Episteme (2025), 1–14
doi:10.1017/epi.2025.10087

episteme

ARTICLE

Dretskean Sensitivity and Higher-Level Knowledge

Kelly Becker

University of New Mexico, Albuquerque, NM, USA
Email: kbecker@unm.edu

(Received 17 January 2025; revised 8 September 2025; accepted 17 September 2025)

Abstract

The sensitivity principle in epistemology has faced numerous, considerable, and relentless challenges since it emerged in Nozick's *Philosophical Explanations* (1981). In this paper, I develop a version of sensitivity, based on Dretske's notion of conclusive reasons (1971), that responds to the complaint that sensitivity is either incompatible with or makes an unprincipled mess of higher-level knowledge. There are three key moves in formulating reasons-based Dretskean sensitivity (RDS). First, sensitivity is conceived in terms of reasons, rather than beliefs, that track the truth. Second, focus shifts from whether S would have those reasons in the relevant counterfactual worlds to whether those reasons would be the case. Third, closer attention is paid to the structure of reasons. Critics of Nozick point out that, typically, even when S knows that they do not have a false belief that p, if S were to have a false belief that p, S would nonetheless believe that they do not have a false belief that p, violating Nozickean sensitivity. I explain how this fact does *not* preclude higher-level knowledge according to RDS, even if the false belief that p were based on their actual method.

Keywords: Dretske; higher-level knowledge; Nozick; sensitivity; tracking

Sensitivity-based tracking accounts of knowledge have faced three principal sources of intense resistance ever since their most well-known elaboration and defense in Nozick (1981). The first and perhaps most notorious problem is that sensitivity, construed as a necessary condition for knowing, implies, absent further remedial conditions, that knowledge is not closed under known entailment. A second problem is that Nozickean sensitivity generates incorrect implications concerning inductive knowledge. The third problem is that sensitivity also generates incorrect implications regarding higher-level knowledge – specifically, knowledge that one does not believe falsely. Early articulations of this problem claim that Nozickean sensitivity is simply incompatible with higher-level knowledge. This is the *core* problem about (/objection to sensitivity and) higher-level knowledge. Later articulations claim that, even if (properly construed) sensitivity is compatible with some cases of higher-level knowledge, it is implausibly incompatible with others. Worse still, sensitivity is counterintuitively compatible with knowledge of

strong but not weak higher-level propositions. This is the *heterogeneity* problem (Melchior, 2014, 2015, 2017).

Here I aim to resolve the problems that higher-level knowledge presents for sensitivity.[1] The key, following Dretske (1971), is to shift focus from the questions of whether S's *belief* is sensitive or whether S's *method of forming belief* is sensitive to the question of whether S's *reasons* are sensitive. A common dialectic has emerged with regard to all three of the problems mentioned. First, it is argued that belief sensitivity generates a problem. Second, the response is that closer attention to belief-forming methods is key to answering the problem. Third, the rejoinder is that the cure is as bad as or worse than the disease.[2] In addition, once the pandora's box of belief-forming methods is flung open, broader problems emerge.[3] A better approach is to ask whether one's actual reasons are sensitive to the truth that p, regardless of whether, if p were false, one's belief would be *based on* those reasons, or one would *believe* that p, or one would have used their actual belief-forming *method*.

In §I, I set out the core and heterogeneity problems for sensitivity with regard to higher-level knowledge. I then develop Dretskean sensitivity (§II), which requires taking a stand on at least two key issues left unresolved in Dretske's presentation. I then explain how Dretskean sensitivity can handle the higher-level knowledge problems (§III).

1. Higher-level knowledge problems for sensitivity

Nozick first sets down sensitivity of belief as a necessary condition for knowing. S knows that p only if, were p false, S would not believe that p (1981, 172ff. for development).[4]

Belief sensitivity generates the core objection (Sosa 1999, 145; Vogel 2000, 610-11). Suppose that S knows that q. Surely, then, S sometimes knows p: *that [S's] belief that q is not false*. However, if p were false, S would have a false belief that q. But because S would believe that q, S would believe that they do not have a false belief that q, violating sensitivity. The core objection is that belief sensitivity is *logically incompatible* with higher-level knowledge. I pause to emphasize that, obviously, there is logically no way out of this problem, formulated this way, for belief sensitivity.

In a footnote, Vogel points out that the problem involves closure failure, since q entails that the belief that q is not false. (Vogel puts it like this (substituting my "q" for his "X"): q entails (¬(B(q) & ¬q) (2000, n. 15).) As mentioned above, many sensitivity theorists accept closure failure – a point whose relevance to these proceedings is

[1]I believe that the approach taken here, which draws on Dretske's (1971) notion of conclusive reasons, can also help with the problems concerning inductive knowledge and perhaps even with the closure problem, but those are matters for another day.

[2]On closure: Kripke's (2011) version of the fake barns objection – S knows there's a red barn but does not know there's a barn – is thought to be devastating. Adams and Clarke (2005) claim that closer attention to methods solves the closure problem and shows that S knows both. Becker (2023, n. 13) complains that S knows neither! On induction: Vogel (1987, 2000) and Sosa (1999) argue that sensitivity is incompatible with induction. Wallbridge (2018) says that closer attention to methods solves the problem. Melchior (2021) complains that the strategy creates a new problem that is at least as bad because the implications are heterogeneous.

[3]See Alfano (2009), B. Zhao (2024, 2025), and H. Zhao (2023) for various developments of the complaint that there is no way of incorporating methods into the sensitivity condition or individuating methods that doesn't generate untoward consequences.

[4]Nozick used the term "variation" for what is now commonly referred to as sensitivity, reserving the term "sensitivity" for the conjunction of his variation condition and another necessary condition, adherence, which, when first presented, is said to offer "additional sensitivity" (1981, 211). I won't discuss adherence here.

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

addressed further below, §III – but even for those willing to countenance closure failures, the logical incompatibility of belief sensitivity and higher-level knowledge is surely unacceptable. That is, it is one thing to deny closure and say that one is not always in a position to know that q, based on one's knowledge that p and that p entails q; it's another to say that knowing that q, for a specific class of q, is *logically incompatible* with knowing that p and that p entails q.

Becker (2006) and Salerno (2010) respond by pointing out that belief sensitivity is compatible with higher-level knowledge, construed in a different way. Again, suppose that S knows that q. S can know p* that they have a true belief that q. How? If it were false that S has a true belief that q, either (a) S would not believe that q or (b) q would be false. On (b): Given that S knows that q, S would not believe that q if it were false. Either way, (a) or (b), S would not believe that q, and so S would not believe p* that they have a true belief that q.

One concern is that, even if the response reconciles belief sensitivity with some kinds of higher-level knowledge, it remains deeply problematic that one could never know that they do not have a false belief – as in the original formulations of the core problem. Melchior (2014, 2015) rubs this in a little. He points out that to know that one has a true belief is to know the conjunction *c*: B(p) & p. Not knowing that one's belief is not false is failing to know the disjunction *d*: ¬(B(p) & ¬p), i.e. ¬B(p) v p. Whereas Becker (2006) and Salerno (2010) suggest that the former is a more faithful rendering of the relevant higher-level knowledge because it alone correctly implies that S (actually) believes that p, Melchior points out that the result is still embarrassing because disjunctions are weaker propositions and ought to be easier to know. Melchior labels this sort of problem "heterogeneity" not only because of the mixed implications, but also because the implications are *counterintuitively* mixed.[5]

Just a few pages after laying out belief sensitivity, however, Nozick presents the following well-known example to show that the method by which one comes to believe must be taken into account. A grandmother sees that her grandson is well when he comes to visit. But if he were ill, his parents would tell her that he is at home and doing fine, so she would believe that her grandson is well even if he weren't. The problem is that, while surely grandma knows that her grandson is well, her belief is insensitive – she would believe he is well even if he were ill.

Nozick responds to the problem by claiming that, in assessing whether a belief is sensitive, we must hold the believer's actual belief-forming method constant. If it were false that the grandson is fine, grandma would not believe that he is fine via her actual method of seeing his healthy-looking visage. Wallbridge (2016) argues that taking methods into account is key to reconciling sensitivity and higher-level knowledge and, later (2017), that the approach can also resolve the heterogeneity problem. Melchior claims that the heterogeneity only gets worse on Wallbridge's method-relative belief sensitivity approach.

> Let us begin to investigate this dialectic with a revised Nozickean sensitivity condition:
>
> RNS: S knows that p only if, were p false, S would not believe that p via [S's actual belief-forming method] M.

[5]Bjerring and Gundersen (2020) agree with Becker and Salerno that the conjunctive formulation is correct, and they use an appeal to methods, an issue to which I shall come presently, to justify this. While I agree with the gist of this position, my aim is to take the bull by the horns and respond directly to the core problem by showing that one can *sensitively* know that it is not the case that one has a false belief.

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

Wallbridge argues that RNS is compatible even with knowing the weaker, disjunctive higher-level proposition *that it is not the case that one has a false belief*. He returns to Vogel's (2000, 609) original example, where S sees Omar, "a perfectly decent and straightforward sort of person," wearing shiny white shoes, and Omar tells S that he has new shoes. S knows that Omar has new shoes, but, Vogel argues (as we have already seen), S does not know that they do not have a false belief. Wallbridge contests this conclusion in the following way. S believes that Omar has new shoes on the basis that Omar says so and they look new. This is then also part of the basis for their higher-level belief that they do not believe falsely. Given that Omar is *trustworthy*, it follows that "it would not easily have been the case that Omar lied to you or was otherwise mistaken" (Wallbridge 2016, 168). Therefore, if S were to believe falsely that Omar has new shoes, "this would have been on some other, less reliable, basis" (*ibid*.). Otherwise put, the claim is that worlds where S believes falsely by some other method are closer than those where S believes falsely by their actual method. In that case, even though S would falsely believe that their belief is not false, their higher-level belief would not be formed from their actual method, hence would not violate sensitivity, construed in a method-relative way. Wallbridge (2017) then uses this strategy to argue that Melchior (2014, 2015) is wrong to say that sensitivity has a heterogeneity problem with higher-level knowledge.

Melchior responds by formulating "Wallbridge's condition" and arguing that it still implies unwanted heterogeneity.

> Wallbridge's condition: S's belief that she does not falsely believe that *p* formed via method M is sensitive iff . . . *W* is fulfilled. *W*: Possible worlds where S falsely believes that *p* via some other method M* are closer than the nearest possible worlds where S false believes that p via M. (Melchior 2017, 5).

Melchior constructs cases where another method M* is available to S only in worlds that are further out than those where S uses M and forms a false belief. The recipe is simply to construct examples where *any* other method would be hard to come by. Melchior (*ibid*, 3-4) reimagines Omar as Robinson Crusoe stranded on an island alone with Freitag, playing the role of S, such that Freitag has no other way to know that Crusoe has new shoes except through Crusoe's testimony. The upshot is that some higher-level beliefs in the "weaker" disjunctions are sensitive and some are not, even though, "Epistemically, there is no crucial difference between [the] cases" (*ibid*., 5). This only exacerbates the heterogeneity problem.

To sum up this section, Becker's and Salerno's approach makes logical space for sensitive higher-level belief (that one has a true belief) in response to the core problem, and Wallbridge's approach responds both to the core and to the original heterogeneity objections by showing how higher-level beliefs (that one does not have a false belief) can be sensitive on a method-relative belief sensitivity account (RNS), assuming that, if one's belief were false, one would not have formed it by one's actual method. Melchior effectively says, "But that's not always true," and on that basis reissues the heterogeneity concern. A more satisfying dissolution of this updated version of the heterogeneity problem – which essentially asks why it should be that S can achieve higher-level knowledge that p if one's first-level belief that q would have been formed from a different method if p were false, but not otherwise – would explain how one can have higher-level knowledge that p even if one were to use the same method if p were false. It would, of course, also answer the core objection.[6]

[6]It would also forestall a potential concern about the required reading of the counterfactual – if p were false, S would have used a different method in believing q – which is analogous to worries about backtracking

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

A bit of reflection on the preceding reveals that, in pursuit of this kind of response, construing the sensitivity condition either in terms of belief or of method-relative belief is a dead end. However, when the question of sensitivity shifts from whether, if p were false, one would believe that p, or one would believe that p via one's actual method, to whether one's reason for believing that p would be the case, a more plausible and satisfying sensitivity-based account of higher-level knowledge emerges. I turn to that account now and then show how it solves the problem(s) of higher-level knowledge.

2. Reasons-based Dretskean sensitivity

We begin with Dretske's initial presentation of a conclusive reason R:

1) S knows that P and he knows this on the basis (simply) of R entails
2) R would not be the case unless P were the case (Dretske 1971, 1).

We set aside the question of whether having such a reason is also sufficient for knowing. Dretske discusses this in detail and concludes with a qualified "yes" (*ibid*., 16ff.). We shall treat such reasons only as necessary for knowing, since that is the source of the higher-level knowledge problem.

For the remainder of this section, I will draw on Dretske's own words to develop an account of reasons-based sensitivity, though I do not claim that I am interpreting Dretske as he intended. I aim to be clear but relatively brief, since I have presented a fuller explication of Dretskean reasons-based sensitivity in another paper (Becker MS).

There is an ambiguity in Dretske's initial statement of the nature of conclusive reasons (2) and his first gloss on them: "For if (2) is true, we are entitled . . . to deny that . . . given R . . . not-*P* might be the case" (*ibid*., 1). A fairly straightforward reading construes this as a safety condition on knowing: *If R were true, p would be true*. But later, in his treatment of the lottery problem, where one does not know that one's ticket will lose, even given its enormous likelihood, Dretske characterizes a conclusive reason as one that is sensitive: "Adjusting the word order . . . to suit the example we have . . . If S were going to win the lottery [if it were false that S will lose], his chances of winning would not be 1/m (m being the number of tickets sold)" (*ibid*., 4). This is false, which is why S does not know that they will lose the lottery. Here is the sensitivity interpretation of a conclusive reason: *If it were not the case that p, R would not be the case*.

From here on, I take the liberty of construing Dretske's notion of a conclusion reason as one that is sensitive. I do not claim this was clearly his intention. Rather, my aim is to show that doing so positions us to see how Dretskean sensitivity, or at least my development of it, generates acceptable implications regarding higher-level knowledge where RNS does not.[7]

readings of counterfactuals: If p were false, something else would *have to* have been different. (If S falsely believed that p, something else would have been the case, namely, S would have used a different method.) Maybe, maybe not. In another debate about sensitivity, involving the exact same players but on the topic of induction, Wallbridge also defends the appeal to backtracking interpretations of counterfactuals, at least where appropriate. Notably, in the induction debate, a dialectic similar to the one in this paper plays out. Vogel (1987) and Sosa (1999) claim that sensitivity is incompatible with inductive knowledge. Wallbridge (2018) appeals to methods and potentially controversial readings of counterfactuals to save sensitivity from the incompatibility charge. Melchior (2021) responds that the resulting implications are wildly heterogeneous. I introduce RDS (Becker, MS).

[7]"But why sensitivity *at all*? Why not safety?" The brief answer is that sensitivity implies a discrimination requirement on one's knowledgeable reasons that safety does not, which I find welcome and illuminative of the nature of knowledge, regardless of whether sensitivity can diagnose every Gettier case and eliminate

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

Interpreting conclusive reasons as sensitive only tells us something about what makes them *conclusive*, not what they are *qua* reasons or how they contribute to S's knowing that p. To these ends, Dretske elaborates:

> S *has conclusive reasons*, R, for believing P if and only if:
> A. R is a conclusive reason for P [R would not be the case unless p],
> B. S believes without doubt, reservation, or question that P is the case, and he believes this on the basis of R,
> C. (i) S knows that R is the case or
> (ii) R is some experimental [sic] state of S . . . (*ibid*., 12-13)

Two small points. First, on (C)(ii), the overall context of the discussion suggests that it should read "experiential" rather than experimental. Second, since our main interest is in sensitive reasons as a necessary condition for knowledge, we set aside any question about Dretske's characterization of the *strength* of S's belief (B), though S must believe that p on the basis of R.

What kind of thing is R? R is either a proposition (C)(i) or an experiential state (C)(ii) which, when conclusive, stands in a specific modal relation to some known proposition – perhaps p, perhaps some other reason in the inferential chain leading to p. (I'll come back to this, which is important.) To *have* a conclusive reason R to believe that p is for R to be the causal basis of one's belief that p, for example when either S believes that R (occurrently or dispositionally) or when R is "the way things look" (*ibid*., 15) on the basis of which S believes that p (clause (B)).[8] When R is S's reason for p, S may not "be in a position to *give* R as his reason . . . He may find it difficult or impossible to give verbal expression to R. He may have forgotten R" (*ibid*., 14). Whatever the explanation, R need not be cognitively accessible to S. Distinguishing the questions whether R is a conclusive reason (A) for p – that is, whether R is sensitive to or tracks p – and whether S would believe that p on the basis of R if p were false is important for understanding the advantages of Dretskean sensitivity over RNS. We are not concerned with whether S *would have* R if p were false, but if R would be the case.

Another interesting question arises from the two subclauses under (C). To the concern that the appearance of "know" in (C)(i) is circular, Dretske says the circularity "can be eliminated by recursive applications of the three conditions until (C) (ii) is reached" (*ibid*., 13). Dretske requires that, whatever one's reasons, they are ultimately grounded in experience. One's basic reasons, then, are like Nozick's methods – dependent on "upshot in experience" – whether or not one is consciously aware of that upshot.

Importantly, however, the structure of Dretske's account of conclusive reasons can be exploited to solve problems that afflict Nozickean methods. Dretske's characterization of having conclusive reasons is recursive with a "foundationalist" structure, of which experience forms the base. This allows that one's belief that p can be based on R_1, that R_1 can be based on $R_2, \ldots$, that R_{n-1} can be based on R_n, and (finally) R_n on an experiential reason R_e. This in turn opens the possibility that one

every salient element of epistemic luck. While it's true that Nozick is motivated by Gettier cases, the discrimination insight stands by itself. H. Zhao (2022, 2023) is right to say that sensitivity doesn't fully live up to the Gettier motivation and more generally to a complete and correct analysis of knowledge (what theory does?) but wrong to conclude that it lacks motivation.

[8]I assume that Dretske is only interested in requirements for knowing that p and thus ignores reasons one might have that are not part of a causal basis for belief.

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

knows that p on the basis of R_1, knows that R_1 on the basis of R_2, R_2 on the basis . . . of R_e, such that each reason is sensitive to (or simply "tracks") what it is a reason for, *without its being the case that* R_e *is sensitive to, tracks, p*. This is another ambiguity in Dretske's account, which again I will exploit for my purposes. One *can* read Dretske as saying that knowledge that p requires one's basic experience to track p. However, *not* doing so is also consistent with his text, and it opens a promising avenue for sensitivity theorists.

I mentioned at the outset that focusing the question of sensitivity on reasons rather than beliefs and dismissing the question of whether S would *have* (conclusive) reasons R in a counterfactual situation in favor (simply) of whether R would be the case constitute the two central advances of Dretskean sensitivity over Nozickean sensitivity (even in the latter's best higher-level-knowledge-friendly version, RNS). The final key move is to resolve the aforementioned second ambiguity in our favor and explicitly unpack the structure of reasons underlying a single higher-level belief conclusion to reveal the possibility that S can achieve higher-level knowledge that p (that they do not have a false belief) without S's actual experience tracking p, which is anticipated neither in Nozick, nor, from what I can see, among his defenders.[9] In the following section, I illustrate how reasons-based Dretskean sensitivity (hereafter RDS) handles the core problem and the heterogeneity problem.[10]

3. RDS's more satisfying account of higher-level knowledge

To set the stage, recall first that Becker and Salerno have already shown how one can have higher-level knowledge in cases where, if it were false that one has a true belief that p, one would not believe that p. But this does not reconcile sensitivity with knowing that it is not the case that one has a false belief, which is the core problem and, together with the Becker and Salerno approach, raises the original heterogeneity problem. Recall second that Wallbridge has shown how to reconcile RNS sensitivity with knowing that it is not the case that one has a false belief, *provided* that in the closest worlds where one does have a false belief, one uses a different belief-forming method, raising another heterogeneity problem.

An adequate response to the core problem will also alleviate Melchior's concerns about heterogeneity. If it can be shown how one can have a sensitive higher-level reason for believing that one's belief that p is not false, *both* when one's method of believing p *would be the same* if p were false *and* when it would not be the same, then the concern that appealing to alternative methods in counterfactual worlds (in order to rescue sensitivity) generates further heterogeneity does not arise.

[9]Roush (2006) and Topey (2022) both incorporate recursion within a sensitivity framework to maintain closure and resolve other issues with sensitivity. On their accounts (whose differences need not detain us here), an otherwise *insensitive* belief can still be knowledge so long as it is reliably inferred from a sensitive belief. My version of reasons-based sensitivity, while not Nozickean in its conception of *what* needs to be sensitive, does speak to Nozick's ambition to present sensitivity as a perfectly general requirement on knowledge.

[10]Incidentally, the account set out above also handles the original grandmother case that inspired Nozick's appeal to methods. Grandmother's reason for believing that the grandson is well is that he looks healthy. Assuming that this is an adequate basis for knowing (which is not in dispute here), then regardless of what she would believe or what method she would use if the grandson were not healthy, her reason would not be the case.

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

3.1. What is the core problem again?[11]

I take the core problem to be that sensitivity is logically incompatible with knowing that one's belief is not false, explained how Becker and Salerno do not address this problem specifically, and explained Melchior's criticism that Wallbridge's solution to the core problem leads to heterogeneity. I noted early in §I, however, that Vogel sees the core problem as (also?) involving a violation of closure, because p itself, which S knows, entails that S's belief that p is not false.

If one takes closure violation to be the central issue, the RDS approach may offer relief. S knows that Omar has new shoes because, if that were false, S's reason (Omar said so) would not be the case. S knows that their belief that Omar has new shoes is not false because, if it were false and S did have a false belief, S's known, higher-level reason (Omar has new shoes) would not be the case, even though in that counterfactual world S would believe Omar has new shoes, perhaps even because Omar said so!

Sensitivity theorists may be reluctant to accept the offering, however, because it may smack of Moorean insouciance (or worse still, begging the question). "How do you know that Omar has new shoes [/that you have hands]?" *Because he said so [Because look!]*. "But how do you know that your belief is not false [/that you're not being deceived by how things look]?" *Because what I believe is true!* Most sensitivity theorists notoriously reject closure, and I submit that this is because they think that the epistemic requirements on knowing q, even when one knows that it follows from p, are sometimes stronger than knowing that p. According to this line of thinking, knowing that one is not a BIV takes more than knowing that one has hands. Achieving higher-level knowledge sometimes takes more than first-level knowledge that entails it. In the RDS idiom, it takes further reasons. In the remainder of this section, I articulate an alternative RDS account that coheres with these commitments.[12]

3.2. RDS treatment (A): preliminaries

Preliminaries (1): S has first-order knowledge. Assume that S knows that q: Omar has new shoes. (Note that I am back to using "q" for the first-level proposition here, because the central question is how and whether S knows that p: that their belief that q is not false.) S's reason for believing that q, R_q, is that Omar said he has new shoes and the shoes Omar is wearing look new. If q were false, R_q would not be the case. This much is

[11]I am indebted to an anonymous referee for raising the issues taken up in this entire subsection and suggesting the solution set out in its second paragraph. I proceed from there to offer a more complex solution to the higher-level knowledge problem because I imagine many theorists who are attracted to sensitivity objecting that one cannot achieve such knowledge solely through inference from first-level knowledge, and more generally that one cannot always achieve knowledge that p solely through inference from sensitively-believed q. They may insist that higher-level knowledge requires a stronger epistemic position – reasons (broadly construed) that are independent of the reasons for one's first-level belief. ("You've reconciled sensitivity with higher-level knowledge. Fine, thanks, but now higher-level knowledge comes too easily!") Higher-level knowledge and closure come under threat where those independent reasons are lacking. The account here speaks to these concerns. It reconciles higher-level knowledge with the necessity of sensitivity and shows how, if one insists that higher-level knowledge demands more than first-level knowledge, the question of whether the more demanding reasons suffice for knowledge still turns on whether those reasons are sensitive, regardless of whether one would have those reasons or would maintain their belief if the higher-level proposition were false.

[12]It is worth noting that, while Wallbridge's account does not require anything more for higher-level knowledge than the method from which one's first-level knowledge derives, Melchior's heterogeneity objection reveals how that account still admits of closure violations because there are cases where, on Wallbridge's account, S can know that p but not be able to know that their belief that p is not false.

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

granted even by opponents of sensitivity. For S to know that q, it must also be true that Omar is not lying or mistaken, but S need not know that.

Preliminaries (2): stronger requirements on higher-level knowledge. Does S know that their belief that q is not false? Because it is a condition on knowing that q that Omar is not lying or mistaken, for S to achieve higher-level knowledge that their belief is not false, arguably, S must know that the condition is satisfied – or so I think many sensitivity theorists would agree, as explained in the previous subsection. The question – But how do you know your belief is not false? – can be understood as generating the implicature that maybe Omar is lying or mistaken, to which the reply "Because Omar has new shoes" will ring hollow. Extant responses to the higher-level knowledge do not try to account for this. Becker's and Salerno's reconciliations of sensitivity with higher-level knowledge do not depend on one's being in an epistemic position any stronger than what is required for first-order knowledge, for example that q: Omar has new shoes. Their view is that, if S's belief that they had a true belief that q were false, S would either not believe that q anyway, or q would be false, but in the latter case, because, *ex hypothesi*, S actually knows that q, S would not believe that q. Either way, S would not believe that q and so would not believe that they believe that q. Wallbridge's response to the core problem is, essentially, "Maybe S doesn't know that Omar is not lying or mistaken, for if Omar were lying or mistaken, S would still believe him, but that world is further away than one where S uses a different method." Sensitivity theorists who countenance closure failures in cases where requirements on knowing that q are stronger than those on knowing that p and that p entails q should welcome the incorporation of this thought in their response to the higher-level knowledge problems.

3.3. RDS treatment (B): application of RDS

What is S's (Dretskean) reason for believing p? Let's ask S: You believe Omar has new shoes because he said so, but how do you know your belief is not false? S gives their reason, R_p: Omar is not lying or mistaken. (1) If p were false, Omar would be lying or mistaken. R_p is sensitive to, or tracks, p. Crucially, I am not claiming (and RDS does not imply) that if p were false, S would not believe R_p or would have some other basis for belief. That would be just another version of Nozickean sensitivity. Applying RDS as developed above, I am only pointing out that R_p tracks p. Still, for R_p to be a Dretskean conclusive reason for p, and to deliver the advertised more compelling response to critics, S must know R_p.[13]

Does S know that R_p, that Omar is not lying or mistaken? Let us ask S: What is your reason for believing R_p (that Omar is not lying or mistaken)? S replies with R_{Rp}: Omar is trustworthy. (2) If R_p were false, R_{Rp} would not be the case. R_{Rp} tracks R_p. If Omar were lying or mistaken, it would be false that Omar is trustworthy. Here again, S may still believe that Omar is trustworthy if Omar were lying or mistaken, but S's actual reason R_{Rp} tracks R_p. The closest worlds[14] where Omar is lying are worlds where Omar is not trustworthy. Does this not contradict the previous paragraph? No. The worlds relevant to assessing the counterfactual "If S's belief that p were false," are not those relevant to assessing "If Omar were lying or mistaken." This is an important and potentially controversial point. I will come back to it after rounding off the RDS treatment.[15]

[13]Note accompanying (1): That (something like) *Omar is trustworthy* must surely be part of S's reasons, a point I am about to address, but citing Omar's trustworthiness in response to "How do you know your belief is not false?" does not answer the question. The questioner will respond, "Okay but his trustworthiness does not entail that your belief is not false."

[14]Or most normal? Skip ahead to note 18.

[15]Note accompanying (2): S's having accumulated experience with Omar's asserting (relevant) truths is not sufficient for knowing that Omar is not lying or mistaken. S must at least tacitly come to believe, at least

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

Which takes us, finally, to the question: Does S know R_{Rp}, that Omar is trustworthy? S believes Omar is trustworthy based on R_e, experience with Omar asserting truths. (3) If R_{Rp} were false, R_e would not be the case. R_e tracks R_{Rp}. ((3) is true *if* S knows R_{Rp}. It could otherwise be false. When (3) is false, it just shows that S does not know that Omar is trustworthy, and therefore does not know that Omar is not lying or mistaken, and therefore does not know that their belief is not false.)

Now for a summary to make the results explicit. S knows that p (they do not have a false belief that Omar has new shoes) on the basis of R_p (Omar is not lying or mistaken). That is, (1) if p were false, Omar would be lying or mistaken. R_p tracks p. S knows that R_p (Omar is not lying or mistaken) on the basis of R_{Rp} (Omar is trustworthy). That is, (2) if R_p were false, R_{Rp} would not be the case. R_{Rp} tracks R_p. S knows that R_{Rp} (Omar is trustworthy) on the basis of R_e (experience with Omar asserting truths). That is, (3) if R_{Rp} were false, R_e would not be the case. R_e tracks R_{Rp}. And yet, it could easily be the case that (4): if p were false, R_e would still be the case and that S would *believe* that it is not the case that p is false. Indeed, it could even be that if R_p were false – if Omar were lying or mistaken – R_e would still be the case. What matters is that each reason in the chain of reasons tracks that for which it is the immediate reason, not that one's actual experience tracks the truth.[16] It is not required that S would not have believed that p, or would have had a different method or different set of experiences, if p were false. The recursive structure of conclusive reasons and focus on the sensitivity of the reasons themselves allow the RDS account to succeed where belief sensitivity and RNS fail to reconcile sensitivity and higher-level knowledge.

3.4. RDS treatment (C): explication and clarification

The RDS response now complete, I turn to two concerns that spring to mind with the RDS treatment of higher-level knowledge. The most prominent is a question about the evaluation of the counterfactuals. The other concern is whether achieving higher-level knowledge requires such an elaborate structure of reasons. I'll take the second concern first.

One might wonder why there should be so many layers of reasons in a typical case of higher-level knowledge. The first thing to note is that higher-level knowledge is conceived here as a greater epistemic achievement than simply knowing that p. The second thing is that, while perhaps not all cases of higher-level knowledge require as many layers of reasons, in the case illustrated, a compelling response to both the core problem and heterogeneity

dispositionally, that Omar is trustworthy (or some other relevant generalization that justifies their believing that Omar is not lying or mistaken this time). There are two related points here. First, Humean skepticism about induction to the side, if one doesn't at least *believe* that the future will be like the past, one has no reason for inductive beliefs about specific instances. Second, the relevant generalization must be somewhat robust. It cannot just be something merely like *It is unlikely that Omar would lie* but must be something stronger, akin to *Omar does not lie*. Otherwise, S is in an epistemic position similar to one in which one believes one will not win the lottery. (See Becker (MS) on how this applies to Melchior's heterogeneity charge with regard to sensitivity and inductive knowledge.)

[16]Does that not generate another worry about heterogeneity? Suppose S's reason for believing that they don't falsely believe that Omar has new shoes is not that Omar is not lying or mistaken, but rather S simply infers it from Omar's saying that he has new shoes. In this case, if S's higher-level belief were false, S would still have that higher-level belief, Omar would say he has new shoes (he'd be lying or mistaken), and so S lacks higher-level knowledge. The results may appear to be heterogeneous – higher-level knowledge in some cases but not others – but the thrust of my account hinges on the insight that higher-level knowledge depends on Dretskean reasons. When they are of the right kind, one knows, when not, not. Since the implications of the account are principled, the heterogeneity (if that's what it is) is benign. Thanks once again to an anonymous referee for pressing this concern.

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

problem requires that S knows that Omar is not lying or mistaken (or can know it – keep in mind that sensitivity, even as construed as RDS, is a necessary but not sufficient condition). S must have some reason for believing this, and simply having had a set of past experiences, or observations, by itself does not constitute a reason. Typically, S must generalize in some way, for example, infer that Omar is trustworthy, and use it as a reason for believing that Omar is not lying or mistaken. The case under consideration has the structure of an inductive inference (conscious or otherwise) to a particular conclusion: observation→generalization→application to particular instance. In short, that there should be a somewhat sophisticated structure of reasons for what is a fairly sophisticated epistemic achievement should not, on final consideration, be surprising.

I imagine two separate but perhaps related questions concerning the evaluation of the counterfactuals. The first is this: How can R_p track p, R_{Rp} track R_p, and R_e track R_{Rp}, while R_e does not track p? (A similar question can be raised with regard to the potentially closure-saving strategy briefly articulated near the beginning of this section.) The basis of the answer is that counterfactuals are widely regarded as not transitive. (See Haze 2023 and Brogard and Salerno 2008 for possible explanations.) A simple example: (a) If I were a fireman, I would wear red suspenders. (b) If I were to wear red suspenders, my outfit would look funny. (c) If I were a fireman, my outfit would look funny. (a) and (b) are (or can be) true compatibly with (c) being false.[17] When the set of worlds relevant to assessing the conclusion in a hypothetical syllogism (e.g. worlds where I am a fireman) is not a subset of the set of worlds relevant to assessing one or more premises (e.g. worlds where I am not a fireman, but I wear my typical clothes, as in (b)), it can turn out that the syllogism is not valid (Haze 2023, 241).

As applied to the Omar case, worlds relevant to (4), where p is false (S does believe falsely that q), may not be among the set of worlds relevant to (3), where Omar is untrustworthy. In the closest worlds where p is false, somehow, what Omar said was incorrect. Period. (Omar is nevertheless trustworthy and S has their actual experience with Omar's truthfulness, which explains why S comes to believe what Omar says, albeit falsely.) But in worlds where Omar is untrustworthy, S doesn't have their experience with Omar's truthfulness.

The explanation of lack of transitivity just adumbrated is somewhat brief, perhaps too abstract, and arguably also ducks the second, more pressing question, which has more to do with (2) than (3). That is, the explanation implies (and relies upon) the claims that (1) if S had a false belief that q, Omar would have lied or been mistaken but *would still have been trustworthy* (and S would have had their experiences with Omar's truthfulness), whereas (2) if Omar would have lied or been mistaken, he *would not have been trustworthy* (and S would *not* have had their experiences with Omar's truthfulness). Some will find this contentious. Because the fireman example is less contentious, perhaps we can gain insight from it.

In the not-so-distant worlds where I wear red suspenders, my outfit looks funny. The explanation is that my actual clothes don't go with red suspenders. But in the more distant worlds where I am a fireman, I wear red suspenders *and my outfit would not look funny*. The explanation is that in such faraway worlds I also wear fireman gear. In the not-so-distant worlds wear Omar speaks falsely (about plain facts in relation to which he is in excellent epistemic position), he is untrustworthy. The explanation is that trustworthy people do not utter falsities (about plain facts in relation to which they are in excellent epistemic position). But worlds where S believes falsely that Omar has new shoes are more distant, and in them Omar lies or is mistaken *and is trustworthy*. S believes falsely *because* Omar is trustworthy.

[17]Another well-known example (from Stalnaker, cited in Starr, 2019). If J. Edgar Hoover were a communist, he would be a traitor. If he were Russian, he would be a communist. It's false, though, that if he were Russian, he would be a traitor. His being Russian would explain his being a communist, which implies he would not be a traitor.

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

Consider the illustration below. In the actual world: my outfit does not look funny (a matter of controversy I suppose, but let it stand) and I don't wear red suspenders; Omar tells the truth and is trustworthy. As one moves out from the actual world, my outfit's looking funny naturally "comes with" the red suspenders; Omar's untrustworthiness naturally "comes with" his lying or being mistaken. But as we move further out still, something else is different: I am a fireman, which naturally "comes with" red suspenders so my outfit no longer looks funny; S has a false belief, where S's believing Omar's false utterance naturally "comes with" or is explained by Omar's being trustworthy (at least not being untrustworthy).[18]

Illustration

I am a Fireman/S has a False Belief that their belief is not false, F/FB
I wear Red Suspenders/Omar Lies or is Mistaken, R/OLM
My outfit Looks Funny/Omar is Untrustworthy, LF/OU

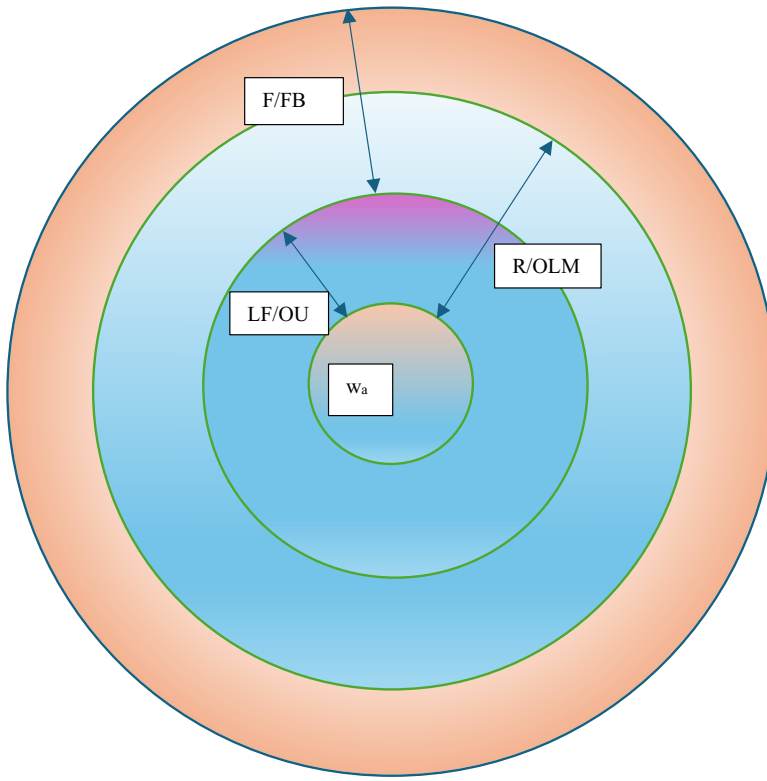

[18]This manner of *explaining* my readings of the counterfactuals may seem to fit more naturally with a "normal worlds" framework of possible worlds (as in, e.g. Smith 2016) rather than the more standard "distance" framework. In normal worlds where Omar is mistaken, Omar is untrustworthy. In normal worlds where S believes what Omar says though it is false, Omar is trustworthy (which is why S believes him). I adhere to the more standard framework in the *explanation* because I want to exploit the analogy with uncontroversial cases like the fireman and Hoover (previous note). But at the end of the day, what matters is not necessarily near or far, normal or abnormal, but the Haze point cited, according to which the fundamental explanation of intransitivity is that worlds relevant to assessing one counterfactual are not relevant to assessing another.

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

I do not claim that my reading of the counterfactuals is unquestionably true, but I do claim that it is defensible, that other cases of intransitivity of counterfactual conditionals support it, and that the explanations of intransitivity are analogous.

4. Concluding remarks

Sifting through all the details, the main objective of this paper is to provide a framework for answering the various criticisms that sensitivity makes a mess of higher-level knowledge. Taking the bull by the horns, I grant, in accord with Melchior and in deference to the spirit of the core objection, that whether or not one achieves higher-level knowledge should not depend on whether they would have used a different method in coming to a first-order belief if their higher-level belief were false. The more specific goal then is to show how sensitivity is compatible with higher-level knowledge even though the closest worlds where one falsely believes that their belief is not false are worlds where their first-order belief has the same basis as in the actual world.

The goal is achieved through two Dretske-inspired moves. The first is the shift in focus from the sensitivity of either belief or method-relative belief to the sensitivity of reasons. The second involves unpacking the recursive structure of reasons. Putting these ideas together, I have shown how S can know p (that their belief is not false), when they have a reason R_p that tracks p, can know R_p when they have experiences that track R_p, even when their experiences (and methods and beliefs) do not track p. The very possibility derives in general from the intransitivity of counterfactuals. At the end of the previous section, I offered an explanation of how intransitivity applies to cases of higher-level knowledge.[19]

References

Adams F. and Clarke M. (2005). 'Resurrecting the Tracking Theories.' *Australasian Journal of Philosophy* **83**(2), 207–21.

Alfano M. (2009). 'Sensitivity Theory and the Individuation of Belief-Formation Methods.' *Erkenntnis* **70**(2), 271–81.

Becker K. (2006). 'Is Counterfactual Reliabilism Compatible with Higher-Level Knowledge?' *Dialectica* **60**(1), 79–84.

Becker K. (2023). 'Sensitivity, Discrimination, and Perspective.' In L.R.G. Oliveira, (ed), *Externalism about Knowledge*. pp. 75–98. Oxford University Press.

Becker K. (MS). 'Dretskean Sensitivity and Induction'.

Bjerring J.C. and Gundersen L.B.. (2020). 'Higher-Order Knowledge and Sensitivity.' *Canadian Journal of Philosophy* **50**(3), 339–49.

Brogard B. and Salerno J. (2008). 'Counterfactuals and Content.' *Analysis* **68**(297), 39–46.

Dretske F. (1971). 'Conclusive Reasons.' *Australasian Journal of Philosophy* **49**(1), 1–22.

Haze T.G. (2023). 'Sufficient Conditions for Counterfactual Transitivity and Antecedent Strengthening.' *American Philosophical Quarterly*, **60**(3), 237–47.

Kripke S. (2011). 'Nozick on Knowledge.' In *Philosophical Troubles: Collected Papers*, *vol.* **1**. pp. 161–224. Oxford University Press.

Melchior G. (2014). 'A Generality Problem for Bootstrapping and Sensitivity.' *Croatian Journal of Philosophy* **14**(1), 31–47.

Melchior G. (2015). 'The Heterogeneity Problem for Sensitivity Accounts.' *Episteme* 12(4), 479–96.

Melchior G. (2017). 'Sensitivity Has Multiple Heterogeneity Problems: A Reply to Wallbridge.' *Philosophia* **45**(4), 1741–47.

Melchior G. (2021). 'Sensitivity and Inductive Knowledge Revisited.' *Dialectica* **75**(3), 441–62.

[19]Many thanks to participants at the Bled Epistemology Conference 2023 and to anonymous referees for helping to make this paper *much* better than it otherwise would have been.

https://doi.org/10.1017/epi.2025.10087 Published online by Cambridge University Press

Nozick R. (1981). *Philosophical Explanations*. Harvard University Press.
Roush S. (2006). *Tracking Truth: Knowledge, Evidence and Science*. Oxford University Press.
Salerno J. (2010). 'Truth Tracking and the Problem of Reflective Knowledge.' In J. K. Campbell, M. O'Rourke, and H. S. Silverstein, (eds), *Knowledge and Skepticism*. pp. 72–81 MIT Press.
Smith M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford University Press.
Sosa E. (1999). 'How to Defeat Opposition to Moore.' *Philosophical Perspectives* **13**, 141–53.
Starr W. (2019). 'Counterfactuals.' *The Stanford Encyclopedia of Philosophy* (Winter 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.), <https://plato.stanford.edu/archives/win2022/entries/counterfactuals/>.
Topey B. (2022). 'Saving Sensitivity.' *Philosophical Quarterly* **72**(1), 177–96.
Vogel J. (1987). 'Tracking, Closure, and Inductive Knowledge.' In *The Possibility of Knowledge: Nozick and His Critics*. pp. 197–215. Rowman & Littlefield.
Vogel J. (2000). 'Reliabilism Leveled.' *Journal of Philosophy* **97**(11), 602–23.
Wallbridge K. (2016). 'Sensitivity and Higher-Order Knowledge.' *Pacific Philosophical Quarterly* **99**, 164–73.
Wallbridge K. (2017). 'Sensitivity Hasn't Got a Heterogeneity Problem: A Reply to Melchior.' *Philosophia* **45**(2), 835–41.
Wallbridge K. (2018). 'Sensitivity, Induction, and Miracles.' *Australasian Journal of Philosophy* **96**(1), 118–26.
Zhao B. (2024). 'On Relativizing the Sensitivity Condition to Belief-Formation Methods.' *American Philosophical Quarterly* **61**(2), 165–75.
Zhao B. (2025). 'On Mentioning Belief-Formation Methods in the Sensitivity Subjunctives.' *Ergo* **12**, 9, doi: https://doi.org/10.3998/ergo.4634.
Zhao H. (2022). 'Sensitivity Unmotivated.' *Acta Analytica* **37**(4), 507–17.
Zhao H. (2023). 'Can Sensitivity Preserve Inductive Knowledge?" *Philosophia* **51**, 1865–82.

Kelly Becker has taught at the University of New Mexico since 2002. Before that, he was Assistant Professor of Philosophy at Centre College from 1999 to 2002, after receiving his PhD from UC San Diego in 1999. He works primarily in the Theory of Knowledge but maintains long-standing interests in the Philosophies of Language and Mind.

Cite this article: Becker K. (2025). "Dretskean Sensitivity and Higher-Level Knowledge." *Episteme* 1–14. https://doi.org/10.1017/epi.2025.10087