Saving Sensitivity[†]

Brett Topey
✉ brett.topey@plus.ac.at

Abstract

Sensitivity has sometimes been thought to be a highly epistemologically significant property, serving as a proxy for a kind of responsiveness to the facts that ensures that the truth of our beliefs isn't just a lucky coincidence. But it's an imperfect proxy: there are various well-known cases in which sensitivity-based anti-luck conditions return the wrong verdicts. And as a result of these failures, contemporary theorists often dismiss such conditions out of hand. I show here, though, that a sensitivity-based understanding of epistemic luck can be developed that respects what was attractive about sensitivity-based approaches in the first place but that's immune to these failures.

1 Introduction

It's natural to suppose that *sensitivity*—where *S*'s belief that *p* counts as sensitive just in case *S* wouldn't still believe *p* were it not the case that *p*—is a highly epistemologically significant property. In particular, the sensitivity of a thinker's belief appears to serve as a proxy for a sort of responsiveness to the features of the part of the world the belief is about, a connectedness between that part of the world being the way it is and the belief being what it is that ensures that the truth of the belief isn't accidental, isn't just a lucky coincidence. And it's a useful proxy: while it's difficult to say directly just what sort of responsiveness is relevant to whether a belief is merely accidentally true, the counterfactual semantics required to analyze sensitivity conditionals is, by comparison, relatively well understood.

[†]This is a pre-copyedited, author-produced version of an article accepted for publication in *Philosophical Quarterly* following peer review. The version of record is available online at doi.org/10.1093/pq/pqab015.

This, I take it, goes some way toward explaining why sensitivity conditionals have so often played a starring role in post-Gettier attempts to offer an analysis of knowledge in the externalist tradition. Very briefly: It's plausible that what it takes for a belief to count as knowledge is for it to be both justified and nonaccidentally true (or, equivalently, for it to be both justified and not veritically lucky),[1] and it's also plausible that what it takes for the truth of a belief to be nonaccidental is for that belief to be responsive in the right way to the features of the world that are relevant to its truth.[2] So, insofar as whether a belief is sensitive tells us whether it exhibits the relevant sort of responsiveness, it turns out that what it takes for a belief to count as knowledge is for it to be justified and sensitive. (The thesis that sensitivity is necessary for knowledge was introduced by Nozick (1981: chap. 3), as a part of his tracking theory of knowledge.[3]) Call this the *simple account* of the epistemological significance of sensitivity.

Notoriously, sensitivity conditionals turn out not to be perfectly suited for the job the simple account requires them to do: there are well-known failures in both directions, insensitive beliefs that seem nevertheless to be knowledge as well as beliefs that clearly fail to be knowledge despite being both justified and sensitive. My belief that I don't falsely believe I have hands (see, e.g., DeRose 1995, Kripke

[1] A belief is *veritically lucky* if it merely happens to be true, if its truth is a lucky coincidence—this terminology is due to Engel (1992). That knowledge excludes veritic luck is arguably the central lesson of Gettier 1963. For further discussion, see, e.g., Unger 1967 and Pritchard 2015.

[2] For explicit discussion of the intuitive relationships between sensitivity, responsiveness, and accidental or lucky truth, see, e.g., Roush 2005: chap. 1, Becker 2012, and Murphy and Black 2012.

[3] Note that Nozick's tracking theory supplements the sensitivity condition with an *adherence* condition, where *S*'s belief that *p* counts as adherent roughly when there aren't nearby worlds where *S* fails to believe that *p* despite its being the case that *p*. But there's compelling reason to think adherence can't be an epistemologically significant property—see, e.g., Sosa 2002, Kripke 2011, Luper 2012, and Topey 2020. Note also that Nozick doesn't include justification as a separate condition on knowledge. Instead, he adopts a radically externalist account of justification and suggests that any belief that meets the sensitivity and adherence conditions will thereby be justified—see his 1981: 267. But the thesis that sensitivity and adherence suffice for justification entails that agents' beliefs are justified in a wide variety of situations in which it's intuitively obvious that they're not—see Kripke 2011: 170ff. Some later tracking theorists—see, e.g., Adams 1986 and Roush 2005: chap. 1 and app. 1.1—explicitly reject the claim that knowledge requires justification. See §3 for further discussion of this aspect of Roush's view.

2011, Vogel 1987),[4] for instance, is a belief of the former sort,[5] and cases with the structure of Kripke's red barn case (again, see his 2011) provide examples of beliefs of the latter sort.[6] We'll discuss both sorts of failure in some detail below. What's important for now is just that, given these failures, we can be sure that the simple account isn't correct—if sensitivity really is epistemologically significant, some more complicated story of that significance is going to be needed.

This paper explores the prospects for providing such a story. What we'll see is that sensitivity is indeed epistemologically significant, and we can indeed describe that significance in a way that allows us to avoid the sorts of failures mentioned above. But gaining a clear picture of what that significance is will require abandoning the idea, common among sensitivity theorists, that the primary epistemological

[4]It's sometimes suggested—see, e.g., Salerno 2010 and Bjerring and Gundersen 2020—that, in fact, the simple account gets cases like this one right: my belief that I don't falsely believe I have hands *doesn't* count as knowledge, but a closely related belief that's easily confused with this one—i.e., my belief that I *do truly* believe I have hands—*is* sensitive and so *does* count as knowledge. (Note: though Becker (2006) anticipates this response by pointing out that beliefs like the latter one are sensitive, he doesn't go so far as to suggest that the simple account is correct in delivering the verdict that beliefs like the former one don't count as knowledge.) But I take it that an account on which one of these beliefs counts as knowledge while the other doesn't is, as Melchior puts it, "far too heterogeneous to provide a plausible picture of higher-level knowledge in terms of sensitivity" (2015: 480).

[5]It's sometimes suggested that true beliefs gained by induction are further examples of beliefs of the former sort—see, e.g., Vogel 1987, 2007 and Sosa 1999b. But this isn't obviously correct, as has been pointed out by, e.g., Cross (2010).

[6]Metaphysically necessary truths are sometimes taken to be a further source of failures of the latter sort: counterfactuals with necessarily false antecedents are vacuously true, and so any case in which *S* believes that *p* and it's necessarily true that *p* is trivially a case in which *S*'s belief is sensitive, even if it's intuitively quite clear that the belief fails to be knowledge (see, e.g., Blome-Tillmann 2017). But there's a sense in which this belief is *not* trivially sensitive. As I suggest elsewhere (see my 2020), if *S* is considering, from a first-personal perspective, whether the truth of her own belief is a matter of luck, the worlds she should take to be relevant aren't the metaphysically possible worlds but the *epistemically* possible worlds, where a world counts as epistemically possible for *S* just in case she isn't able to definitively rule out a priori that it will turn out to be the actual world—what she's considering, after all, is whether it's just a matter of luck that, from all the candidates for being her actual situation, she's successfully picked out the situation she's really in. When what's in question is whether a belief is sensitive, where the worlds relevant to sensitivity are the worlds that are epistemically possible in this sense, let us say that what's in question is whether the belief is *first-personally sensitive*. Then even if it's metaphysically necessary that *p*, it's not trivial that *S*'s belief that *p* is first-personally sensitive: a metaphysically impossible world may be epistemically possible in the relevant sense. And at least one theorist—i.e., Collin (2018)—suggests that the sensitivity conditionals that appear in analyses of knowledge should be understood first-personally. But even if this isn't correct, there's a case to be made that first-personal sensitivity is relevant to whether our beliefs are *justified*. (See the discussion of the Benacerraf–Field challenge in §2 below.) So, insofar as knowledge requires justification, a belief with necessarily true content can fail to be knowledge by failing to be first-personally sensitive.

role for sensitivity to play is as a constituent notion in a Gettier-proof analysis of knowledge. In the next section I explain why that is and motivate a broader conception of the role of sensitivity, and then in the rest of the paper I show that, given this broader conception, we can motivate a new sensitivity-based condition that's counterexample-free.

2 How to conceive of the role of sensitivity

Attempts by sensitivity theorists to avoid the sorts of failures mentioned above aren't hard to come by, and they all tend to proceed in roughly the same way: after the existence of a particular species of counterexample is noted, some patch is offered, some complication that has been designed explicitly to allow a sensitivity-based condition on knowledge to avoid those counterexamples.[7] This way of proceeding is summed up neatly by Cross (2010: 40), in his description of his own project:

> I aim to defend [sensitivity] against the onslaught of counterexamples in the usual manner of analytic epistemologists: namely, by modifying the condition stepwise, ever so slightly, so as to accommodate the cases while leaving some non-baroque remainder that retains at least as much intuitive appeal as sensitivity itself.

The idea, then, is to design a sensitivity-based condition on knowledge that avoids counterexamples without being so convoluted that it seems obviously gerrymandered.

I suggest, though, that this way of proceeding isn't fit for purpose. Consider again the motivation I offered above for a sensitivity condition on knowledge: the thought was that whether a belief is sensitive can tell us whether it's responsive in a particular way to the features of the world relevant to its truth, where this sort of

[7] Examples of theorists who present patched conditions in roughly this way include DeRose (1995; DeRose's condition is also endorsed in Murphy and Black 2012), Roush (2005), and Cross (2010; Cross's condition is also endorsed in DeRose 2010), to name a few—see below for further discussion of all of these patched conditions. Even Nozick's own discussion has this form: after introducing the idea of sensitivity, he notes the existence of certain counterexamples and then, in order to avoid them, relativizes his sensitivity condition to methods of belief formation, so that *S*'s belief that *p*, arrived at via method *M*, counts as sensitive just in case, if it weren't the case that *p* and if *S* used *M* to arrive at a belief about whether *p*, *S* wouldn't believe that *p*. (See his 1981: 179ff. Incidentally, it's not clear that relativization to methods really helps Nozick handle problematic cases of the relevant sorts (see Goldman 1983), and, at any rate, complicating the condition in this way turns out to be unnecessary: an unrelativized sensitivity condition, appropriately revised so as to handle counterexamples that uncontroversially are *not* handled by relativization to methods, turns out to have no problem handling the problematic cases that worried Nozick in the first place. For further discussion, see, e.g., Roush 2005: chap. 2 and Cross 2010.)

responsiveness is what's relevant to whether the truth of the belief is nonaccidental. What the sorts of cases mentioned above show, it seems to me, is that there's something flawed about this motivation.

Consider, for instance, the case of my belief in the following proposition:

(1) I don't falsely believe I have hands.

The reason this belief counts as knowledge, I take it, is that, despite its insensitivity, it *is* responsive to the features of the world relevant to its truth and so *is* nonaccidentally true. To see why, note first that I have a belief closely related to this one that *is* sensitive: my belief in

(2) I have hands.

And this belief certainly is responsive to the features of the world relevant to its truth: my having hands—i.e., the truth of (2)—is what causes me to have this belief. But my having hands is, of course, *also* sufficient for the truth of (1). Furthermore, I *believe* that my having hands is sufficient for the truth of (1), and so, for any grounds I take to suffice for believing (2), I take those grounds to suffice also for believing (1). The (subjective) grounds on which I in fact believe (2), then, are, by my own lights, sufficient to ground a belief in (1) as well. And this fact plausibly plays a role in explaining why I maintain a belief in (1).[8] But if that's right, I do indeed believe (1) on the basis of responding to a feature of the world—my having hands—that's sufficient for the truth of (1).[9]

If this diagnosis is correct, what cases of this sort tells us is that there are beliefs that, despite being insensitive, turn out, in virtue of their connections to beliefs that *are* sensitive, to be responsive (albeit in an inherited way) to the features of the world relevant to their truth, in which case sensitivity's relationship to the sort of responsiveness relevant to whether a belief is nonaccidentally true isn't as straightforward as we might have thought. (And cases with the structure of Kripke's red barn case, I take it, tell us that this relationship is more complicated than we might have thought in yet a different way—in §4 I diagnose these as cases in which a belief happens to be sensitive despite failing to be responsive in the relevant way.) The way to proceed, then, surely isn't to try out ad hoc modifications of the sensitivity condition in the hope of avoiding counterexamples—the theory that results from

[8] This isn't to say that I've formed my belief in (1) by inferring it from (2)—indeed, it isn't to say anything at all about the precise mechanisms by which I formed my beliefs. It's to say only that, whatever the (subjective) grounds are on which I believe (2), I believe (1) at least partly on those same grounds.

[9] Note that the claims I'm making here are descriptive, not normative: what's relevant to whether my belief in (1) is responsive to my having hands is not what justifies my beliefs in any objective sense but what the connections are between my own subjective reasons for believing (1) and my own subjective reasons for believing (2).

such a procedure is bound to be unprincipled even if not obviously gerrymandered. The way to proceed is instead to reexamine the relationship between sensitivity and the relevant sort of responsiveness.[10]

I take it that the lesson here, more broadly, is this: if we want to understand the epistemological significance of sensitivity, we must keep in mind that its primary epistemological role is as a constituent notion, not in the analysis of knowledge, but in an account of nonaccidental truth; it may play a role in the analysis of knowledge, but if so, it does so only derivatively, via its role in some condition ruling out veritic luck.[11] My proposal, then, is that we set the analysis of knowledge aside and instead try directly, by thinking about the relevant sort of responsiveness, to give a sensitivity-based account of what it takes for a belief to be nonaccidentally true.

Approaching the question of the epistemological significance of sensitivity in this direct way is beneficial for at least two reasons. First, an account of sensitivity's relationship to nonaccidental truth in general has the potential to be more theoretically fruitful than an account of its role in the analysis of knowledge: if we can give a sensitivity-based account of nonaccidental truth, we'll thereby have given an account of sensitivity's role in the analysis of knowledge (since we'll have shown what role sensitivity must play in a condition ruling out veritic luck), but we'll also have left open the possibility of using our account to answer epistemological questions that have nothing to do with the analysis of knowledge.

Consider, for instance, the Benacerraf–Field challenge for mathematical platonism (and relevantly similar realist theories of other domains, such as logic and morality).[12] The epistemological principle underlying that challenge can be stated as follows:

FIELD'S THESIS. Insofar as our own theory of the nature of (say) the mathematical facts seems to make it impossible in principle to explain the coincidence between

[10] Note that abandoning sensitivity in favor of a *safety* condition of the sort developed by, e.g., Sosa (1999a, 1999b), Williamson (2000), and Pritchard (2005, 2009, 2015) amounts to abandoning altogether a responsiveness-based understanding of what it takes for a belief to be nonaccidentally true. The problem is that safety, unlike sensitivity, is defined just in terms of what is the case in *nearby* worlds, and so, if all nearby worlds are worlds where p and also are worlds where S believes that p, then, trivially, S's belief that p is safe regardless whether that belief is in any way responsive to its being that case that p. So, insofar as we're interested in a responsiveness-based understanding of nonaccidental truth, there's reason not to abandon sensitivity in favor of safety. This paper proceeds under the assumption that nonaccidental truth *is* to be understood in terms of responsiveness.

[11] Though sensitivity theorists do sometimes acknowledge this fact—see, e.g., Roush's description of her view as "an anti-luck epistemology" (2005: 5)—they tend to ignore it when the time comes to make adjustments in response to counterexamples, opting instead for the methodology described above by Cross.

[12] Benacerraf's version of the challenge appears in his 1973, but what I'm describing here is Field's improved version—see, e.g., his 1989 and his 2005.

our mathematical beliefs and those facts, we can't with justification hold on both to that theory and to our mathematical beliefs.

And the reason this principle is plausible is just that, if it really is impossible in principle to explain this coincidence, then it's just a matter of luck that the coincidence obtains. To accept a theory on which it's impossible in principle to explain this coincidence, then, is thereby to accept a theory according to which our own beliefs are accidentally true if true at all. So it would be unreasonable to accept such a theory while also insisting on the truth of those beliefs. Hence Field's suggestion that sensitivity has an important role to play here, that the challenge "seems to arise from the thought that we would have had exactly the same mathematical or logical beliefs, even if the mathematical or logical facts were different"; given this insensitivity, he says, "it can only be a coincidence if our mathematical or logical beliefs are right, and this undermines those beliefs" (2005: 81).

(It's sometimes suggested (see, e.g., Clarke-Doane 2016) that the challenge, construed in this way, is misguided: given the necessity of the truths of (say) mathematics, our mathematical beliefs are trivially sensitive if true at all. But note: the context here is a context in which we're considering, from a first-personal perspective, whether our own beliefs are nonaccidentally true, and as I explain elsewhere, such contexts are exactly the contexts in which it's plausible that certain metaphysically impossible worlds—i.e., those that are epistemically possible, in the sense that I can't definitively rule out their turning out to be actual—are relevant to sensitivity.[13] And if that's right, the metaphysical necessity of the truths of mathematics doesn't trivialize the question of whether our mathematical beliefs are sensitive in the sense relevant to the Benacerraf–Field challenge.)

The point is that the notion of nonaccidental truth has a crucial role to play here, a role that isn't directly relevant to the analysis of knowledge, and so the Benacerraf–Field challenge is at least one place where an account of the relationship between sensitivity and nonaccidental truth in general may pay epistemological dividends unrelated to the analysis of knowledge.[14]

The second benefit is that, by focusing directly on nonaccidental truth, we can avoid being led astray by the perceived need to give an account that respects every intuition, no matter how equivocal or idiosyncratic, about what cases of belief are correctly described as cases of knowledge.

[13] See my 2020 along with fn. 6 above.

[14] Incidentally, the connection of nonaccidental truth to the Benacerraf–Field challenge is the source of my own interest in sensitivity—my hope is to provide a sensitivity-based account of the sorts of first-personal judgments about our epistemic luckiness that are relevant to the challenge. The analysis of knowledge, on the other hand, is to my mind not particularly epistemologically interesting, for reasons discussed by, e.g., Kaplan (1985) and Schechter (2017).

For example, it's purportedly intuitive that we aren't correctly described as having knowledge of the negations of traditional skeptical hypotheses such as the brain-in-a-vat hypothesis, and many sensitivity theorists have been motivated, at least in part, by a desire to respect this intuition. It's by design, for instance, that, according to Nozick's original analysis, knowledge isn't closed under known implication; this feature allows the analysis to return the verdict that I don't count as having knowledge of

(3) I'm not a handless brain in a vat who's been electrochemically stimulated to have exactly the experiences I've in fact had.

despite the fact that I do have knowledge of (2)—i.e., the proposition that I have hands.[15] And some later patched conditions designed to be more lenient than Nozick's own are nevertheless *also* designed to respect the intuitiveness of this verdict: DeRose (1995), for instance, suggests a condition that, unlike Nozick's, can return the verdict that my knowledge of (2) gives me knowledge of (1), but this condition is also designed, like Nozick's, to return the verdict that I *don't* have knowledge of (3).[16] I'm not convinced that the intuition here is as robust (or as widespread) as these theorists take it to be. But whatever our intuitions are, the fact remains, as DeRose himself points out (see his 1995: sec. 10), that I'm in just as strong a position, in all epistemologically important respects, with respect to my belief in (3) as I am with respect to my belief in (2), for the simple reason that I myself recognize that any world where the former belief is false is thereby a world where the latter one is false as well. And focusing on nonaccidental truth rather than knowledge allows us to capture this fact: regardless of whether I'm correctly described as knowing (3), it's clear enough that my belief in (3) is indeed responsive to the features of the world relevant to its truth and so is nonaccidentally true, for the same reason my belief in (1) is: I believe (3) on the very same (subjective) grounds on which I (sensitively) believe (2), the truth of which is sufficient for the truth of (3).

This is all to say that, even independently of the fact that the failures mentioned in §1 are best seen, fundamentally, as counterexamples to the thesis that a belief is

[15] See also, e.g., Becker 2012: secs. 1–2.

[16] The idea is that the hypothesis that I falsely believe I have hands "doesn't explain how I went wrong with respect to my having hands" (1995: 23), while the brain-in-a-vat hypothesis does. (But see Black 2002 and Murphy and Black 2012 for an argument that DeRose's condition doesn't necessarily return the verdict about (3) that DeRose takes it to.) Note: DeRose doesn't take himself to be offering a genuine necessary condition on knowledge; he's instead concerned with the possibility of using sensitivity to provide an explanation of why, in particular cases, we judge subjects to know or not to know. This distinction doesn't matter for our purposes here. Note also that DeRose acknowledges in his 2010 that there are problems with the condition suggested in his 1995 and endorses instead a slightly different condition: the relative sensitivity condition introduced by Cross (2010). But this condition, too, is designed to return the verdict that I don't have knowledge of (3).

nonaccidentally true just in case it's sensitive, there are good reasons to conceive of our task here, not as the task of providing a sensitivity-based analysis of knowledge, but as the task of explaining what role sensitivity plays in an account of what it takes for the truth of a belief to be nonaccidental. In the rest of this paper, then, I develop a sensitivity-based necessary and sufficient condition for nonaccidental truth and demonstrate its immunity to counterexample, showing in §3 that *all* nonaccidentally true beliefs satisfy the condition and in §4 that *only* nonaccidentally true beliefs satisfy it.

3 Insensitive but nonaccidentally true beliefs

As suggested above, the following simple thesis:

Sensitivity. A belief is nonaccidentally true just in case it's sensitive.

has counterexamples in both directions—there are beliefs that are insensitive but nonaccidentally true as well as beliefs that are sensitive but merely accidentally true. In order to develop a counterexample-free sensitivity-based account of what it takes for a belief to be nonaccidentally true, it will be useful to have a diagnosis of both sorts of counterexample. We begin with the former.

Again, both my belief in (1) and my belief in (3) are nonaccidentally true despite being insensitive, since they're responsive to the features of the world relevant to their truth: what explains (at least in part) why I maintain these beliefs is that I sensitively believe (2) on (subjective) grounds that I take to suffice also to ground my beliefs in (1) and (3), since I take the truth of (2) to be sufficient for the truth of those two propositions. And something analogous seems to be true of other insensitive but nonaccidentally true beliefs. For example, in White's case in which jury members believe, in a case where compelling evidence against the defendant is abundant, that there wasn't an elaborate conspiracy in which "the defendant was framed and all the evidence planted" (2010: 581), this belief is insensitive, since, were there such a conspiracy, the jury members would have exactly the evidence they in fact have. But the belief is nevertheless nonaccidentally true: what explains why the jury members maintain this belief is that they sensitively believe the defendant is guilty and also take the truth of this latter belief to be sufficient for the truth of the former one.

These examples of insensitive but nonaccidentally true belief share the same general structure, which we can describe as follows: For some proposition q, S sensitively believes that q—i.e., in the nearest worlds where q isn't true, S refrains from believing that q—but there are nevertheless some other worlds, worlds farther from the actual world, where q isn't true but where S has the same (subjective) grounds for believing that q that she has in the actual world and so incorrectly believes that

q. Call such a world a *q-skeptical scenario*. Insofar as there are *q*-skeptical scenarios, there are going to be *p*s such that

(i) the truth of *q* is sufficient for the truth of *p* and
(ii) the nearest worlds where *p* isn't true are *q*-skeptical scenarios.[17]

And for any such *p*, if *S* takes the grounds on which she believes that *q* to suffice also to ground a belief that *p*, so that she believes that *p* and is disposed to do so in any case in which she has those grounds, then her belief that *p* is going to be an example of an insensitive but nonaccidentally true belief. Why nonaccidentally true? Because it inherits its responsiveness from *S*'s belief that *q*: by (i) along with the sensitivity of *S*'s belief that *q*, *S* does indeed believe that *p* for reasons that are responsive to the features of the world relevant to whether *p* is true. Why insensitive? Because, by (ii), the nearest worlds where *p* isn't true are *q*-skeptical scenarios and so are worlds where she has the same grounds for believing *p* that she has in the actual world.

Now for the surprising part: as it turns out, we can show that something like this structure is shared by *every* case of insensitive but nonaccidentally true belief—in particular, that, if *S* has an insensitive but nonaccidentally true belief that *p*, then *S* will believe some proposition *q* such that

(a) the truth of *q* is sufficient for the truth of *p*,
(b) *S*'s belief that *q* is nonaccidentally true, and
(c) *S* takes the grounds on which she believes that *q* to suffice also to ground her belief that *p*.[18]

Let *p* be a proposition such that *S* has an insensitive but nonaccidentally true belief that *p*. Then this belief, since it's nonaccidentally true, is responsive to the features of the world relevant to its truth: there are some features of the world such that, first, the world having those features is sufficient for the truth of *p*, and second, *S* believes that *p* on the basis of *responding* to those features, where responding involves discriminating between cases in which the world has those features and cases in which it doesn't. But insofar as *S* believes that *p* because she's discriminated in this way, she'll believe that *q* because she's discriminated in this way, where *q* is just the proposition that the world does indeed have the features in question. And this *q* satisfies all three of the above conditions: (a) because the world having the features in question is sufficient for the truth of *p*, (b) because *S*'s belief that *q* is sensitive, since *S* can discriminate between cases in which the world has the features

[17] The simplest *p*s with these properties are just propositions like (1)—i.e., propositions to the effect that *S* isn't in a *q*-skeptical scenario.

[18] Note that this isn't a normative condition but a descriptive one—again, the grounds on which she has a given belief are to be understood as her *subjective* grounds.

in question and cases in which it doesn't, and (c) because *S*'s basis for believing that *q* is such that she believes that *p* on the very same basis.

What this tells us is that we can narrow our focus significantly: in order to devise a sensitivity-based condition for nonaccidental truth that isn't overly strict—i.e., that *never* incorrectly classifies a belief as merely accidentally true—we need only ensure that our condition, in addition to classifying sensitive beliefs as nonaccidentally true, also classifies as nonaccidentally true insensitive beliefs in cases that have the structure just outlined. Our question, then, is how we might ensure that our condition does indeed return the right verdicts in cases with this structure.

The most straightforward way to do so is simply to let our condition for nonaccidental truth be recursive, where what it takes for a belief to satisfy the base clause is for it to be sensitive and where what it takes for a belief to satisfy the inductive clause is for there to be some other belief from which it inherits its responsiveness in the way captured by (a), (b), and (c) above, as follows:

R-Sensitivity. The truth of *S*'s belief that *p* is nonaccidental just in case *either S* wouldn't believe *p* were it not the case that *p or*, for some *q* the truth of which is sufficient for the truth of *p*, *S* believes that *q*, where the truth of this belief independently counts as nonaccidental and where *S* takes the grounds on which she has this belief to suffice also to ground her belief that *p*.

To see why this condition correctly classifies as nonaccidentally true my belief in, e.g., (1), consider the following. My belief in (2) is sensitive—it's caused, via my perceptual apparatus, by my having hands—and so that belief satisfies the base clause, which means it independently counts as nonaccidentally true. And the truth of (2) certainly is sufficient for the truth of (1): any world that's arranged in such a way that I have hands is thereby arranged in such a way that I don't falsely believe I do. Furthermore, I take the grounds on which I believe (2) to suffice also to ground my belief in (1): my visual experience as of hands is the grounds on which I believe the former, and I take that experience also to ground a belief in the latter (since I take the truth of the former to be sufficient for the truth of the latter). So (2) can play the role of *q* in the inductive clause, in which case my belief in (1) satisfies that clause.

More generally: By construction, R-Sensitivity, via its inductive clause, returns the verdict that an insensitive belief is nonaccidentally true in any case with the structure outlined above. And as we've seen, every case of insensitive but nonaccidentally true belief has this structure. Any insensitive but nonaccidentally true belief, then, is guaranteed to be such that this condition correctly classifies it as nonaccidentally true. So, since the condition, via its base clause, also classifies all sensitive beliefs as nonaccidentally true, we can be sure that it isn't overly strict: if a belief is in fact nonaccidentally true, the condition will classify it as such.

A clarification is in order here: R-Sensitivity is in some respects strikingly reminiscent of the recursive tracking analysis of knowledge developed by Roush (2005), according to which, very roughly, *S*'s belief that *p* (where *p* is contingent) counts as knowledge just in case either that belief tracks the truth (i.e., is sensitive[19]) or, for some *q*,

(i) *S* independently counts as knowing that *q*,
(ii) *q* implies *p*, and
(iii) *S* believes that *q* implies *p*[20] and is robustly disposed to arrange her beliefs in such a way as to respect this implication.[21]

So it's worth taking a moment to explain how the two accounts differ.

The most most obvious difference is that what Roush is offering is an analysis of knowledge, while R-Sensitivity is only a condition for nonaccidental truth. This is significant for reasons related to our discussion in §2 above: Roush's proposed analysis is vulnerable to certain powerful objections, objections that work by appeal to cases in which that analysis returns intuitively unacceptable verdicts about whether a given belief counts as knowledge, but since R-Sensitivity doesn't commit us to any particular knowledge claims, these sorts of objections have no purchase against that condition. Consider, for instance, the fact, pointed out by Brueckner (2012: sec. 6), that Roush's analysis, as a result of its radically externalist character—it counts as knowledge any belief meeting the above condition, regardless of whether the agent is being even minimally reasonable in holding that belief—commits her, absurdly, to the possibility of gaining knowledge by bootstrapping.[22] By contrast, R-Sensitivity, since it's only a condition for nonaccidental truth, doesn't share this radically externalist character and so doesn't commit us to this possibility: we're free, consistently with R-Sensitivity, to adopt a view of knowledge on which unreasonably held beliefs, even if nonaccidentally true, aren't knowledge, in which case we can say that beliefs arrived at via bootstrapping fail to be knowledge on the grounds that such beliefs are manifestly unreasonable.

[19] Actually, Roush, following Nozick, takes tracking to require adherence along with sensitivity—see fn. 3 above. This complication won't be relevant to our discussion.

[20] Here Roush includes an additional clause intended to ensure that *S*'s belief in this implication isn't based on mistaken inferences. Again, this complication won't be relevant to our discussion.

[21] Note that, though Roush often uses sensitivity counterfactuals in informal discussion, her official view is that the base clause here is to be stated not in terms of counterfactuals but in terms of conditional probabilities. Evaluating this modification of Nozick's approach is beyond the scope of this paper, but see Arló-Costa 2006 and Goldman 2009 for some concerns.

[22] Roush (2012: 263ff.) claims that this isn't so, but nothing she says explains how she can avoid the verdict that that, in a paradigmatic case of bootstrapping like Vogel's gas gauge case (see his 2000), the agent can come to have knowledge of at least the "weakened conclusion" mentioned by Brueckner (2012: 240): that the gas gauge is probably reliable. See also Melchior 2015: sec. 5.4.

A second difference between the accounts is one of scope: while R-Sensitivity applies to all beliefs, Roush's condition is intended to apply only to beliefs with contingent contents. Sensitivity conditionals for beliefs with necessarily true contents, Roush (2005: 43) says, are "very difficult to evaluate: who of us knows what happens in a logically impossible world, such as the one where a q that logically implies p is true but p is false?"; this, she suggests, is why a separate treatment of such beliefs is needed. The details of this treatment are beyond the scope of our discussion here, but note that, in the particular case where the necessary content in question is a proposition to the effect that some proposition implies another, what it takes for *S* to have knowledge of that implication, on Roush's view, is more or less just for her to satisfy (iii) above, to believe that the implication holds and to be disposed, in a very wide variety of circumstances, to arrange her beliefs so as to respect the implication. And the satisfaction of (iii), despite what Roush suggests, doesn't guarantee responsiveness: (iii) places restrictions on how *S*'s dispositions may be structured but not on what the explanation may be for why *S*'s dispositions are structured in that way, in which case *S* might satisfy that condition for reasons having nothing to do with whether the implication in fact holds.[23]

I take it that a unified treatment can do better on this score. Note first that, though Roush suggests that sensitivity conditionals for beliefs with necessarily true contents are unevaluable, they are in fact vacuously true, at least if we interpret sensitivity conditionals in the orthodox way, in terms of metaphysically possible worlds. If sensitivity conditionals are interpreted in this way, R-Sensitivity returns the verdict that, in any case in which a belief has a necessarily true content, it's not an accident, not just a lucky coincidence, that the belief is true. And though there's a sense in which this is intuitively the correct result—as Lewis has pointed out, "If what I believe is a necessary truth, then there is no possibility of being wrong" (1986: 114–115)—it must be admitted that this sort of trivial sensitivity is no guarantee of responsiveness: being responsive to certain features of the world involves doing different things in cases where the world has those features than in cases where it doesn't, and so there's no way to manifest such responsiveness if there just aren't any cases where the world fails to have the features in question. But remember: we might interpret sensitivity conditionals not in terms of metaphysically possible worlds but in terms of worlds that are epistemically possible in the sense described in §2, and if sensitivity conditionals are interpreted in *this*

[23] This is related to Brueckner's objection (see his 2012: sec. 2) that, on Roush's view, an agent can count as having knowledge of an implication even if her belief that the implication holds is based on wildly unreliable testimony. Roush (2012: sec. 2) claims that this objection is based on a confusion, but it's not obvious what that confusion is supposed to be—she herself acknowledges elsewhere (see her 2005: 143–145, especially the discussions of the Odysseus case and the premature logician case) that, on her account, an agent can know a logical truth by having the right dispositions even if she came to have those dispositions for paradigmatically irrelevant reasons.

way, they aren't rendered vacuous in cases in which a belief has a necessarily true content. Furthermore, there are reasons to think sensitivity conditionals, so interpreted, aren't especially difficult to evaluate, as I discuss elsewhere (see my 2020). Given the availability of this sort of interpretation, the metaphysical necessity of a belief's content doesn't prevent a sensitivity-based condition like R-Sensitivity from providing a nontrivial test of that belief's responsiveness.

There's one more difference to discuss here: where the q in R-Sensitivity's inductive clause is a proposition whose truth *is sufficient for* the truth of p, the q in Roush's analogous clause is a proposition that *implies* p. Whether this is a genuine difference, though, depends on what, exactly, the expressions "is sufficient for" and "implies" mean in this context. For my own part, I have to this point been relying on an intuitive grasp of "is sufficient for", but I can say here that what I have in mind is something like the traditional notion of conceptual or analytic entailment: what it is for the truth of q to be sufficient for the truth of p, in the sense relevant to R-Sensitivity, is just for the latter to *require nothing more of the world* than the former does. This is what explains why S's belief that p can inherit its responsiveness from her belief that q: since the truth of p requires nothing more of the world than the truth of q does, S's responsiveness to the features of the world relevant to the latter guarantees her responsiveness to the features of the world relevant to the former. As for Roush's use of "implies", it's not entirely transparent what she has in mind here, but I do want to note that she takes logical entailment to be a subspecies of implication. So, if facts about logical entailment are substantial facts requiring something of the world—as they are taken to be by proponents of the view McSweeney (2019) calls *metaphysical logical realism*, according to which there's a single true logic that captures the structure of mind-independent reality—then Roush's view makes unavailable the above explanation of why insensitive beliefs can inherit their responsiveness from sensitive ones.[24] This, I take it, is another way in which R-Sensitivity improves on Roush's condition.

Despite all these differences, though, the two conditions are similar enough that there's no reason here to work through individual cases of insensitive but nonaccidentally true belief, discussing in detail how R-Sensitivity handles, e.g., Gold-

[24]Suppose, for instance, that p is $\neg\neg q$, and suppose that S knows that q, believes that q implies p, and is disposed to arrange her beliefs so as to respect this implication. Then S counts as having knowledge of p, on Roush's view. But suppose metaphysical logical realism is true, so that the truth of q is not sufficient for the truth of p: the truth of p requires that q be true *and* that the world have certain logical features. Then it may well be that S, despite being responsive to the features of the world relevant to whether q, is not responsive to the relevant logical features and so is not responsive to the features of the world relevant to whether p. Roush's account ignores this possibility. Perhaps Roush thinks it's just obvious that facts about logical entailment *don't* require anything of the world—if so, it may well be that her use of "implies" is just equivalent to my use of "is sufficient for". But in that case her analysis of knowledge is hostage to a controversial view about the metaphysics of logic.

man's dachshund case (see his 1983: 84) or Williamson's distance underestimator case (see his 2000: sec. 7.5)—Roush (2005: sec. 2.3) has already explained how her own condition handles such cases, and exactly the same explanations are going to be available for R-Sensitivity. And in any case, we've already demonstrated, via a highly general argument, that R-Sensitivity *never* incorrectly classifies a belief as merely accidentally true. We can safely move on.

4 Sensitive but accidentally true beliefs

We've seen that the inductive clause of R-Sensitivity ensures that that condition isn't overly strict in the way Sensitivity is, but recall that Sensitivity is overly lenient in addition to being overly strict: there are beliefs that are sensitive but merely accidentally true. And R-Sensitivity's inductive clause is of no help with these counterexamples. Some other adjustment is needed.

It will again be useful to have a diagnosis of what goes wrong in these cases. Consider, then, a version of what is perhaps the most well-known such case: Kripke's red barn case (see his 2011: sec. 4(b)). Suppose I see a red barn in a field. Unbeknownst to me, though, I've stumbled into a strange region in which genuine barns are always painted red and in which counterfeit barns, constructed of papier-mâché and painted green, are extremely common, so that, if there were no genuine red barn in the field, a green papier-mâché counterfeit would have been placed there instead. Sensitivity, then, will classify my belief in

(4) There's a barn before me.

as merely accidentally true, since the nearest worlds where (4) isn't true are worlds where I see a counterfeit barn and so still believe (4). And this verdict seems to be correct: despite my visual experience of the barn, my belief isn't properly responsive to the features of the world that are relevant to whether there's a (genuine) barn before me. Presumably, though, I also believe

(5) There's a red barn before me.

And Sensitivity will return the verdict that the truth of this belief is *not* merely accidental: in the nearest worlds where (5) isn't true, what's in the field is a *green* counterfeit, and so, in such worlds, I don't believe (5). But this verdict, unlike the previous one, seems incorrect. After all, if I were responding in the right way to the features of the world relevant to the truth of (5), I certainly would also be responding to the features of the world relevant to the truth of (4), since the truth of (5) is sufficient for the truth of (4).[25] But we've already noted that I'm *not* responding in

[25] In a defense of Nozick's theory of knowledge, Adams and Clarke (2005: 215) accept that, if the

the right way to the features of the world relevant to the truth of (4). (To make this problem especially vivid, consider that R-Sensitivity, via its inductive clause, will have to classify my belief in (4) as nonaccidentally true if it classifies my belief in (5) as nonaccidentally true. But we've already determined that the former belief is merely accidentally true.) My belief in (5), then, appears to be merely accidentally true despite being sensitive.

What goes wrong in this case is the following. My belief in (5) requires for its truth that the world have two distinct sets of features: the facade of the object before me must be red, and the interior of the object must have the structure of a genuine barn. And that belief, though it's responsive to the former features, isn't responsive to the latter ones. But because (5)'s sensitivity conditional is true as long as the *nearest* worlds where (5) isn't true aren't worlds where I believe it, and because the example is engineered in such a way that the nearest worlds where (5) isn't true are worlds where the object's facade isn't red, the belief turns out to be sensitive despite its failure to be responsive to what the object's interior is like.

Broadly the same thing seems to be what goes wrong in other paradigm cases of sensitive but merely accidentally true belief as well. Consider, for instance, a case in which *S* believes that $p \wedge q$, and suppose that, in the nearest worlds where q isn't true, *S* doesn't believe that $p \wedge q$. Then if the nearest worlds where q isn't true are closer to the actual world than are the nearest worlds where p isn't true, *S*'s belief turns out to be sensitive, regardless of what proposition p is. Even if p is some unknowable proposition about the precise number of stars in the Milky Way galaxy, in which case the truth of *S*'s belief that $p \wedge q$ is certainly an accident, that belief is nevertheless going to be sensitive as long as q is some more modally fragile proposition such as the proposition that *S*'s left shoelace is untied.[26] Here again the problem seems to be that *S*'s belief, despite its failure to be responsive to certain features of the world that are required for its truth—in this case, those features relevant to the truth of p—nevertheless turns out to be sensitive just in virtue of its responsiveness to some other features of the world that are *also* required for its truth—in this case, those features relevant to the truth of q.

belief in (5) is knowledge, the belief in (4) must be knowledge as well. But they argue that, when relativization to method of belief formation is taken into account, it turns out that both beliefs are in fact knowledge—both beliefs are formed by the method of using the "reddish barnish look" of the object to determine what that object is, and if I were looking at a green counterfeit instead, neither belief would be formed by that method. As Becker (2012: 94) points out, though, this response "solve[s] the wrong problem": intuitively, the correct verdict is that *neither* belief is knowledge, not that *both* are.

[26] This feature of conjunctive beliefs, now well known (see, e.g., Roush 2005: sec. 3.1, Murphy and Black 2012: sec. 3, and Vogel 2012: sec. 2), was, as far as I know, first pointed out by Kripke (2011: sec. 4(b)), who called it "the absorption phenomenon". (Kripke's critique of Nozick's theory, thought it was only published in 2011, seems to have been in circulation, in manuscript form, since the 1980s.)

Let's say that, when a proposition requires for its truth that the world have two or more distinct sets of features,[27] it's an *agglomerative* proposition, and let's say also that, for each distinct set of features that the world must have in order for a proposition p to be true, the proposition that the world has those features is an *ingredient* of p.[28] Then we can describe the general structure shared by these cases as follows: *S* believes an agglomerative proposition p such that

(a) for some ingredient p_x of p, *S* doesn't believe p in the nearest worlds where p_x isn't true,
(b) the nearest worlds where p isn't true are worlds where p_x isn't true, and
(c) for some ingredient p_y of p, *S*'s belief that p isn't responsive to the truth of p_y.

And the reason Sensitivity returns the wrong verdicts in cases with this structure is simply that, by (c), *S*'s belief that p fails to be responsive to the features of the world relevant to its truth despite the fact that, by (a) and (b), that belief is sensitive.[29]

Much as before, it turns out, surprisingly, that *every* case of sensitive but merely accidentally true belief has this general structure. Let p be a proposition such that *S* has a sensitive but merely accidentally true belief that p. Then this belief, since it's merely accidentally true, isn't responsive to the features of the world relevant to its truth: there are some features of the world such that, first, the world having those features is required for the truth of p, and second, *S*'s belief that p fails to be responsive to those features. Let p_y be the proposition that the world has those features. Now, since *S*'s belief is sensitive, there are also some features of the world such that, first, the world having those features is required for the truth of p, second, the nearest worlds that don't have those features are worlds where *S* doesn't believe that p, and third, the nearest worlds where p isn't true are worlds that don't have those features. Let p_x be the proposition that the world has those features. Then p is an agglomerative proposition satisfying (a), (b), and (c) above.

What this tells us is that we can once again narrow our focus: we can ensure that our condition isn't overly lenient—i.e., that it *never* incorrectly classifies a belief as nonaccidentally true—just by ensuring that it returns the right verdicts in cases

[27] One might wonder just what this comes to. There isn't space here for a fully worked-out answer to this question, but the picture I have in mind is one on which, roughly, for a proposition to require for its truth that the world have multiple sets of features is for the truth of that proposition to require that multiple *parts of the world*—in the case of a proposition about the physical world, multiple *spatiotemporal regions* (where these might be divided up as finely as we like)—be arranged in a particular way.

[28] An agglomerative proposition, then, will just be a proposition with at least two ingredients. Note that not all agglomerative propositions are conjunctions—(5), for example, is not.

[29] If it's not immediately obvious why (a) and (b) guarantee the belief's sensitivity, consider that the truth of p requires the truth of p_x, in which case, by (b), the nearest worlds where p isn't true must also be the nearest worlds where p_x isn't true.

with this structure. What's needed, then, is just to generate a strengthened notion of sensitivity that guarantees responsiveness to *all* the features the world must have in order for a belief to be true, and there's a straightforward way to do so, based on a suggestion independently offered by Becker (2012: 95) and by Murphy and Black (2012: 36–37).[30] Let's say that *S*'s belief that *p* is *strongly sensitive* just in case every ingredient p_i of *p* is such that *S* wouldn't believe *p* were it not the case that p_i. Then, in any case in which *p* is an agglomerative proposition satisfying the conditions above, *S*'s belief in *p*—though sensitive, by (a) and (b)—will fail to be strongly sensitive, by (c). In the case of my belief in (5), for instance, the proposition that plays the role of p_y in (c)—namely, that the interior of the object before me has the structure of a genuine barn—is an ingredient of (5) such that, if that ingredient weren't true, I'd still believe (5), in which case my belief isn't strongly sensitive.

In short, in order to ensure that our condition isn't overly lenient, we need only design it so that it checks for strong sensitivity rather than its weaker cousin. Applying this insight to R-Sensitivity gives us the following:

Strong R-Sensitivity. The truth of *S*'s belief that *p* is nonaccidental just in case *either* every ingredient p_i of *p* is such that *S* wouldn't believe *p* were it not the case that p_i *or*, for some *q* the truth of which is sufficient for the truth of *p*, *S* believes that *q*, where the truth of this belief independently counts as nonaccidental and where *S* takes the grounds on which she has this belief to suffice also to ground her belief that *p*.

And we can confirm that this condition indeed isn't overly lenient. It's clear enough that, for reasons we've just noted, no accidentally true belief can satisfy the condition's base clause. Furthermore, given this fact about the base clause, there's no way for an accidentally true belief to satisfy the inductive clause either: insofar as *S*'s belief that *q* is indeed nonaccidentally true, *S*'s belief that *p*, where the truth of *q* is sufficient for the truth of *p* and where *S* takes the grounds on which she believes that *q* to suffice also to ground her belief that *p*, will be nonaccidentally true as well, for reasons discussed in §3 above. So Strong R-Sensitivity never incorrectly classifies a belief as nonaccidentally true.

The only remaining question here, then, is whether, by strengthening the base clause in the way we have, we've somehow introduced new counterexamples in the other direction and so have made our condition overly strict in a way that R-Sensitivity was not. And it's clear enough, I take it, that we've done no such thing. Strong R-Sensitivity is stricter than R-Sensitivity, to be sure, but not in a way that's going to generate new counterexamples—the only added strictness, after all, is that Strong R-Sensitivity guarantees, correctly, that when a belief requires for its truth

[30] Both Becker and Murphy and Black are discussing conjunctive propositions in particular, but the basic idea generalizes straightforwardly to agglomerative propositions of other sorts.

that the world have two or more distinct sets of features, the truth of the belief counts as nonaccidental only if the belief is responsive to *all* of those features. If this is right, Strong R-Sensitivity is is entirely free of counterexamples in either direction: it classifies as nonaccidentally true all and only those beliefs that are in fact nonaccidentally true.

5 Conclusion

Let's take stock. Our goal has been to give a satisfying account of the importance of sensitivity, an explanation of just why it is that sensitivity is an epistemologically significant property. We began by noting that the standard story here, on which sensitivity's primary epistemological role is as a constituent notion in the analysis of knowledge, is vulnerable to familiar counterexamples, and we also determined that, though attempts have been made to avoid these counterexamples by offering a patched version of a sensitivity-based analysis of knowledge, the lesson to be drawn from these counterexamples is in fact that sensitivity doesn't directly play a role in the analysis of knowledge at all. Instead, it has a crucial role to play, via its connection to responsiveness, in an account of what it takes for the truth of a belief to be nonaccidental, to be more than just a lucky coincidence. This may mean sensitivity has a role to play, albeit an indirect one, in the analysis of knowledge, since knowledge plausibly excludes accidental truth. But it also means sensitivity is of much broader importance—nonaccidental truth, after all, is highly epistemologically significant for a variety of reasons, only some of which have anything to do with the analysis of knowledge.

Taking this lesson on board allowed us to meet our goal. By examining in detail the structural features of certain problem cases, we were able to determine what those cases could tell us about the nature of sensitivity's connection to responsiveness, and this in turn allowed us to motivate and develop, in Strong R-Sensitivity, a necessary and sufficient condition for nonaccidental truth that's demonstrably counterexample-free. What explains why sensitivity is indeed an epistemologically significant property, then, is just that it has a crucial role to play in that condition.[31]

[31] This research was completed over the course of two research fellowships, generously funded by, respectively, the Lehigh University Department of Philosophy and the Austrian Science Fund (FWF), Grant No. P29716-G24. I'm also grateful to Gordon Bearn, Ricki Bliss, Patrick Connolly, Christopher Hill, David Christensen, Josh Schechter, Riki Heck, Paul Horwich, Jamie Dreier, Phil Galligan, Miquel Miralbés del Pino, Iain Laidley, Mary Renaud, Leo Yan, Zach Barnett, Han Li, Jonathan Courtney, Louis Gularte, Lorenzo Rossi, Julien Murzi, and several anonymous referees for helpful discussion of the ideas presented herein.

References

Adams, F. 1986. The function of epistemic justification. *Canadian Journal of Philosophy*, *16*, 465–492.

Adams, F., and Clarke, M. 2005. Resurrecting the tracking theories. *Australasian Journal of Philosophy*, *83*, 207–221.

Arló-Costa, H. 2006. Review of Sherrilyn Roush, *Tracking truth: Knowledge, evidence, and science*. *Notre Dame Philosophical Reviews*, *2006.07.20*.

Becker, K. 2006. Is counterfactual reliabilism compatible with higher-level knowledge? *Dialectica*, *60*, 79–84.

Becker, K. 2012. Methods and how to individuate them. In K. Becker and T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 81–97). Cambridge University Press.

Benacerraf, P. 1973. Mathematical truth. *Journal of Philosophy*, *70*, 661–679.

Bjerring, J. C., and Gundersen, L. B. 2020. Higher-order knowledge and sensitivity. *Canadian Journal of Philosophy*, *50*, 339–349.

Black, T. 2002. A Moorean response to brain-in-a-vat skepticism. *Australasian Journal of Philosophy*, *80*, 148–163.

Blome-Tillmann, M. 2017. Sensitivity actually. *Philosophy and Phenomenological Research*, *94*, 606–625.

Brueckner, A. 2012. Roush on knowledge: Tracking redux? In K. Becker and T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 229–241). Cambridge University Press.

Clarke-Doane, J. 2016. What is the Benacerraf Problem? In F. Pataut (Ed.), *Truth, objects, infinity: New perspectives on the philosophy of Paul Benacerraf* (pp. 17–43). Springer.

Collin, J. H. 2018. Towards an account of epistemic luck for necessary truths. *Acta Analytica*, *33*, 483–504.

Cross, T. 2010. Skeptical success. In T. S. Gendler and J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 3, pp. 35–62). Oxford University Press.

DeRose, K. 1995. Solving the skeptical problem. *Philosophical Review*, *104*, 1–52.

DeRose, K. 2010. Insensitivity is back, baby! *Philosophical Perspectives*, *24*, 161–187.

Engel, M., Jr. 1992. Is epistemic luck compatible with knowledge? *Southern Journal of Philosophy*, *30*, 59–75.

Field, H. H. 1989. *Realism, mathematics and modality*. Basil Blackwell.

Field, H. H. 2005. Recent debates about the a priori. In T. S. Gendler and J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 1, pp. 69–88). Oxford University Press.

Gettier, E. L. 1963. Is justified true belief knowledge? *Analysis*, *23*, 121–123.

Goldman, A. I. 1983. Review of Robert Nozick, *Philosophical explanations*. *Philosophical Review*, *92*, 81–89.

Goldman, A. I. 2009. Recursive tracking versus process reliabilism. *Philosophy and Phenomenological Research*, *79*, 223–230.

Kaplan, M. 1985. It's not what you know that counts. *Journal of Philosophy*, *82*, 350–363.

Kripke, S. A. 2011. Nozick on knowledge. In *Philosophical troubles: Collected papers* (Vol. 1, pp. 162–224). Oxford University Press.

Lewis, D. 1986. *On the plurality of worlds*. Basil Blackwell.

Luper, S. 2012. False negatives. In K. Becker and T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 207–226). Cambridge University Press.

McSweeney, M. M. 2019. Logical realism and the metaphysics of logic. *Philosophy Compass*, *14*, 1–10.

Melchior, G. 2015. The heterogeneity problem for sensitivity accounts. *Episteme*, *12*, 479–496.

Murphy, P., and Black, T. 2012. Sensitivity meets explanation: An improved counterfactual condition on knowledge. In K. Becker and T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 28–42). Cambridge University Press.

Nozick, R. 1981. *Philosophical explanations*. Harvard University Press.

Pritchard, D. 2005. *Epistemic luck*. Oxford University Press.

Pritchard, D. 2009. Safety-based epistemology: Whither now? *Journal of Philosophical Research*, *34*, 33–45.

Pritchard, D. 2015. Anti-luck epistemology and the Gettier problem. *Philosophical Studies*, *172*, 93–111.

Roush, S. 2005. *Tracking truth: Knowledge, evidence, and science*. Oxford University Press.

Roush, S. 2012. Sensitivity and closure. In K. Becker and T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 242–268). Cambridge 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. 73–83). MIT Press.

Schechter, J. 2017. No need for excuses: Against knowledge-first epistemology and the knowledge norm of assertion. In J. A. Carter, E. C. Gordon, and B. Jarvis (Eds.), *Knowledge first: Approaches in epistemology and mind* (pp. 132–159). Oxford University Press.

Sosa, E. 1999a. How must knowledge be modally related to what is known? *Philosophical Topics*, *26*, 373–384.

Sosa, E. 1999b. How to defeat opposition to Moore. *Philosophical Perspectives*, *13*, 141–153.

Sosa, E. 2002. Tracking, competence, and knowledge. In P. K. Moser (Ed.), *The Oxford handbook of epistemology* (pp. 264–286). Oxford University Press.

Topey, B. 2020. Realism, reliability, and epistemic possibility: On modally interpreting the Benacerraf-Field challenge. *Synthese*. Advance online publication. doi: 10.1007/s11229-020-02984-7

Unger, P. 1967. Experience and factual knowledge. *Journal of Philosophy*, *64*, 152–173.

Vogel, J. 1987. Tracking, closure, and inductive knowledge. In S. Luper-Foy (Ed.), *The possibility of knowledge: Nozick and his critics* (pp. 197–215). Rowman & Littlefield.

Vogel, J. 2000. Reliabilism leveled. *Journal of Philosophy*, *97*, 602–623.

Vogel, J. 2007. Subjunctivitis. *Philosophical Studies*, *134*, 73–88.

Vogel, J. 2012. The enduring trouble with tracking. In K. Becker and T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 122–151). Cambridge University Press.

White, R. 2010. You just believe that because…. *Philosophical Perspectives*, *24*, 573–615.

Williamson, T. 2000. *Knowledge and its limits*. Oxford University Press.