I don't know which squares are blue.
Mary: # I agree, but you know which squares are blue.

Mary's response in (11) is infelicitous. Nonetheless, the possibility of such an exchange is predicted by Theiler et al.; Mary's agreement with John's utterance in (11), under the internal interpretation, does not preclude Mary from believing that John knows which squares are blue under its external interpretation. In other words, Mary can believe that John doesn't fulfill the introspection requirement, while asserting that he knows which squares are blue. If introspection is encoded in only one of the two entries for *know*, then the infelicity of Mary's utterance in (11) remains unexplained. Note that Mary uses *but*, instead of *and*. Therefore, the source of infelicity/oddness of (11) cannot be due to the lack of a contrasting marker. However, same infelicity stays when *and* is used.

(12) I know which squares are blue and so does John.

VP-ellipsis in knowledge-*wh* ascriptions such as (12) constitutes another argument against the ambiguity proposal. The elided VP in the second conjunct in (12) is supposed to have the same interpretation as the non-elided one in the first conjunct, just as in *John went to a bank and so did Mary*. To put (12) into context, let's consider C&C's experimental setting above. There, Mary, as an external observer, ascribes knowledge to John under the IE conditions. From this, we conclude that Mary's ascription to John is under the IE reading. Moreover, note that if we were to ask Mary if she knows which squares are blue, she would say she does—as she has epistemic access to the actual card (cf. Figure 1(a)). Subsequently, Mary fulfills introspection; her self-ascription either uses $[\![know]\!]_{int}$ or the SE complement. Therefore, the VP in the first conjunct in (12) is interpreted differently from the elided VP in the second conjunct under ambiguity.

Apart from not being simple, the data and arguments presented above make the ambiguity proposal an unattractive approach to adopt. Additionally, the puzzle becomes even more interesting, as the arguments in this section can be taken to pose a constraint: that to solve the puzzle above, *know* must not be taken to be ambiguous. In the next section, we propose a novel semantics for *know*, and argue that it is a promising solution to our puzzle.

4 Relativist *know*

We propose to make the interpretation of *know* relative to an information state, construed as part of the index. Without shift, the information state is provided by the context of utterance, which happens to be that of the speaker.[8] This maneuver allows us to keep Introspection a part of the semantics for *know*, which needs to be fulfilled only if the information state is identical to the doxastic state of the subject. We implement this informally sketched idea more formally below.

[8]This makes the proposal contextualist. The relativist tweak is suggested below.

The interpretation function $[\![\ ]\!]$ is defined w.r.t. two coordinates: c, which models the actual context of speech, and an index, $\langle w, i\rangle$, which is an n-tuple of all shiftable parameters.[9] For our purposes, it suffices to define the index to the extent of populating it with w, a world and i, an information state. The i in the index is initialized by the context, s.t. without any shift, i is the set of worlds not ruled out by what the speaker/agent of c knows.[10] With all the specifications, the semantics for *know* is defined in (13).

(13) $[\![know]\!]^{c}_{\langle w,i\rangle} = \lambda f.\lambda x.\ (Dox^{w}_{x} \in f(w)\ \ \wedge\ \ ((i = Dox^{w}_{x}) \rightarrow \forall v \in Dox^{w}_{x} : Dox^{v}_{x} \in f(v)))$

Breaking (13) down, where f is the denotation of the interrogative complement and x is the subject, x knows f iff two conditions are met.

(14) $Dox^{w}_{x} \in f(w)$

(15) $((Dox^{w}_{x} = i) \rightarrow \forall v \in Dox^{w}_{x} : Dox^{v}_{x} \in f(v))$

(14) is explained above with the help of the toy model in Figure 2. (15) is a conditional, where the consequent is Introspection. The crucial point here is that when $Dox^{w}_{x} = i$, for a knowledge-*wh* ascription to come out as true, Introspection must hold. $Dox^{w}_{x} = i$ when the speaker is ascribing knowledge to herself, as in first person ascriptions. Then, for first person ascriptions to come out as true, it is required that Introspection holds. As explained above, Introspection cannot be true without the subject knowing that her information state is such that she believes only of the true answers to some question that they are true. Therefore, John doesn't self-ascribe knowledge w.r.t. *which squares are blue?* when he is agnostic about whether one of the false answers to *which squares are blue?* is true or not (cf. Figure 1(c)).

Now take Scenario 2 (cf. Figure 1(c)) in Cremers and Chemla's experiment. External observers know what the actual card looks like (cf. Figure 1(a)), which makes $Dox^{w}_{x} = i$ false. Then, (15) is true even when Introspection, $\forall v \in Dox^{w}_{x} : Dox^{v}_{x} \in f(v)$, is false. In Scenario 2, that's exactly the case. John is agnostic about the color of the bottom right square; although it is true that $Dox^{w}_{x} \in f(w)$, $\forall v \in Dox^{w}_{x} : Dox^{v}_{x} \in f(v)$ is false. Yet, the external observers consider it to be true that John knows which squares are blue. Introducing a conditional in the semantics for *know* makes it possible to vary the truth of a knowledge-*wh* ascriptions w.r.t the information states of the agents. We can say that whether an agent ascribes knowledge to a subject covaries with the agent's information state. This is captured in (13) by relativizing the interpretation of *know* to an information state. Thus, we have a solution to our puzzle, without proposing *know* to be ambiguous.

[9] We follow Lewis (1980) in populating the index with shiftable parameters; the need for this has been challenged recently however (see Santorio (2017) and Rabern and Ball (2019)).

[10] Broadly, our proposal is comparable to Yalcin's (2007) semantics for epistemic modals. However, unlike Yalcin, we let i be initialized by the context. See discussion in MacFarlane (2014, 277).

Above, we set i to be initialized by the context of utterance. Here, we propose a tweak. To make sense of disagreement as in (10), while keeping our compositional semantics as in (13) the same, we propose i to be initialized by the context of assessment. In other words, we propose that it is the assessor's information state relative to which a knowledge ascription is interpreted. This can be implemented using MacFarlane's assessment-sensitivity framework. Traditionally, truth of a sentence is defined at a context. Once we bring MacFarlane's recent work into the picture, truth of a sentence is defined as uttered at a context c_1 and assessed from a context c_2. This allows i to be determined as follows: it is the set of worlds not ruled out given the information at c_2 — which without any shift happens to be that of the agent at c_2.[11]

In (10), let c_m be the context where Mary is the agent. According to our proposal, at c_m, *know* is interpreted relative to Mary's information state. Then, under IE conditions, at c_m, John's self-ascription would be true, due to the antecedent of (15) being false – making (15) true. Accordingly, John's negation of self-ascription is false at c_m. Therefore, we are able to explain why, in response to John's *I don't know which squares are blue*, Mary disagrees. Although our compositional semantics is compatible with i being initialized by the context of use, to account for disagreement cases, we think that i must be initialized by the context of assessment. This makes our theory for *know* relativist.

For reasons of space, we briefly consider only one possible objection to our proposal here. *Prima facie*, the scenario, where, by coincidence, the external observer has the same i as the subject's, is a counterexample. To see why, note that in such a scenario, (13) predicts that the subject needs to fulfill introspection for truth of a third person ascription; otherwise, the ascription comes out false. Although *prima facie*, a bad prediction, the worry is allayed on a closer look. Note that, under such a scenario, the observer truthfully ascribes knowledge to the subject, under the IE reading, only if the observer's information state is such that she fulfills introspection. Otherwise, the observer would not know if the subject believes of all those who called that they did. Then, by virtue of the information states of both the observer and the subject being identical (due to the purported counterexample), the subject needs to fulfill introspection too—which is exactly what (13) requires. Thus, peculiar at first, this scenario turns out to corroborate that the truth of knowledge ascriptions is sensitive to information states.

5 Taking stock and looking ahead

We have presented what seems to us an interesting problem. We have considered one solution to the problem, and presented our case against it. We have presented a more parsimonious proposal, which takes *know*'s interpretation to be relativized to the information state of the agent at the context of assessment. Our proposal is thus couched in two recent frameworks: Inquisitive Semantics

[11] See MacFarlane (2014). Also see Lasersohn (2005) and Stephenson (2007) for other versions of relativism.

and MacFarlane's assessment-sensitivity framework. In the following two paragraphs, we consider the scope of our proposal and its implications respectively.
Scope: As we haven't considered alternative pragmatic solutions at length, we think that our proposal should be taken to show what a *semantic* solution to the puzzle should look like. We didn't expand on our reasons for choosing the semantic route; however, what seems to us the most obvious pragmatic explanation is one that takes knowing that p as a necessary condition for asserting that p. This solution doesn't cut it, as it doesn't explain why, in spite of fulfilling the IE conditions for *who called?*, the subject asserts the negation of the knowledge ascription, i.e. *I don't know who called*. The said pragmatic route would work if the subject instead asserts *I don't know if I know who called*.
Looking ahead: Given that it is at the level of postsemantics that we let the information state be initialized by the context of assessment, a possible consequence is that epistemic modals are relativist too. However, this consequence obtains when the relativist postsemantics is coupled with a compositional semantics for epistemic modals that uses an information state parameter. In addition to offering a semantics for *know*, our proposal might have epistemological consequences. If indeed knowledge attributions are relative to information states of the assessor, then we should be wary of drawing conclusions about knowledge simpliciter from *our* intuitions about the usual thought experiments in epistemology.

References

Barker, C. 2002. "The Dynamics of Vagueness", *Linguistics and Philosophy* 25: 1-36.

Ciardelli, I., Groenendijk, J., & Roelofsen, F. 2019. *Inquisitive Semantics*. Oxford: Oxford University Press.

Cremers, A. & Chemla, E. 2016. "A psycholinguistic study of the exhaustive readings of embedded questions", *Journal of Semantics* 33: 49–85.

Cremers, A. Tieu, L., & Chemla, E. 2017. "Children's exhaustive readings of questions", *Language Acquisition* 24: 343-360.

George, B. 2011. *Question Embedding and the Semantics of Answers*, Ph.D. thesis, University of California, Los Angeles.

George, B. 2013. "Knowing-wh, mention-some readings, and non-reducibility", *Thought: A Journal of Philosophy* 2: 166–77.

Groenendijk, J. & Stokhof, M. 1982. "Semantic analysis of wh-complements", *Linguistics and Philosophy* 5: 175–233.

Groenendijk, J. & Stokhof, M. 1984. *Studies on the Semantics of Questions and the Pragmatics of Answers*, Ph.D. thesis, University of Amsterdam.

Heim, I. 1994. "Interrogative semantics and Karttunen's semantics for know", In R. Buchalla and A. Mittwoch (eds.), *The Proceedings of the Ninth Annual Conference and the Worshop on Discourse of the Israel Association for Theoretical Linguistics*. Academon: Jerusalem.

Karttunen, L. 1977. "Syntax and semantics of questions", *Linguistics and Philosophy* 1: 3-44.

Khoo, J. 2015. "Modal Disagreements", *Inquiry* 58 (5).

Klinedinst, N. & Rothschild, D. 2011. "Exhaustivity in questions with non-factives", *Semantics and Pragmatics* 4:1–23.

Lasersohn, P. 2005. "Context dependence, disagreement, and predicates of personal taste", *Linguistics and Philosophy* 28: 643–686.

Lewis, D. 1980. "Index, context, and content". In S. Kanger and S. Ohman (eds.), *Philosophy and Grammar*, D. Reidel, Holland. 79–100.

MacFarlane, J. 2014. *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.

Phillips, J. & George, B. 2018. "Knowledge-*wh* and False Beliefs: Experimental Investigations", *Journal of Semantics*. 35: 467-94.

Plunkett, D. & Sundell, T. 2013. "Disagreement and the Semantics of Normative and Evaluative Terms", *Philosophers' Imprint* 13 (23).

Rabern, B. & Ball, D. 2019. "Monsters and the Theoretical Role of Context", *Philosophy and Phenomenological Research*. 98: 392-416.

Rysiew, P. 2001. "The context-sensitivity of knowledge attributions", *Noûs*, 35: 477–514.

Santorio, P. 2017. "Context-Free Semantics", In E. Lepore and D. Sosa (eds.) *Oxford Studies in Philosophy of Language, Volume 1*. Oxford: Oxford University Press. 208-239.

Stephenson, T. 2007. "Judge dependence, epistemic modals, and predicates of personal taste", *Linguistics and Philosophy* 30: 487-525.

Theiler, N., Roelofsen, F., & Aloni, M. 2018. "A Uniform Semantics for Declar-

ative and Interrogative Complements", *Journal of Semantics* 35: 409-66.

Uegaki, W. 2015. *Interpreting Questions under Attitudes*. Ph.D. thesis, Massachusetts Institute of Technology.

Yalcin, S. 2007. "Epistemic Modals", *Mind* 116: 983-1026.