

The Ambiguity Theory of ‘Knows’

Mark Satta

(This is a final unpublished draft of an article forthcoming in *Acta Analytica*. The article can be found at <http://dx.doi.org/10.1007/s12136-017-0324-5>)

Abstract: The ambiguity theory of ‘knows’ is the view that ‘knows’ and its cognates have more than one propositional sense—i.e. more than one sense that can properly be used in ‘knows that’ etc. constructions. The ambiguity theory of ‘know’ has received relatively little attention as an account of the truth-conditions for knowledge ascriptions and denials—especially compared to views like classical, moderate invariantism and epistemic contextualism. In this paper, it is argued that the ambiguity theory of ‘knows’ has an advantage over both classical, moderate invariantism and epistemic contextualism. This advantage is that it is the only one of these views that can account for “diverging knowledge responses without inconsistency” (DRWI)—i.e. cases in which, for the same subject S and proposition p, one and the same speaker says truly “S knows p” but instead could have truly said “S does not know p” and vice versa. This paper argues both for the existence of DRWI scenarios and the ability of the ambiguity theory of ‘knows’ to best explain their existence.

The ambiguity theory of ‘knows’ (hereafter referred to simply as the ambiguity theory), is the view that ‘knows’ and its cognates have more than one propositional sense—i.e. more than one sense that can properly be used in ‘knows that’ etc. constructions. An immediate upshot of the view is that which sense of ‘knows’ is used in a knowledge ascription or denial determines, in part, the semantic content (and, as a result, the truth-conditions) of that knowledge ascription or denial.¹

In this paper, I make a case for the ambiguity theory by arguing that only the ambiguity theory can comfortably account for instances of what I call acceptable diverging knowledge responses. I make this case in Section 2. In Section 1, I begin by making clear what the ambiguity theory says and how it differs from its most popular competitors—epistemic contextualism and classical, moderate invariantism. I take the considerations offered in Sections

¹ From here on out the phrases ‘and its cognates’ from “‘knows’ and its cognates’ and ‘and denials’ from “ascriptions and denials” will typically be omitted, but should be understood as implied where appropriate.

2 to be enough to get the ambiguity theory up and running as a viable contender for explaining certain puzzles that have arisen concerning the truth-conditions of knowledge ascriptions, but this argument, on its own, does not constitute a full-fledged case for the ambiguity theory, which is not something that I think can be accomplished in a single paper.²

1. The Views Defined, Explained, and Compared

The ambiguity theory says that ‘knows’ has more than one propositional sense—i.e. more than one sense that can properly be used in ‘knows that’ constructions. A related, but distinct view, epistemic contextualism (hereafter simply contextualism) says that the truth-conditions of knowledge ascriptions and denials, in virtue of making a knowledge claim, shift in a non-trivial way in accordance with certain changes in the context of utterance. Contextualism and the ambiguity theory share several features in common—for example, they both allow for change in semantic content of knowledge ascriptions across contexts, and both seek to solve certain epistemological problems pertaining to skepticism.³ Despite these similarities, the views differ in important ways.

One key difference is that, according to contextualism, context plays an essentially *semantic* role which helps *determine* or *fix* the truth-conditions of knowledge ascriptions, while on the ambiguity theory context plays an essentially *pragmatic* role and only *guides* speakers in

² In Satta (2016), I argue that the ambiguity theory has the same major advantages as epistemic contextualism, and has a variety of additional strengths which epistemic contextualism lacks. In Satta (2017), I offer several linguistic arguments for an ambiguity theory of ‘knows’ (more specifically a polysemy theory) using standard tests for ambiguity such as those found in Cruse (1982) and Sennet (2011). For additional defenses of specific types of ambiguity theories see Malcolm (1952), Steup (2005), and Reed (2013).

³ As I’ve characterized contextualism here, it includes all of the more traditional forms of contextualism on offer—e.g. Lewis (1979), Cohen (1999, 2010) and DeRose (1999, 2009). There are ways in which this characterization elides certain elements of newer views which may also be considered contextualist like Brogaard (2008), MacFarlane (2007, 2009), and Pynn (2015a). That said, I take the basic structure of my comments on contextualism more broadly to apply to these views as well.

their sense selection and *assists* hearers in their recognition of what a speaker's ascription means.⁴ The general point can be made again using John Perry's useful distinction between the *semantic* and the *presemantic* which he describes in the following way,

“Sometimes we use context to figure out with which meaning a word is being used, or which of several words that look or sound alike is being used, or even which language is being spoken. These are *presemantic* uses of context: context helps us to figure out meaning. In the case of indexicals, however, context is used *semantically*. It remains relevant after the language, words and meanings are all known; the meaning directs us to certain aspects of context.”⁵

For Perry, rightfully so I think, ambiguity is a presemantic feature of a word. After one has determined which sense of the word is being used, context ceases to play an essential role in understanding the meaning of the word and how it functions within the larger lexical units it is embedded within. (Or at least context stops playing a greater role than it plays for any word whatsoever.) For example, when we encounter the word ‘bank’, we use context to figure out what meaning a speaker has employed, but once we’ve successfully done that, there is nothing especially context-sensitive about ‘bank’ that remains. This is in contrast with clearly context-sensitive words such as indexicals like ‘I’ in which context remains relevant after the language, words, and meanings are all known. ‘I’ remains tethered to the context even after presemantic tasks are completed, in a way that context does not for ambiguous words.

Both contextualism and the ambiguity theory allow for the variation of semantic content (and truth-conditions) of knowledge ascriptions *across contexts*. But the semantic versus

⁴ ‘Meaning’ is, of course, itself an ambiguous term. As used in this paper, ‘meaning’ will refer to that which the speaker literally says (with ‘says’ understood in a narrow way—for a discussion on differing senses of ‘says’, see, e.g. Saul (2015)) and is a product of the intentions and proper use of a language by a competent speaker of that language. The term ‘meaning’ here is interchangeable with the term ‘sense’ as used technically in a roughly Fregean way (*sinn*).

⁵ Perry (2001, 39-40).

presemantic roles that context plays respectively on the views leads to a second key difference between the two views. This difference is that *within a particular context* contextualism does not allow for variation of semantic content (nor for variation of truth-conditions) of knowledge ascriptions, while the ambiguity theory allows for variation of semantic content (and variation of truth-conditions) for knowledge ascriptions even within a particular context. The reason this is so is because on contextualism it is a change in context that is responsible for the variation in semantic content and truth-conditions, while in the case of the ambiguity theory it is the need for sense disambiguation that allows for the variation.⁶ On contextualism, each context only allows for one meaning, which is fixed within that context and which fixes the truth-conditions of the ascription in that context. But all the senses of an ambiguous word remain available for selection in all contexts (so long as the appropriate syntax is present), even if some senses are more felicitous to use than others.

An upshot of this situation is that on both contextualism and the ambiguity theory the following principle is true:

Surface Conflict Without Inconsistency (SCWI): There are cases in which, for the same subject S and proposition p, one speaker can truly say “S knows that p” and another speaker, at the same time, can truly say “S does not know that p”.

However, only on the ambiguity theory and not on contextualism is this second principle true:

Diverging Responses Without Inconsistency (DRWI): There are cases in which, for the same subject S and proposition p, at a given time t, one and the same speaker says truly “S knows that p” but instead could have truly said “S does not know that p” and vice versa.⁷

⁶ Thus, on the ambiguity theory, the source of the variation is whatever it is that determines which sense of a word has been employed. I favor a view on which this source is the intention of a speaker (so long as the speaker is a sufficiently competent user of the language who is using the words within the bounds of what the language allows for). Moving forward, I will employ this understanding of what accounts for sense disambiguation. However, this view about what accounts for sense disambiguation is not part of the ambiguity theory itself. One can reject my view of how senses get disambiguated and still accept the ambiguity theory of ‘knows’.

⁷ It is also the case that only on the ambiguity theory and not on epistemic contextualism is the following true. **DRWI*:** There are cases in which, for the same subject S and proposition p, at a given time t, one and the same speaker can truly choose to say “S knows that p” and can also truly chose to say “S does not know that p.” If one

Contextualism must reject DRWI because context-sensitivity (via whatever the semantic mechanism may be) allows sentences and/or utterances to pick out different propositions and to have varying truth-conditions across different contexts, but doesn't allow for a sentence to pick out varying propositions or to have varying truth-conditions within a particular context. In contrast, the ambiguity theory allows for DRWI because the semantic mechanism of ambiguity allows sentences and/or utterances to pick out different propositions and to have vary truth-conditions *both across different contexts and within a particular context*. SCWI requires only what context-sensitivity and ambiguity have in common, namely, the ability for varying truth-conditions for a sentence and/or utterance across contexts. But DRWI requires what only ambiguity allows for, namely, the ability for varying truth-conditions for a sentence and/or utterance within a particular context.⁸ This difference will play an important role in Section 2, where I argue that the ambiguity theory has an important advantage over both contextualism and moderate classical invariantism.

Invariantism, in the epistemological sphere, is often simply taken to be the view that contextualism is false.⁹ However, invariantism is frequently described in a way that makes it more than simply the denial of contextualism. For example, DeRose writes that,

“Thus, the contextualist will allow that one speaker can truthfully say ‘S knows that P,’ while another speaker, in a different context where higher standards are in place, can truthfully say ‘S doesn’t know P’ at the same time. The “invariantist” – Peter Unger’s good name for one who denies contextualism – will have none of this.

prefers to consider a formulation that avoids ‘could’, the arguments that follow can just as easily be run substituting DRWI* for DRWI.

⁸ From here on out the relevant difference between the truth-conditions of sentences versus the truth-conditions of utterances will be suppressed with phrases like ‘the truth-conditions of knowledge ascribing sentences’ and ‘the truth-conditions of knowledge ascribing utterances’ being used interchangeably unless otherwise specified. On certain views, such as Geoff Pynn’s pragmatic contextualism (2015a, 2015b), this distinction between the truth-conditions of utterances and sentences is relevant. But the distinction is not relevant for our purposes.

⁹ For example, Rysiew (2011) writes that “one is an invariantist (with respect to a given class of statements) just in case one rejects contextualism (with respect to statements of that class)”.

According to her, there's a single, invariant set of standards which, at least as far as truth-conditions go, govern the use of knowledge attributions regardless of the context in which they're uttered; thus, our two speakers can't both be speaking a truth."¹⁰

This description implies that, in addition to being the denial of contextualism, invariantism is also the denial of SCWI. If I am correct in thinking both that the ambiguity theory is not a type of contextualism and that the ambiguity theory entails SCWI, the extension of what counts as invariantism will differ depending on whether or not invariantism is seen simply as the denial of contextualism or as the denial of contextualism and SCWI. I am not committed to one understanding over another, but in order to keep things clear in the discussion ahead, in this paper by 'invariantism' I mean the view that (1) contextualism is false and that (2) SCWI is false. Thus defined, the ambiguity theory is not a type of invariantism (although on other reasonable definition of the term it would be a type of invariantism). Given than the falsity of SCWI entails the falsity of DRWI, invariantism, as understood in this paper, also results in the denial of DRWI.

In the next section, I will argue that the ambiguity theory has an advantage over both contextualism and the form of invariantism that DeRose has called "the real threat to contextualism"¹¹ and "the most influential source of resistance to contextualism."¹² This view is a form of classical, moderately non-skeptical invariantism that employs what DeRose calls a "warranted assertability maneuver".¹³ Thus, it will be important to get clear on the extra commitments one holding this view has.

¹⁰ DeRose (1999, 188). John Turri conveys a similar understanding writing that "Invariantism denies that the truth-conditions of knowledge ascriptions are context-sensitive. They remain constant, invariantists say, across contexts." (Turri (2010), 78). John Hawthorne may be committed to something similar when he writes that "[a]ccording to the *invariantist*, the semantic value of 'know', and accordingly the extension of 'know', does not vary across contexts of use" (Hawthorne 2004, 113).

¹¹ DeRose (2004b, 349-350).

¹² DeRose (2002, 170-171).

¹³ DeRose (1999, 2002, 2005, 2009).

What Jessica Brown and Christoph Kelp (among others) have called classical invariantism¹⁴ adds to invariantism the denial of the following claim.

Pragmatic-Encroachment (PE): A difference in pragmatic circumstances can constitute a difference in knowledge.¹⁵

Therefore, classical invariantism does not include views like the interest-relative invariantism of Jason Stanley or the subject-sensitive invariantism of John Hawthorne, Jeremy Fantl, and Matthew McGrath.¹⁶ Rather, the classical invariantist seeks a view more akin to that advocated in the initial post-Gettier program, where a single concept of knowledge was typically assumed and it was assumed not to have any purely pragmatic features.

As has been noted by Brown, DeRose, and Hawthorne (among others), there can be different types of classical invariantism depending on the level of stringency required for knowledge.¹⁷ In general, two broad categories of classical invariantism have been put forward. First, skeptical invariantism, a view held perhaps most famously in the post-Gettier era by Peter Unger in his younger years,¹⁸ refers to a type of invariantism on which the standards for knowledge are high enough that most of our everyday knowledge attributions are false, strictly speaking. Due to the relative unpopularity of this view, it will not be focused on in this chapter.¹⁹

Instead, we will focus on the second view, which John Hawthorne has labeled moderate invariantism.²⁰ Moderate invariantism is a form of classical invariantism on which the standard

¹⁴ Brown (2005, 2006), Kelp (2008).

¹⁵The phrasing of this claim comes from Ichikawa and Steup (2014). Like many others in the literature on knowledge ascriptions, in this context, I take the terms 'pragmatic' and 'practical' to be interchangeable.

¹⁶ See, for example, Stanley (2005), Hawthorne (2004), and Fantl and McGrath (2009).

¹⁷ See Brown (2006), DeRose (1999), and Hawthorne (2004, ch. 3).

¹⁸ Unger (1975).

¹⁹ For a more modern (and in my mind quite plausible) defense of this view see Wayne Davis (2007). Davis argues that most of our everyday knowledge claims are cases of loose use in which the claims are strictly speaking false but appear both true and appropriate. This is roughly the skeptical parallel of the invariantist view being considered in this chapter.

²⁰ Hawthorne (2004, ch 3).

for knowledge is low enough such that most of our everyday knowledge attributions are true. (What exactly those standards are isn't an essential element of the family of views picked out by the term moderate invariantism and what exactly those standards are won't be important for our purposes.) Thus, in short, classical moderate invariantism (hereafter referred to simply as moderate invariantism or MI) is the view that contextualism, SCWI, DRWI, and pragmatic encroachment are all false, and that the single invariant standard of justification/warrant for knowledge is low enough such that most of our everyday knowledge ascriptions turn out true.

But in order to understand fully the type of invariantism we'll be comparing the ambiguity theory to, we also need to understand the response it gives to the typical contextualist data, which initially don't seem to jive well with moderate invariantism. (By contextualist data, I mean the things that contextualists usually put forward to support their view—things like DeRose's banks cases, Cohen's airport case, inconsistent brain-in-a-vat/hands triads, etc.). The dominant (and I believe most compelling) method put forward is to make, what DeRose calls, a *warranted assertability maneuver*. DeRose explain a warranted assertability maneuver (hereafter a WAM) as follows.

“A WAM involves explaining why an assertion can seem false (or at least not true) in certain circumstances in which it is in fact true by appeal to the fact that the utterance would be improper or unwarranted in the circumstances in question.”²¹

In the case of the debate between the contextualists and the classical invariantists, the specific WAM employed by the proponent of MI is the following: “what the contextualist takes to be a variation in the truth-conditions of knowledge attributions is in reality only a variation in the conditions for the warranted assertability of those claims.”²²

²¹ DeRose (2002, 171).

²² DeRose (2002, 171).

In order to see more clearly how the WAM operates, let's look at a concrete example of the WAM as used in response to DeRose's popular bank cases.

Bank Case A: My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says 'Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays.' I reply, 'No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon.'

Bank Case B: My wife and I drive past the bank on a Friday afternoon as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before Monday morning, the important check will bounce, leaving us in a *very* bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, 'Banks do change their hours. Do you know the bank will be open tomorrow?' Remaining as confident as I was before that the bank will be open then, still, I reply, 'Well, no, I don't know. I'd better go in and check.'²³

DeRose then adds the following: "Assume in each case that the bank will in fact be open on the Saturday in question, and that there is nothing unusual about either case that has not been included in the description of it."²⁴

The intuition that the contextualist expects and desires of listeners (and that listeners often seem to have) is that both the speaker's first person knowledge attribution in Case A and the speaker's knowledge denial in Case B are true and appropriately made statements, and that this is so despite the fact that it remains true in both cases that the bank will be open, the degree of confidence of belief remains the same, and the level of justification for the belief remains the same. The contextualist argues that these intuitions are evidence in favor of contextualism. I have

²³These versions of the cases come from DeRose (2009). Similar iterations of these cases have been presented by DeRose and others on many occasions.

²⁴ DeRose (2009, 2).

argued elsewhere that this move is too quick, and that this evidence just as strongly supports the ambiguity theory. However, that's not the point here.

Moderate invariantists must offer an explanation of these intuitions. For, on moderate (or any form of classical) invariantism the knowledge assertion in Case A and the knowledge denial in Case B cannot both be true, given the circumstances of the case. For the moderate invariantist who employs a WAM, the basic response goes as follows. The speaker's knowledge assertion in Case A is both true and pragmatically appropriate to assert. In Case A, the truth of the claim and the appropriateness of assertion are in tandem. However, in Case B, the warranted assertability conditions of what it is pragmatically appropriate to assert come apart from the truth of the claim. In Case B, according to our moderate invariantist, the knowledge denial is strictly speaking false, but it is appropriate to assert given the circumstances. In short, the invariantist who uses a WAM claims that the truth-conditions for knowledge assertions and denials remain constant across contexts, but that what it's appropriate to assert varies contextually. From here on out the only form of classical invariantism that will concern us will be a form of moderate invariantism (MI) that employs a warranted assertability maneuver (WAM) in order to explain the contextualist data. I'll refer to this view simply as MI-WAM.²⁵

2. The Case for the Ambiguity Theory via DRWI

Recall that the ambiguity theory entails the following principle,

Diverging Responses Without Inconsistency (DRWI): There are cases in which, for the same subject S and proposition p, at a given time t, one and the same speaker says truly "S knows that p" but instead could have truly said "S does not know that p" and vice versa.

²⁵ The narrowing of our focus to this form of WAM is primarily because this is the form of response that has been most prominent in the literature, but this is not to say that there aren't other responses a moderate invariantist could employ.

Recall also that this principle is incompatible with both contextualism and MI-WAM. Thus, if we have a good reason to accept DRWI, then we have a good reason to accept the ambiguity theory over contextualism and MI-WAM. But do we have a good reason to accept DRWI?

I want to offer a couple of cases for which I think an acceptance of DRWI is the best explanation of certain linguistic intuitions. The base for the first case comes from Jessica Brown. The base for the second comes from Baron Reed. I've added to both cases in order to modify their original purposes to better serve my own. Let's begin with Brown's *Surgeon* case. The case as originally presented goes as follows.

“A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is lying anaesthetized on the operating table. The operation hasn't started as the surgeon is consulting the patient's notes. The student is puzzled and asks one of the nurses what's going on:

Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She should not operate before checking the patient's records.”²⁶

In this case the nurse's claim that the doctor knows which of the patient's kidneys is to be removed is meant to be taken intuitively as a true statement and appropriate response. Initially, it likely is not clear why this case should favor either DRWI or its rejection. However, the relevance of the case comes into focus once we consider Fantl and McGrath's claim that “the nurse in Brown's surgeon case could just as easily and just as properly have said, ‘Well, of course she's checking the chart; it's not enough to rely on her memory that it's the left kidney.

²⁶ Brown (2008, 176).

She needs to know it is’.”²⁷ The implication of such a claim on the nurse’s behalf amounts to an implicit denial that the doctor knows which of the kidneys is to be removed. Let us label the response that Fantl and McGrath claim the nurse could give (along with an explicit knowledge denial for good measure) *Nurse Response 2 (NR-2)* and the original response that Brown claims the nurse could give *Nurse Response 1 (NR-1)*. Our options are then as follows:

NR-1: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She should not operate before checking the patient’s records

NR-2: Well, of course, she’s checking the chart; it’s not enough to rely on her memory that it’s the left kidney. She needs to know it is, and without double-checking she doesn’t know it.

Both NR-1 and NR-2 deal with the same subject (the surgeon), the same proposition (it is the left kidney of the patient that is to be removed) and are given in response to the same context.²⁸ NR-1 contains an explicit knowledge ascription while NR-2 contains an explicit knowledge denial concerning the same subject, proposition and context. For those with the intuition that both NR-1 and NR-2 represent plausibly true and appropriate responses, this poses a problem for views that deny DRWI like contextualism and MI-WAM, because the context prompting each response is identical to the other,²⁹ yet both claims seem very plausibly true. Thus, on contextualism and MI-WAM some explanation as to why both the knowledge ascription and knowledge denial appear

²⁷ Fantl and McGrath (2012, 71).

²⁸ Obviously, once NR-1 or NR-2 is asserted the context is different than one in which the other claim had been asserted instead. But prior to the assertion the context for the two scenarios is the same—or at least the context is the same if we take a view of context as a sharable thing created by the conversational partners together, and not merely what is going on in one party’s head. (I take this to be how context is typically understood by epistemic contextualists. See footnote 29.)

²⁹ One might protest that surely these contexts aren’t *identical* to each other, for presumably, even if externally the contexts appear identical, there must be some kind of internal difference in the mental states of the nurse that prompt her to respond differently in the two scenarios. My response is that on the dominant pictures of conversational context (e.g. Lewis (1979), Stalnaker (2014)—the relevant sort of context at hand—context appears to be something that exists between persons in conversation. This is further verified by the “single-scoreboard semantics” that seems to have been either explicitly or implicitly been adopted by contextualism’s most prominent defenders (e.g. DeRose 2004, Cohen 1999). Thus, even if there is a necessary difference in the mental states of the nurse leading up to her utterance, this mental state isn’t a relevant part of the context—although any altered behavior this leads to (in this case her diverging response) is part of the context and explains why the two contexts diverge once different responses are posited.

true (even though one of them is not) is needed. In a moment, I consider several attempts at giving such an explanation, but before turning to those possible explanations, let's look at another scenario.

This following case comes from Baron Reed:

Punishment/reward case 1. You are participating in a psychological study intended to measure the effect of stress on memory. The researcher asks you questions about Roman history—a subject with which you are well acquainted. For every correct answer you give, the researcher will reward you with a jelly bean; for every incorrect answer, you are punished by an extremely painful electric shock. There is neither reward nor punishment for failing to give an answer. The first question is: when was Julius Caesar born? You are confident, though not absolutely certain, that the answer is 100 BC. You also know that, given that Caesar was born in 100 BC, the best thing to do is to provide this answer (i.e., this course of action will have the best consequences—you will be one jelly bean richer!).³⁰

Reed then points out, however, that due to the significant stakes if you are wrong, you may not want to answer. Reed says “it would be natural for a subject to say” the following:

Subject Response 1 (SR-1): I know this one, but I'm not going to risk an answer.³¹

Aside from those who have a commitment to a knowledge-action principle who, for theoretical reasons, may want to resist the conjunction, most of us will likely find this an appropriate and accurate response. However, it also seems that the following would have been a true, appropriate, and natural response as well:

Subject Response 2 (SR-2): I'm pretty confident about this one, but I don't know this one, so I'm not going to risk an answer.

³⁰ Reed (2010, 228-229). As with the *Surgeon* case, I am not using this example for its original purpose.

³¹ Reed (2010, 229). An assumption here is that one can be in a highly justified position concerning the truth of the proposition that Caesar was born in 100 B.C.E. If one doesn't think we can be in such a position with regard to the truth of that proposition, they are encouraged to replace the historical proposition with one for which it is possible to have that high level of justification.

Once again, in both SR-1 and in SR-2 we have the same subject (you via your first person attribution or denial of your own knowledge), the same proposition (“this one” referring to the proposition that Julius Caesar was born in 100 BC), and the same context (including subjective factors like your confidence level in the truth of the proposition). Once again this looks initially like it causes a problem for contextualism and MI-WAM.

Intuitions that both NR-1 and NR-2 are true and/or that both SR-1 and SR-2 are true support the ambiguity theory over contextualism and MI-WAM. We can imagine the nurse to have in mind the ‘beyond a reasonable doubt’ standard of knowledge when giving NR-1, and we can imagine her to have in mind a more demanding standard of knowledge according to which knowledge requires double or even triple checking in NR-2, thus grounding her differing claims in the two scenarios. On a form of the ambiguity theory that recognizes both these senses of ‘knowledge’, the nurse’s knowledge ascription and her knowledge denial can both be true. I refer to the nurse’s ability to given either a knowledge ascription or a knowledge denial as a case of acceptable diverging knowledge responses. Positing a switch between similar knowledge-standards in the Punishment/Reward case could also ground acceptable diverging knowledge responses on an ambiguity theory as well.

So long as the context is the same leading up to both utterances, and so long as both the ascription and the denial are true, contextualism and MI-WAM cannot account for the data. Of course, the proponents of contextualism and MI-WAM could try to explain away the seeming truth of one of the responses in both pairs by appeal to the MI-WAMmer’s preferred strategy of trying to explain why both claims strike us as *appropriate* responses, even though one or the other of the responses is not true. For the proponent of MI-WAM, for example, in the case of SR-1 the appropriateness is grounded in the truth of the claim. And in the case of SR-2, a story is

told about how SR-2 implies something true (even though it is strictly speaking false) and as a result strikes us as appropriate. However, a WAM cannot account well for why both responses might still seem true even after careful reflection, because, if this is what is going on, the true implication generated by the false statement should be cancellable. And on a WAM one response from each pair must be false. But it seems that a reasonable case can be made that neither claim from either pair is false. Thus, cases of this sort pose a problem for MI-WAM and contextualism because they seem to imply the truth of DRWI.³² The ambiguity theory can avoid this difficulty entirely by positing that ‘know’ is being using in a different sense in SR-1 than it is in SR-2.

Before continuing, it seems worth pausing to consider what other sorts of explanations, aside from a WAM, might be given in order to offset this initial advantage that the ambiguity theory seems to have. I’ll consider two other explanations for the intuitions of the truth of both NR-1 and NR-2 and of SR-1 and SR-2. The first potential explanation is that ‘knows’ is not ambiguous, but vague, and that these cases pull us in inconsistent directions because they are in the fuzzy boundary area of ‘knows’. The second potential explanation is that while ‘knows’ is not vague, the situations in *Surgeon* and *Punishment/Reward* represent “close calls”, where the threshold for knowledge is either just met or just not quite met. On this account, it is because it is close call we are apt to be swayed that both the knowledge ascription and the knowledge denial are true because neither is obviously false in virtue of it being a close call.

³² A potential additional problem in using a WAM to explain away the appearance of acceptable diverging knowledge responses in these cases comes from the fact that in each case they will need to be employed to explain away “an intuition of truth” (i.e. they will be needed to explain why something false appears true). DeRose has argued a length that a WAM should only explain away an “intuition of falsehood”, but not an “intuition of truth” (DeRose 1999, 198-9). See also DeRose (2002, 2009). If this criticism is on target, it would apply in cases of diverging acceptable knowledge responses as well. For responses to this line of criticism see Brown (2006) and Rysiew (2001, 2005, 2007).

Let's start with the first potential alternative explanation, namely, that 'knows' is vague.³³ Given that the intended role that *Surgeon* and *Punishment/Reward* are meant to play in this paper (namely, as evidence for the ambiguity theory over contextualism or MI-WAM), I think the following two questions are the appropriate ones to ask. Can vagueness function as a plausible explanation for our intuitions from *Surgeon* and *Punishment/Reward* in a way that is compatible with MI-WAM? Alternatively, can vagueness so function in a way compatible with contextualism? We'll start with the question about MI-WAM and then turn to the question about contextualism.

Let's starting by make more explicit how 'knows' being vague could serve as an explanation on an invariantist account for the intuitions one may have that both responses in *Surgeon* and/or in *Punishment/Reward* are true. I'm assuming that if 'knows' is vague that this vagueness enters via vagueness about the level of justification required to count as knowing.³⁴ How one accounts for the vagueness in the level of justification required for knowledge will change depending on other epistemological assumptions. For example, on an evidentialist account, there will be a fuzzy strength of evidence threshold marking the difference between knowing and not knowing; on a reliabilist account, there will be a fuzzy degree of reliability threshold that marks that difference; etc. The important point unifying all these views is this: the justification threshold required for knowledge is not precise but fuzzy. On any such account combined with invariantism, the resulting single sense of 'knows' is one with a fuzzy boundary, meaning that there likely will be cases where there is no clear matter of fact about whether or not one knows. In order for *Surgeon* and *Punishment/Reward* to be cases where the vagueness

³³ That 'knows' is vague is the view that Stephen Schiffer seems to me to favor in response to contextualist data in his influential and oft cited "Contextualist Solutions to Scepticism" (1996).

³⁴ I think this assumption is warranted, given that I think most invariantist proponents of a vagueness view of 'knows' would say the same.

affects the truth of the utterances, the level of justification that the nurse and you have respectively must be in that grey area where it is not clear that the level of justification is high enough.

Thus, even though the claims made in the responses to *Surgeon* and *Punishment/Reward* may not be true (whether they are will depend on your theory of the semantics of vague terms), given that the state of the putative knower in question is in the penumbra, one can easily understand both the intuition that NR-1 is true and that NR-2 is true. Similarly, one may have the intuition that a speaker referencing a man with one hundred hairs on his head may be saying something true if she says of the man that he is bald (due to the low number of hairs on his head) or if she denies that he is bald (due to her focus on the fact that he does indeed still have one hundred hairs on his head).

This on its own strikes me as a plausible explanation of the intuitions one may have about *Surgeon* and *Punishment/Reward*. However, I don't think this explanation is available for the defender of MI-WAM. The reason being that a constitutive part of MI-WAM is that the standard of knowledge is low enough that most of our everyday knowledge ascriptions turn out true. However, recall that in *Surgeon* the doctor has strong justification for her belief that it is the left kidney that is to be removed (after all she had consulted with that same patient earlier that very day). In order for the standard of knowledge to be low enough for most of our everyday knowledge ascriptions to turn out true, it would seem that the surgeon's evidence ought to be strong enough to be a clear case of knowledge. If that level of justification is part of the penumbra, it seems many of our other everyday knowledge claims will also be in that penumbra—or worse, beyond the penumbra and no longer instances of knowledge. Note also that the defender of moderate invariantism, if she is appealing to vagueness, cannot cite anything

about the practical importance of getting the matter right in the case of *Surgeon* as an explanation for why this is an instance of a knowledge ascription in the penumbra. That practical concerns like this affect whether one knows is one of the claims that the moderate invariantist denies. In short, the *Surgeon* case is not a plausible case of a fuzzy area on the moderate invariantist picture.

This problem presents itself all the more clearly in *Punishment/Reward* where we are told that “[y]ou are confident, though not absolutely certain, that the answer is 100 BC.” Assuming that your confidence is warranted, it seems like it is the lack of absolute certainty which is motivating the knowledge denial in SR-2. And for the moderate invariantist, this should clearly put the level of justification past the fuzzy boundary area. Trying to explain *Punishment/Reward* as a case of a fuzzy boundary for knowledge will lead to either an implausibly large boundary area or to something akin to skepticism. And neither option is desirable on moderate invariantism. Thus, while vagueness may make a plausible response in its own right to what’s going on in *Surgeon* and *Punishment/Reward*, it won’t be of much help to the moderate invariantist.

The situation is different when it comes to combining vagueness and contextualism. Because the standard of justification required for ‘knowledge’ can change across contexts, the contextualist doesn’t encounter the problem of *Surgeon* and *Punishment/Reward* being situations in which the standard of justification is too high for the claims to be in the fuzzy boundary area on contextualism. However, there is another problem for the contextualist—namely, the presence of vagueness seems to undermine the function that contextualism is meant to play in the first place. It does this in two ways.

First, one way to understand contextualism is that it is the view that context takes the acontextual penumbral region in which things might count as knowledge and removes the vagueness by fixing a standard when knowledge ascriptions are made within particular contexts. The idea, at least on some traditional conceptions of contextualism, is that by providing a fixed standard of knowledge, one removes the semantic and truth-conditional indeterminacy (of the type that plausibly often accompanies vagueness) in order to create determinate truth-conditions for knowledge ascriptions. Thus, positing vagueness in addition to context-sensitivity for ‘knows’ is theoretically awkward.

Second, positing vagueness as an explanation for our intuitions in cases like *Surgeon* and *Punishment/Reward* undermines contextualism by suggesting that a second semantic mechanism, in addition to context-sensitivity is required to explain the data. One runs the risk of an over-determinacy that renders context-sensitivity superfluous. I don’t think either of these issues completely removes the helpful force that combining a vagueness theory with contextualism might serve in helping the contextualist respond to cases of seeming acceptable diverging knowledge responses, but these considerations do weaken the case in such a way that the ambiguity theory retains its advantage as a simpler, more intuitive explanation of the data provided by cases like *Surgeon* and *Punishment/Reward*.

Let’s turn now to our second potential explanation, that this is the case of a close call where, because these cases are on the border of what counts as knowledge, our intuitions are easily pulled both ways. Given what we’ve already established, this suggestion when combined with MI-WAM can be treated much more quickly. This is because the suggestion that these are close calls, while plausible in its own right, cannot be paired up successfully with moderate invariantism. The reason for this is that, unless the category of “close calls” is very large (which

seems perhaps antithetical to the concept of a close call in the first place), *Surgeon* and *Punishment/Reward* being cases of close calls won't allow for most of our everyday knowledge ascriptions to be true. This is because if *Surgeon* and *Punishment/Reward* were close calls, the boundary region where close calls happen would be at a justification level so high that many of our everyday ordinary language assertions would cleanly fall outside of this close call region and into the region of the clearly false. That this is so is once again a product of the fairly high level of justification had by the speakers in both *Surgeon* and *Punishment/Reward*. Thus, this suggestion does no favors for MI-WAM.

What about the combining a close calls view with contextualism? Unlike the case of combining contextualism with vagueness, there is not a problem of the justification level being too high on *Surgeon* and *Punishment/Reward* to rule out the close call explanation. However, combining contextualism with a close call account does have its weaknesses. First, unlike on the ambiguity theory, on a contextualism-combined-with-close-calls view, it is no longer a single semantic mechanism which explains the data. Context-sensitivity is no longer sufficient to explain what's going on. An appeal to something extra (namely, the close call and its ability to mislead our intuitions) is needed in order to explain the data. As a result, contextualism fails to be as simple or straightforward an explanation as the ambiguity theory in explaining situations that seems to allow for diverging acceptable knowledge responses.

Second, in his discussion of WAMs, DeRose puts forward a general principle that a WAM ought only to explain away "intuitions of falsehood" and not also "intuitions of truth", where an intuition of falsehood is the intuition that a claim appears false despite actually being true and an intuition of truth is an intuition that a claims appears truth despite actually being

false.³⁵ However, a close calls explanation of *Surgeon* and *Punishment/Reward* on contextualism seems to amount more or less to a WAM on which intuitions of truth are being explained away. (Or at the very least, it amounts to something quite similar.) This is because on the contextualist's close call account of *Surgeon* and *Punishment/Reward*, the contextualist uses the fact that NR-2 and SR-2 are close calls to explain why the claims are perceived as appropriate and true despite being false. Perhaps there is something refinement between close call cases and traditional WAMs that can limit DeRose's principle such that it doesn't apply to *Surgeon* and *Punishment/Reward*. But minus some clear delineation, if DeRose is right, a weakness he charges MI-WAM with is also a weakness that contextualism faces—as least when it comes to cases of seeming acceptable diverging knowledge responses where an appeal to close-calls is made. While these difficulties don't strike me as insurmountable, at this juncture, these difficulties are great enough such that the ambiguity theory retains its advantage over MI-WAM and contextualism when it comes to explaining the data from what appears to be cases of acceptable diverging knowledge responses. Because the ambiguity theory has this advantage, I think it a view worth considering seriously, and fuller comparative treatments between the ambiguity theory and the current leading accounts for the truth-conditions of knowledge ascriptions and denials are warranted.

³⁵The rationale for this position is laid out in in DeRose (1999, 198-9) DeRose (2002, 192-3). See also DeRose (2009). For responses to this line of criticism see Brown (2006) and Rysiew (2001, 2005, 2007).

Works Cited

- Brogaard, Berit. "In Defence of a Perspectival Semantics for 'Know'." *Australasian Journal of Philosophy*, 86, no. 3 (2008): 439-459.
- Brown, Jessica. "Adapt or Die: The Death of Invariantism?" *Philosophical Quarterly*, 55, (2005): 263-85.
- Brown, Jessica. "Contextualism and Warranted Assertability Manoeuvres" *Philosophical Studies* 130, no. 3 (2006): 407-35.
- Cohen, Stewart. "Is Knowledge Contextual?: Contextualism Defended" in *Contemporary Debates in Epistemology* (eds.) Matthias Steup and Ernest Sosa, Oxford: Blackwell Publishers, 2010: 55-62.
- Cohen, Stewart. "Contextualism, Skepticism, and the Structure of Reasons," in *Philosophical Perspectives, 13: Epistemology, 1999* (ed.) James E. Tomberlin, Oxford: Blackwell Publishers, 1999.
- Cruse, Alan D. "On Lexical Ambiguity," *Nottingham Linguistic Circular* 11, 2 (1982): 65-80.
- Davis, Wayne. "Knowledge Claims and Context: Loose Use." *Philosophical Studies*, 132 (2007):395-438.
- DeRose, Keith. *The Case for Contextualism*. Oxford: Clarendon Press, 2009.
- DeRose, Keith. "The Ordinary Language Basis for Contextualism, and the New Invariantism." *Philosophical Quarterly*, 55, 219 (2005): 172-98.
- DeRose, Keith. "The Problem with Subject-Sensitive Invariantism." *Philosophy and Phenomenological Research*, 68, no. 2 (2004a): 346-350.
- DeRose, Keith. "Assertion, Knowledge, and Context." *The Philosophical Review*, 111, no. 2 (2002): 167-203.
- DeRose, Keith. "Contextualism: An Explanation and Defense," in *The Blackwell Guide to Epistemology* (eds.) John Greco and Ernest Sosa, Oxford: Blackwell Publishers, 1999, 187-204.
- Fantl, Jeremy and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford: Oxford University Press, 2009.
- Fantl, Jeremy and Matthew McGrath. "Arguing for Shifty Epistemology," in *Knowledge Ascriptions* (eds.) Jessica Brown and Mikkel Gerken, Oxford: Oxford University Press, 2012, 55-74.
- Hawthorne, John. *Knowledge and Lotteries*, Oxford: Oxford University Press, 2004.

Ichikawa, Jonathan Jenkins and Matthias Steup, "The Analysis of Knowledge", *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), <<http://plato.stanford.edu/archives/spr2014/entries/knowledge-analysis/>>.

Kelp, Christoph. "Classical Invariantism and the Puzzle of Fallibilism." *The Southern Journal of Philosophy*, 46, (2008): 221-44.

Lewis, David. "Elusive Knowledge," *Australasian Journal of Philosophy*, 74, 4 (1996): 549-67.

Lewis, David. "Scorekeeping in a Language Game." *Journal of Philosophical Logic*, 8, 3 (1979): 339-359.

MacFarlane, John. "Semantic Minimalism and Nonindexical Contextualism" in *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics* Gerhard Preyer and Georg Peter (eds.), Oxford: Oxford University Press, 2009.

Macfarlane, John. "Nonindexical Contextualism." *Synthese* 166, no. 2 (2007): 231-50.

Malcolm, Norman. "Knowledge and Belief," *Mind* 61, 242, (1952): 178-89.

Perry, John. *Reference and Reflexivity*. Center for the Study of Language and Information: Stanford University, 2001.

Pynn, Geoff. "Contextualism in Epistemology" in *Oxford Handbooks Online*. New York: Oxford University Press, 2015b.

Pynn, Geoff. "Pragmatic Contextualism." *Metaphilosophy* 46, no. 1 (2015a): 26-51.

Reed, Baron. "A Defense of Stable Invariantism." *Nous* 44, no. 2 (2010): 224-244.

Reed, Baron. "Fallibilism, Epistemic Possibility, and Epistemic Agency," *Philosophical Issues* 23 (2013): 40-69.

Rysiew, Patrick, "Epistemic Contextualism", *The Stanford Encyclopedia of Philosophy* (Winter 2011 Edition), (ed.) Edward N. Zalta <<http://plato.stanford.edu/archives/win2011/entries/contextualism-epistemology/>>.

Rysiew, Patrick. "Speaking of Knowing." *Nous* 41, no. 4 (2007): 627-662.

Rysiew, Patrick. "Contesting Contextualism." *Grazer Philosophische Studien*, 69 (2005): 51-69.

Rysiew, Patrick. "The Context-Sensitivity of Knowledge Attributions." *Nous* 35, no. 4 (2001): 477-514.

Satta, Mark. "A Linguistic Grounding for a Polysemy Theory of 'Knows'." *Philosophical Studies* (forthcoming, 2017).

Satta, Mark. "A Defense of the Ambiguity Theory of 'Knows'." PhD diss., Purdue University. (2016).

Saul, Jennifer. *Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics*, Oxford: Oxford University Press, 2015.

Schiffer, Stephen. "Contextualist Solutions to Skepticism," *Proceedings of the Aristotelian Society*, 96, (1996): 317-333.

Sennet, Adam. (2011). "Ambiguity", *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/sum2011/entries/ambiguity/>>.

Stalnaker, Robert. *Context*. Oxford: Oxford University Press, 2014.

Stanley, Jason. *Knowledge and Practical Interests*, Oxford University Press, 2005.

Steup, Matthias. "Contextualism and Conceptual Disambiguation," *Acta Analytica*, 20, (2005): 3-15.

Unger, Peter. *Ignorance*. Oxford: Oxford University Press, 1975.