

The dizzying turn of epistemic contextualism

Article

Accepted Version

Grindrod, J. ORCID: <https://orcid.org/0000-0001-8684-974X>
(2020) The dizzying turn of epistemic contextualism.
Metaphilosophy, 51 (1). pp. 87-96. ISSN 1467-9973 doi:
<https://doi.org/10.1111/meta.12406> Available at
<https://centaur.reading.ac.uk/88334/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1111/meta.12406>

Publisher: Blackwell Publishing

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

The dizzying turn of epistemic contextualism

Jumbly Grindrod

Abstract

The debate concerning epistemic contextualism represents a kind of linguistic turn in epistemology, where the focus has shifted from theorising about knowledge to theorising about knowledge attributions. Such a shift may well prove valuable, but only if we are clear on what the relationship is between a semantic analysis of knowledge attributions and a philosophical analysis of knowledge. One plausible approach is to claim that the semantic analysis entails and is entailed by the philosophical analysis. Yet this view - referred to here as the *default view* - has been explicitly adopted by few in the contextualism debate. This paper considers a form of argument in favour of the default view, and then considers the challenges that arise from either accepting or rejecting the default view.

Keywords: epistemic contextualism; invariantism; semantics; metaphilosophy

1. The problem

A philosophical analysis of knowledge will at least aim to provide necessary and sufficient conditions for knowledge. As Ichikawa and Steup (2018) note, philosophical analyses often seek to do more than that. For example, a metaphysical analysis of knowledge will seek to give an account of knowledge in terms of the parts that constitute it, while a conceptual analysis of knowledge will seek to give an account of what it is to possess and apply the concept KNOWLEDGE, and perhaps a full-blown account will attempt to do all of these things. But the minimum would be that an analysis of knowledge provides necessary and sufficient conditions for knowledge in a non-trivial manner (i.e. something more than “S knows p iff S knows p”). Many will only seek to provide a partial analysis of knowledge whereby conditions for knowledge

are outlined that are not jointly necessary and sufficient, and such an analysis can of course be of tremendous philosophical value.

A semantic account of “know”, on the other hand, will provide an account of the truth-conditional content of the expression. In other words, it will provide an account of when sentences of the form “S knows p” are true and when they are false. Such accounts have, of course, generated a huge amount of discussion in recent years, particularly with regard to the contextualism/invariantism debate.ⁱ As a result, now perhaps as much as ever, epistemologists discuss epistemic analyses and semantic analyses in tandem. This can be viewed as a kind of linguistic turn within epistemology. But whether this linguistic turn towards semantics helps us with the original project of providing a philosophical analysis of knowledge depends on what exactly the relationship is between the two kinds of theory. So what is the relationship between a philosophical analysis of knowledge and a semantic analysis of knowledge attributions?

One position is to claim that the two projects are extremely closely related to one another. Any account of the necessary and sufficient conditions for when S knows p will also thereby provide an account of when it is true to say “S knows p”. Conversely, any account of the truth conditions of sentences of the form “S knows p” will entail an account of the necessary and sufficient conditions for knowledge. Let’s call this the *default view*. The default view is in many respects the natural position to hold. After all, if an agent meets all of the necessary and sufficient conditions for knowledge regarding a particular proposition, then how could it not be true to say of that agent that he knows? And if it is true to say of an agent that he knows, then surely that agent must meet all the conditions for knowledge. In this respect, the default view posits a very tight, albeit natural, link between semantic and philosophical analyses.

One might expect that those who have taken part in this linguistic turn towards the semantics of knowledge attributions would naturally hold the default view, as it provides a clear link between semantic and philosophical theorising. Yet, particularly in the contextualism/invariantism

literature, we find surprisingly few people who explicitly endorse the default view. There is instead a tendency to be clear that contextualism or invariantism is merely a semantic view about the truth conditions of knowledge attributions. For example, DeRose considers the possible objection that his contextualist view is not relevant to epistemological theorising because it is merely a linguistic claim:

Sadly and all too commonly, this has been known to give rise to the following type of outburst: “Your contextualism isn’t a theory about knowledge at all; it’s just a theory about knowledge attributions. As such, it’s not a piece of epistemology at all, but of the philosophy of language. (DeRose 2009, 18)

DeRose responds to this objection by accepting that contextualism is strictly a view in philosophy of language, but he argues that linguistic facts are sometimes crucial to philosophical theorising because they can provide us with ways to solve certain philosophical puzzles (he particularly has in mind the sceptical arguments that contextualism is thought to provide an attractive explanation of). But notice that, if DeRose held the default view, the response to this objection would have to be quite different. DeRose would simply respond that in giving an account of when ‘S knows p’ is true or false, he is also thereby providing an analysis of what knowledge is.

In more recent work, DeRose comes closer to addressing this question of how semantic analyses could be relevant to philosophical analyses:

I believe that separating out the factors which affect whether claims involving philosophically important terms are true or false from when such claims have other properties that might be confused with truth and falsity is an important tool for understanding the meanings of those key terms [...] and that, in turn, while it is certainly very far from all that is involved, understanding the meaning of the key terms involved – or perhaps more precisely, avoiding radical and

profound misunderstandings of those meanings – is important to addressing philosophical problems. (DeRose 2017, 106)

He then continues in a footnote:

I am sometimes mentioned as a current philosopher (such creatures were apparently more plentiful, or at least more prominent, in earlier times) who has “taken the linguistic turn” in philosophy – sometimes as a particularly clear example of such. I’m not very confident in my understanding of this description, but I suppose that the sentence to which this note attaches would be the kind of spot to look at in deciding whether the description fits me. I’m guessing that my inclusion of the limiting phrase “while it is certainly very far from all that is involved” may harm my candidacy. But if thinking that understanding the meaning of the key terms involved, and especially avoiding profound misunderstandings of the meanings of the key terms involved, can be important to addressing philosophical problems suffices for one to count as having taken the “turn”, then I would certainly plead guilty”. (DeRose 2017, 106, fn. 21)

Although DeRose does attempt to be more explicit here on what the relationship is between semantic and philosophical analyses, much is still left unsaid. As he notes, he is certainly not committed to the idea that philosophical theorising only ever amounts to providing a linguistic analysis of the relevant terms, but he is committed to the idea that understanding the semantics and pragmatics of key terms “can be important to addressing philosophical problems”, although nothing is said specifically about why. Again, if DeRose were to hold the default view in this instance, he could be far more specific here – he could state that providing the truth conditions of knowledge attributions will entail necessary and sufficient conditions for knowledge.

We also find reluctance to accept the default position from the other side of the epistemological coin, from those providing a theory of knowledge. For example, Goldberg has recently provided

a novel account of *epistemic propriety* as a necessary condition for knowledge and which he takes to mirror certain semantic proposals:

Although I am not offering my proposal as a semantic one – my claim concerns knowledge and epistemic propriety in the way of belief, not the semantics of the expressions ‘knowledge’ or ‘epistemic propriety’ – even so, the position mirrors SSI in that there are contextual parameters that fix the standards by which subjects are assessed, even as the property under assessment is not regarded as a relation to the context. (Goldberg 2018, 241)

In saying this, Goldberg indicates that he does not accept the default view, because if he were, then his proposal of epistemic propriety would directly entail the mirroring semantic proposal. Instead, he intends to leave it open whether mirroring semantic proposal is the right one. We find, then, that there is a resistance to the default view both from those, such as DeRose, who are defending a semantic view, and from those, such as Goldberg, who are defending a particular analysis of knowledge.

Yet in other areas of the debate, something like the default view seems to be implicitly assumed. To see this, consider the invariantist view whereby the truth conditional content of knowledge attributions remain invariant across contexts of utterance. As such, it is a linguistic view (as Goldberg rightly notes). Yet some of the most prominent views that fall under the invariantist banner are first-and-foremost views about the metaphysical structure of knowledge. For example, Stanley’s (2005) interest-relative invariantism is the view that knowledge is a ternary relation between a knower, a proposition, and the knower’s practical interests in being right about that proposition, or put less metaphysically, whether someone knows depends in part on the practical interests they have in their belief being true. But this is part of an analysis of knowledge, not a semantic account of knowledge terms. So in order for this to be a form of

invariantism about knowledge attributions, it seems that we have to assume that the semantic facts will reflect the nature of knowledge much in the way that the default view suggests.

I suggest that there is a fundamental lack of clarity and consistency across the contextualism/invariantism debate regarding the relation between a semantic analysis of “know” and an epistemological analysis of knowledge. This issue is clearly crucial to understanding what is at stake in this debate, and so further attention needs to be paid to this issue. As the default view represents a natural starting position that is easy to outline, I suggest that the best way to approach this question is by considering the plausibility of the default view. In the next section, I will confront a form of argument that could be used to establish the default view. Then, in section 3, I will map challenges faced by those who reject and those who accept the default view.

2. A Disquotational argument

To see a form of argument that could be used to defend the default view, let’s first assume that knowledge attributions are not context-sensitive. If that is the case, then we can provide a simple argument that appeals to a disquotational principle and establishes a biconditional relationship between philosophical and semantic analyses:

1. *Disquotational principle*: ‘S knows p’ is true iff S knows p.
2. *Semantic analysis*: ‘S knows p’ is true iff α .
3. Therefore, S knows p iff α .

Here α acts as a placeholder for whatever semantic analysis of knowledge attributions is given.

The point of this argument form is that once we accept the disquotational principle outlined in 1 – which is an instance of Tarski’s T-schema – then we can infer necessary and sufficient conditions for knowledge from the truth conditions of knowledge attributions.ⁱⁱ And of course, this argument could be run in the other direction – once you have necessary and sufficient conditions for knowledge then you can use the disquotational principle to deduce truth

conditions for sentences of the form “S knows p”. In this way, the disquotational principle can be used to argue in favour of the default view.

However, this argument form is not available if contextualism is true. Disquotational principles do not hold for sentences containing context-sensitive items. For example, “‘I am a man’ is true iff I am a man” is a false biconditional simply because the two tokens of ‘I’ may not have the same referent. The same applies for “know”: the contextualist will claim that 1 is false because the two tokens of ‘know’ may refer to different epistemic relations. Equally, a contextualist will reject the idea that you can give a semantic analysis in the form of 1. Instead, they will argue that a semantic analysis of the truth conditions of knowledge attributions must make some appeal to the context of utterance i.e. “‘S knows p’ is true in c iff α ” (where α will also have to make some appeal to c).

Perhaps it is the unavailability of the disquotational argument that makes those who discuss contextualism so reluctant to endorse something like the default view. This would partially explain why invariantists seem willing to assume the default view while contextualists are less willing. Yet this in itself is a surprising result. This would then mean that the invariantist is committed to a particular picture of the relationship between semantic and philosophical theorising that the contextualist can avoid. If you are an invariantist, in giving the semantic analysis that you do, you are also thereby giving a philosophical analysis of knowledge (and vice versa). Whereas if you are a contextualist, you can avoid the disquotational argument and thus avoid any commitment to necessary and sufficient conditions for knowledge. However, if this is the current state of play, this needs to be made explicit within the literature.

At the same time, it might be thought that while the disquotational argument cannot be applied to the contextualist view, there may be a similar form of argument that does the same job. In particular, if we can replace the disquotational principle with some similar form of principle, then we may be able to establish the link in a similar way. In order to outline what I have in mind, I

will use some set-theoretic parlance. Let's represent the knowledge relation – as in, the thing that is the subject of epistemological analysis – as the set K containing ordered pairs of knowers and propositions. Now, if contextualism is true, it seems that we talk of different knowledge relations in different contexts, and so the natural way to link up semantic and philosophical analysis is to allow for many different knowledge relations in our epistemology. So rather than having a single relation K , we could have relations $K_1 \dots K_n$, which are indexed appropriately to contexts $c_1 \dots c_n$.ⁱⁱⁱ Using this terminology, we could then outline a biconditional principle similar in form to the disquotational principle:

4. $(S,p) \in K_n$ iff 'S knows p' is true in c_n .

This principle takes us from the fact that some knowledge relation holds between S and p to the truth of a knowledge attribution within a particular kind of context, and back again. We could then use this principle to form a similar kind of argument to the disquotational argument in the following manner:

5. $(S,p) \in K_n$ iff 'S knows p' is true in c_n .
6. *Semantic analysis*: 'S knows p' is true in c_n iff α .
7. Therefore, $(S,p) \in K_n$ iff α .

Just as with the disquotational argument, we can of course run this argument the other way by replacing the semantic analysis with a philosophical analysis and thereby deduce a semantic analysis. This way, we could provide an argument in favour of the default view that allows for the context-sensitivity of knowledge attributions.^{iv}

The important difference, however, between 5-7 and the disquotational argument is that the disquotational principle is independently motivated as being an instance of the T-schema, which itself is supposed to provide some account of what it takes for a sentence to be true. 5, on the other hand, has no similar kind of independent backing, and so it really just stands as a way of

outlining the default view. But if some way of defending 5 could be found, this would constitute a way of arguing in favour of the default view from the contextualist perspective.

3. The default view: reject or accept?

Whether we accept or reject the default view is crucial to understanding what is at stake in the contextualism debate. As I will argue, both options face challenges, and so this can be seen as something of a dilemma within this debate. I don't believe, however, that this is an insurmountable dilemma. But whichever horn is chosen, it will lead to repercussions in the way that semantic analyses are viewed in epistemological theorising.

3.1. Accepting the default view

First, the default view could be accepted. This could of course, be due to the arguments outlined earlier, or it may just be because the default view is considered the most plausible view of the relation between semantic and philosophical theorising. But of course, accepting the default view means accepting that there is this incredibly close, biconditional relationship between semantic and philosophical analyses. A philosophical analysis would entail a semantic analysis and vice versa. So when DeRose argues for contextualism, he would effectively be arguing for a certain epistemological view regarding the necessary and sufficient conditions for knowledge. And when Goldberg argues for a particular account of epistemic propriety, he would also effectively be arguing for a particular semantic analysis.

We might think that this picture of the relationship between semantic and philosophical analyses is a little too close. After all, the two projects seem to differ in their methodology, as Chalmers notes:

...linguists' analyses often aim to capture the inferences that are held to be valid by most or all competent users, whereas a philosopher might instead aim to capture inferences that are in fact valid. Likewise, a linguist might freely invoke

abstract objects in their analyses, while some ontologically scrupulous philosophers might not. A philosopher might also give a role to empirical facts about science and naturalness in giving their analyses, while a linguist might not.

(Chalmers 2012, 6)

With this in mind, the challenge facing the default view is in allowing for this incredibly tight relationship between philosophical and semantic theorising, even when it seems that there are in fact such great differences between the two projects. For example, a semanticist giving an account of knowledge attributions will be sensitive to syntactic considerations in a way that an epistemologist will not be, and an epistemologist will be sensitive to considerations of epistemic normativity and value in a way that the semanticist typically will not be.

Perhaps, then accepting the default view requires a form of revisionism about how to proceed in semantics or in epistemology, or in both. However, as a friendly suggestion, I will outline one way that the default view could be accepted and still allow for important differences between the two projects. We could allow that the correct semantic analysis of “S knows p” will in turn provide necessary and sufficient conditions for knowledge, while still being committed to the idea that further philosophical work has to be done. This is because, as Chalmers indicates, the semanticist may rely on notions in their analysis that the philosopher would be content with leaving as given. So a semanticist could claim that “S knows p” is true iff Fa , where “F” is some predicate and “a” is some singular term. According to the default view, we could then deduce that S knows p iff Fa . But a complete philosophical explanation may well require that we further investigate the nature of the property that “F” denotes or the object that “a” denotes to a greater degree of specificity. And the same may hold in the other direction as well, where a complete philosophical analysis would rely on notions that a complete semantic analysis would need to unpack further.^v Allowing this much would allow that the aims and methodologies of the two projects differ even if there is a biconditional relationship between philosophical and semantic

analyses. But whether this would completely explain the differences that Chalmers outlines above I will leave for further consideration.

3.2. Rejecting the default view

The second alternative is that the default view is rejected. We saw earlier that invariantists can appeal to the disquotational argument in defence of the default view, and so the immediate challenge for invariantism is to reject the disquotational argument. And insofar as 4 is a plausible principle (which I'm sure for many, it is), the contextualist likewise has to reject the argument from 5-7. But regardless of the particular view, rejecting the default view means rejecting the intuitive plausibility it possesses. As stated earlier, the default view is attractive insofar as it takes each project – philosophical and semantic – at face value. The semantic project is providing truth conditions for knowledge attributions, and to do this properly it does seem as though you have to provide some account of when a subject knows and when they don't, which ultimately is to provide conditions for knowledge. Instead, a rejection of the default view has to paint some alternative picture on the relation between philosophical and semantic analyses. In particular, it seems to open up the possibility of there being particular cases in which $(S,p) \in K_n$, but in which it would be false to say "S knows p", in c_n . Furthermore, in rejecting the default view, the risk is run of letting "semantics float free from metaphysics" (McKenna 2015, 501) and allowing for no real connection between semantic facts about knowledge attributions and an epistemological account of knowledge. This would then be to claim that the linguistic turn made by contextualists and others was not merely dizzying but entirely wrong-headed insofar as semantic analyses of knowledge attributions are in fact irrelevant to epistemological theorising.

Just as I put forward a suggestion for how the advocate of the default view could meet their challenge, I will finish by providing the briefest of sketches of how an alternative view could be developed. Consider the following famous quote from Austin on the use of ordinary language in philosophical theorising:

Certainly ordinary language has no claim to be the last word, if there is such a thing. It embodies, indeed, something better than the metaphysics of the Stone Age, namely, as was said, the inherited experience and acumen of many generations of men. But then, that acumen has been concentrated primarily upon the practical business of life. If a distinction works well for practical purposes in ordinary life (no mean feat, for even ordinary life is full of hard cases), then there is sure to be something in it, it will not mark nothing: yet this is likely enough to be not the best way of arranging things if our interests are more extensive or intellectual than the ordinary. And again, that experience has been derived only from the sources available to ordinary men throughout most of civilised history: it has not been fed from the resources of the microscope and its successors. And it must be added too, that superstition and error and fantasy of all kinds do become incorporated in ordinary language and even sometimes do stand up to the survival test (only, when they do, why should we not detect it?). Certainly, then, ordinary language is *not* the last word: in principle it can everywhere be supplemented and improved upon and superseded. Only remember, it *is* the *first* word. (Austin 1956, 11)

The ordinary language school are often held to have placed too much weight on the properties of natural language, yet here we find that Austin conveys a cautionary yet optimistic attitude to the study of language for philosophical purposes. Following from Austin we could view the evidence provided by language as defeasible, as something that can be overridden if it makes for a superior epistemological theory. In the case of knowledge attributions, we could view the truth conditions of knowledge attributions as defeasible evidence in favour of a particular analysis of knowledge, but which can be overridden if there strong enough philosophical considerations to do so. This would this make good of the idea that the philosophical and semantic projects are distinct, complete with different methodologies, while also allowing that drawing upon the

semantics of knowledge attributions in seeking to understand the nature of knowledge has not been a wrong-headed turn.

4. Summary

In this paper I have argued that there is a fundamental lack of clarity in the contextualism/invariantism debate regarding the relationship between semantic and epistemological theorising. I have outlined what I call the default view as a natural starting position, and I have assessed arguments in favour of the default view. Finally, I have laid out the challenges that arise from either rejecting or accepting the default view.

Dr Jumbly Grindrod

Philosophy department

Whiteknights campus

University of Reading

RG6 6AH

[*j.grindrod@reading.ac.uk*](mailto:j.grindrod@reading.ac.uk)

References

- Austin, John Langshaw. 1956. "A Plea for Excuses." *Proceedings of the Aristotelian Society* 57:1--30.
- Chalmers, David. 2012. "Eighteenth Excursus: Conceptual Analysis and Ordinary Language Philosophy." In *Constructing the World*. Oxford: Oxford University Press.
- DeRose, Keith. 2009. *The Case for Contextualism*. Oxford: Oxford University Press.
- DeRose, Keith. 2017. *The Appearance of Ignorance: Knowledge, Scepticism, and Context*. Oxford: Oxford University Press.
- Goldberg, Sandford. 2018. *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*. Oxford: Oxford University Press.
- Hansen, Nat. 2014. "Contemporary Ordinary Language Philosophy." *Philosophy Compass* 9/8:556-569.
- Ichikawa, Jonathan Jenkins, and Matthias Steup. 2018. "The Analysis of Knowledge." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta (Winter 2019 edition) <https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/>.
- Jary, Mark, and Robert J. Stainton. 2017. "Contextualism in Epistemology and Relevance Theory." In *Routledge Handbook of Epistemic Contextualism*, edited by Jonathan Jenkins Ichikawa. London: Routledge.
- McKenna, Robin. 2015. "Contextualism in Epistemology." *Analysis* 75 (3):589-503.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.

ⁱ Throughout the remainder of this paper, I will assume that contextualism and invariantism are semantic views that are concerned with the truth conditions of knowledge attributions. However, some have argued for a form of contextualism according to which the context-sensitivity in

question is a pragmatic phenomenon (Jary and Stainton 2017). I put these views to one side for the purpose of this paper.

ⁱⁱ This argument form is discussed in (Chalmers 2012, pp. 4-7) and (Hansen 2014, pp. 561-2)

ⁱⁱⁱ Contexts here are to be understood as abstract types of situation. A full account of contexts might represent them as n-tuples containing every contextual feature required to assign values to context-sensitive items. For my purposes here, we are only dealing with one context-sensitive item (“know”), and so contexts are differentiated purely by the knowledge relation they are indexed with.

^{iv} I said that a natural way to capture the contextualist claim that knowledge attributions are used to talk about different relations in different contexts is to posit a range of relations $K_1 \dots K_n$. An alternative would be to claim that there is a single knowledge relation K , but that it is a ternary relation between a subject, a proposition, and a context (or some contextual feature). On this account, you could still outline a principle akin to 4 in the following way: “ $(S, p, c_n) \in K \leftrightarrow$ ‘S knows p ’ is true in c_n .”

^v By allowing for this, I am remaining neutral as to whether there is any kind of priority between semantics and philosophy.