

IN DEFENSE OF SENSITIVITY**TIM BLACK AND PETER MURPHY****ABSTRACT**

The sensitivity condition on knowledge says that one knows that P only if one would not believe that P if P were false. Difficulties for this condition are now well documented. Keith DeRose has recently suggested a revised sensitivity condition that is designed to avoid some of these difficulties. We argue, however, that there are decisive objections to DeRose's revised condition. Yet rather than simply abandoning his proposed condition, we uncover a rationale for its adoption, a rationale which suggests a further revision that avoids our objections as well as others. The payoff is considerable: along the way to our revision, we learn lessons about the epistemic significance of certain explanatory relations, about how we ought to envisage epistemic closure principles, and about the epistemic significance of methods of belief formation.

Quite a while ago now, Robert Nozick suggested that one knows that P only if one's belief that P is sensitive to the truth, that is, only if one would not believe that P if P were false.¹ In spite of the fact that sensitivity seems an appropriate condition for knowledge, its troubles are well documented.² Keith DeRose has recently proposed a condition on knowledge that assigns sensitivity a leading role while it recruits supporting players that will help sensitivity avoid the difficulties it faces when it stands alone.³ We argue, however, that there are decisive objections to DeRose's proposal. Yet rather than simply abandoning his proposed condition, we uncover a rationale for its adoption, and then revise it so that it avoids objections like those we offer. The payoff is considerable: along the way to our revision, we learn lessons about the epistemic significance of certain explanatory relations, about how we ought to envisage epistemic closure principles, and about the epistemic significance of methods of belief formation.

1. DEROSE'S PROPOSAL

DeRose introduces his conditions on knowledge as a means to solving a problem. The problem arises for theories according to which one knows that P only if one sensitively believes that P. Again, this seems a reasonable requirement for knowledge. For one thing, as Nozick points out, the sensitivity condition, unlike the justified-true-belief account, allows us to get the right result in cases like those described by Gettier. Suppose, for example, that I am justified in believing that

(J) Jones, who works in my office, owns a Ford.

From this, I legitimately infer that

(S) someone in my office owns a Ford.

This belief is true, but it's true not because Jones owns a Ford, but because Brown, who works in my office but is a stranger to me, owns a Ford. I have a justified true belief that S, but I do not know that S. The sensitivity condition, however, yields the right result in this case: I don't know that S because I would believe that S even if it were true that no one in my office, including Brown, owns a Ford.⁴

But certain cases suggest a problem for these theories. The problematic cases are those in which we judge that I know that P even when I do *not* sensitively believe that P. DeRose provides two such cases. First, take my belief that

(F) I don't falsely believe that I have hands.

This belief is insensitive, for I would believe that F even if F were false. Nevertheless, we judge that I know that F.⁵ DeRose's second case centers around the belief that

(D) I'm not an intelligent dog who's always incorrectly thinking that I have hands.

Again, this belief is insensitive, but we judge that I know that D.⁶ These cases suggest a problem: while sensitivity seems to be a reasonable condition on knowledge, some *insensitive* beliefs seem to count as knowledge.

DeRose responds to this problem by replacing the pure sensitivity condition with a weaker condition, one that makes a place for sensitivity while still allowing us to know in the problematic cases. He says,

We *don't* ... judge ourselves ignorant of P where not-P implies something we take ourselves to know to be false, without providing an explanation of how we came to falsely believe this thing we think we know. Thus, *I falsely believe that I have hands* implies that I don't have hands. Since I do take myself to know that I have hands (*this* belief isn't insensitive), and since the above italicized proposition doesn't explain how I went wrong with respect to my having hands, I'll judge that I do know that proposition to be false.⁷

DeRose here suggests a condition on knowledge that is weaker than the pure sensitivity condition. Call it *weakened sensitivity*, or (WES):

(WES) S knows that P only if *either* S sensitively believes that P *or*, where $\sim P$ implies some Q and we take it that S knows that $\sim Q$, $\sim P$ fails to explain how S came falsely to believe that $\sim Q$.

DeRose's strategy, then, is to replace the pure sensitivity condition with a disjunctive condition, one that requires beliefs *either* to be sensitive *or* to meet a second condition. Much of the ensuing discussion will focus on the second disjunct of DeRose's condition, and on whether he is right to think that introducing it allows us to continue to accommodate sensitivity.

The second disjunct, which we call (EXP), has three components:

- (EXP) (i) Where $\sim P$ implies some Q and
(ii) we take it that S knows that $\sim Q$,⁸
(iii) $\sim P$ fails to explain how S came falsely to believe that $\sim Q$.

The introduction of (EXP) represents an improvement because we can now hand down the proper verdict in each of the two problematic cases. In the quoted passage, DeRose explains how (EXP) helps us to get the right result in the first problematic case. And it will help in the second case as well, for (i) *I'm an intelligent dog who's always incorrectly thinking that I have hands* implies *I don't have hands*, (ii) we take it that I know that I do have hands, and (iii) my being such a dog fails to explain how I came falsely to believe that I have hands. Thus, (WES), which represents the disjunction of (EXP) and the pure sensitivity condition, takes care of the problem that was suggested by DeRose's cases.

We are attracted to the idea that sensitivity is conceptually related to knowledge, and we are also inclined to think that DeRose's proposal is on the right track. Nevertheless, DeRose does little to make it clear exactly *why* sensitivity and (EXP) should go together in the way he proposes. What's worse, (EXP) seems merely to have been cobbled onto sensitivity in order to take care of certain unfavorable cases. So, until we discover a sound philosophical motivation for its introduction, (EXP) will seem *ad hoc*. In the next section, we argue that there *is* a sound motivation for its introduction.

2. MOTIVATIONS FOR THE INTRODUCTION OF (EXP)

When the pure sensitivity condition stands on its own and we fail to introduce (EXP) in the way DeRose does, or at least in a similar way, we are faced with a dilemma. Avoiding this dilemma

provides a powerful motivation for (EXP)'s introduction. We fill in the details of this line of thought below.

2.1 How sensitivity leads to a dilemma

Suppose that, in my present environment, I believe that P and P is true. Yet suppose that I would believe that P even if I were in an epistemically relevant environment in which P is false.⁹ Let's say of this case that the method I use in forming a belief as to whether P is not appropriately responsive to my environment: it will lead me to believe that P in environments in which P is true, but it will also lead me to believe that P in certain environments in which P is false. Moreover, it seems quite reasonable to think that I don't know that P when the method I use in forming my belief does not respond appropriately to my environment. The sensitivity condition on knowledge captures this sentiment. We can express the spirit of that condition in this way: S knows that P only if

(SEN) the method she uses in forming that belief is appropriately responsive to her environment, in that it can distinguish environments in which P is true from certain epistemically relevant environments in which P is false.

Yet when we evaluate certain beliefs in terms of (SEN), we run into trouble. As we will see, the natural way to avoid this trouble is to emend (SEN) by disjoining it with a condition like (EXP). To see the trouble, consider my belief that

(F) I don't falsely believe that I have hands.

Even though I take this belief to amount to knowledge, it fails to meet the condition on knowledge expressed in (SEN). Whatever method I use in coming to believe that F will respond inappropriately to my environment: As DeRose notes, even if I were in an environment in which

I *do* falsely believe that I have hands, I would still believe that I *don't* falsely believe that I have hands.¹⁰ Consider, however, my belief that

(H) I have hands.

I take this belief, too, to amount to knowledge, and it *does* meet the condition expressed in (SEN). The method that I use in forming my belief that H responds appropriately to my environment: in no epistemically relevant environments in which H is false do I believe that H.

But we're left with trouble after we filter these evaluations through a very plausible epistemic principle, namely, the closure principle, according to which knowledge is closed under known logical implication:

(CLO) If S knows that P, and if S knows that P implies Q, then S knows that Q.¹¹

Now, given that I know that H implies F, we get the following instance of (CLO):

(CLO₁) If I know that H, then I know that F.

The trouble is now upon us: The sensitivity condition allows the antecedent of (CLO₁) to be true, but renders its consequent false. This means that (CLO₁), and hence the very plausible (CLO), is false. We are thus faced with a dilemma. If we cling to the intuition that I know that F, we must reject (SEN)'s being a condition on knowledge. Yet (SEN) seems quite plausible. If, on the other hand, we cling both to (SEN) and to the highly plausible (CLO), we must reject our entrenched intuition that I know that H.

2.2 Avoiding the dilemma with (EXP)

Yet we can avoid this dilemma. If we include a claim like (EXP) among the conditions for knowledge, we can save (CLO), our firm intuitions about what we know, and (SEN)'s status as a condition on knowledge. We can begin to see this by considering (CLO₁). Note that if I come to

believe that F through an inference, in accordance with (CLO₁), that involves my perceptual belief that H, then my belief that F will have been produced by an inference from a *perceptual* belief.¹² In such cases, when we successfully reason in accordance with (CLO₁), both its antecedent and its consequent will make reference to the same belief-forming method.

In light of this, let's now determine whether S knows that F. Notice first that since S's perceptual belief that H is both sensitive and true, we ought not deny that it amounts to knowledge. Moreover, we take ourselves to be competent and reliable when it comes to performing simple inferences. We therefore have little reason to doubt that S can know that F on the basis of an inference, in accordance with (CLO₁), from her perceptual belief that H.

Yet recall that S's belief that F is *insensitive*. Perhaps we should take this to show, contrary to the results of the preceding paragraph, that S *cannot* know that F on the basis of an inference from her perceptual belief that H. We maintain, however, that the insensitivity of S's belief that F gives us no reason to draw this conclusion. Since S insensitively believes that F, there are environments in which F is *false* but in which she nevertheless believes that F is *true*. In these environments, H is false as well (since H implies F). Yet there is nothing in our description of these environments that *explains how perception* would lead S to hold the false belief that H. So far, we have described these environments simply as environments in which she falsely believes that she has hands. Clearly, such a description fails to explain how S, by using perception, would come to hold the false belief that H. This failure of explanation leaves open the possibility that perception, the method responsible for S's belief that H, responds appropriately to its environment, thus producing beliefs as to whether H that are sensitive. This allows us to continue to claim that S knows that H.

Furthermore, this failure of explanation allows us to continue to maintain that S knows that *F*. We have just seen that, in spite of the fact that S's belief that *F* is insensitive, we may continue to maintain that she knows that *H*. Moreover, nothing at all casts doubt on S's ability to perform simple inferences (like the one from *H* to *F*). Hence, the failure of explanation allows us to maintain that S knows that *F*: she knows that *F* at least partly because her being in certain environments in which *F* is false *fails* to explain how she would come to hold the false belief that *H*. This, of course, is just the requirement expressed in (EXP), according to which, in this instance, I can know that *F* because $\sim F$ implies $\sim H$, we take it that I know that *H*, and $\sim F$ *fails* to explain how I came falsely to believe that *H*. It seems quite reasonable, then, to include (EXP) among the conditions on knowledge. Moreover, doing so allows us to avoid the dilemma we noted earlier. Even though S's insensitive belief that *F* did not count as knowledge when (SEN) was the lone condition on knowledge, we have now seen that disjoining (EXP) to (SEN) allows us to take S to know that *F*. Furthermore, disjoining (EXP) to (SEN) allows us to retain the intuition that S knows that *H* – the sensitivity of her belief that *H* allows it to count as knowledge. (SEN) may therefore keep its place in the conditions on knowledge. Finally, since there is now nothing to keep us from claiming that S knows both that *H* and that *F*, we may also retain (CLO), which tells us in this case that S knows that *F* if she knows that *H*. So, including (EXP) among the conditions for knowledge helps us to steer clear of the dilemma we noted earlier.

3. REVISING (WES) IN LIGHT OF CRITICISMS

Even though there is a sound philosophical motivation for the introduction of (EXP), there are problems with (WES) as it stands. First, there are decisive counterexamples to (WES).

Consider, for instance,

- (C) I am a handless creature who, as a result of using an undetectable (by me) and unreliable form of sonar, comes mistakenly to believe that I have hands.¹³

Now, just as I know that I'm not an intelligent dog who's always incorrectly thinking that I have hands, I know that I'm not the sort of creature described in C. That is, I know that C is false.

Furthermore, I insensitively believe that C is false – if C were true, I would still believe it to be false. The question, then, is whether (EXP) is satisfied. If not, our counterexample to (WES) succeeds.

We can see only two implications of C that are relevant to the evaluation of (EXP).¹⁴

Let's first consider the implication that *there is available to me an undetectable and unreliable form of sonar*. This fails to satisfy component (ii) of (EXP). For the method's undetectability leaves me with no way of knowing either that I possess such a method or that I lack such a method. Thus, (EXP) is false in this case.¹⁵

Next, consider the implication that *I don't have hands*. Here, components (i) and (ii) are true: (i) C implies that I don't have hands, and (ii) we take it that I know that I do have hands. However, (iii) is *false* since C *does* explain how I would come mistakenly to believe that I have hands: in worlds in which C is true, I falsely believe that I have hands because I use the unreliable method described in C in forming a belief as to whether I have hands. So (EXP) is false in this second case as well. This means that we now have a successful counterexample to (WES): I *know* that C is false; I *insensitively* believe that it's false; and my belief that C is false *fails* to satisfy the conditions in (EXP).

Now that we have a counterexample to (WES), it seems that the introduction of (EXP) cannot help those who, like us, hope to find a place for sensitivity among the conditions for knowledge. We think, however, that we can revise (WES) so as to preclude the possibility of such counterexamples. Our revision is motivated by a diagnosis of the problems that make (WES) susceptible to our counterexample.

We begin our diagnosis by noting that we can now see, in light of the connection we brought out above between methods and explanation, how my being in certain situations fails to explain my coming to hold certain false beliefs. Consider DeRose's dog hypothesis, or D*:

(D*) I'm an intelligent dog who's always incorrectly thinking that I have hands. My being a dog of this sort fails to explain how I would come to hold the false belief that I have hands. Note first that D* specifies no method. We are therefore in the dark about which belief-forming method the dog uses. Given this, my being the sort of dog described in D* does not explain how, by using the method that I actually use in coming to believe that I have hands, I would come *mistakenly* to believe that I have hands. And since we take it that I know that I have hands, we may continue to maintain that I know D* to be false.

But we can now also see, in light of the same connection, how my being in certain situations *does* explain my coming to hold certain false beliefs. Consider my belief that C is false. Now, since C specifically describes the method that I use in that situation – as an undetectable and unreliable form of sonar – C precludes the possibility that my method responds appropriately to its environment. We therefore have good reason to suppose that the belief-forming method described in C would lead me mistakenly to believe that I have hands. Thus, my being in the situation described in C explains why I would falsely believe that I have hands.¹⁶ Unfortunately, since I *insensitively* believe that C is false, and since my being the sort of creature

described in C explains why I would falsely believe that I have hands, (WES) now yields the claim that I don't *know* that C is false. This generates a counterexample to (WES), however, since we take it that I *do* know that C is false.

What has gone wrong here? According to (WES), the conditions for knowing that P demand that when I use a particular method in coming to believe that P, we should not be able to explain, in terms that make reference to that very method, how I would come falsely to believe certain epistemically relevant implications of P's not being the case. We claim, however, that it is not enough to make this demand. Even though it demands of the conditions on knowledge that they make reference to the very same belief-forming methods, (WES) can be exploited by our counterexample because the methods in question *are not restricted to those that we actually use in coming to hold the relevant beliefs*. My using the method specified in C – the undetectable and unreliable form of sonar – explains why I would falsely believe that I have hands. But this belief-forming method is utterly different from the method I actually use in forming the belief that I have hands. This leaves room for our counterexample to (WES) – we get such a counterexample because we take ourselves to know, on the basis of the method that we *actually* use, that we are not creatures like the one described in C. Thus, to preclude the possibility of such counterexamples, the conditions on knowledge should make it clear that the methods involved in knowing that P are just those that we actually use in coming to believe that P. This thought motivates the following revised version of (WES), one designed to patch holes that let counterexamples like ours leak through:

(WES*) S knows via *m* that P only if *either* S sensitively believes via *m* that P *or*, where $\sim P$ implies some Q and we take it that S knows that $\sim Q$, $\sim P$ fails to explain how S would come falsely to believe via *m* that $\sim Q$ (where *m* is

the belief-forming method that S actually uses in coming to believe that P).¹⁷

(WES*) is not susceptible to our counterexample. We take ourselves to know, on the basis of familiar perceptual belief-forming methods (or on the basis of an inference from beliefs formed via those methods), that C is false. Yet C's being true does *not* explain how we might come mistakenly to believe, on the basis of the methods that we actually use, that we have hands. Thus, we may continue to judge that we know that C is false.¹⁸

Opponents of our proposal might try to generate a counterexample to (WES*) by revising C. They might try

(C*) I am a handless creature who, as a result of using my actual, unreliable belief-forming methods, comes mistakenly to believe that I have hands.

C* improves on C by making reference to the method that I actually use in forming the belief that I have hands, and it seems obvious that I know that C* is false. Nevertheless, there is no counterexample to (WES*) here, for either C* fails to explain how I would come falsely to believe that I have hands, or my belief that C* is false fails to count as knowledge. The problem lies in C*'s appeal to my *actual* unreliable belief-forming method. Now, my actual method is either reliable or unreliable. If it's reliable, then C* entails something that is necessarily false, namely, that my actual reliable method is unreliable.¹⁹ In this case, then, C* is unfit to play *any* explanatory role. On the other hand, if the method that I actually use is *unreliable*, then it is no longer obvious that I know that I have hands, and therefore no longer obvious that I know that C* is false. In fact, it might now seem obvious that I know *neither* of these things, for it's quite plausible to suppose that these beliefs must be produced by reliable methods if they are to count as knowledge.

We have seen, then, that (WES*) takes care of our proposed counterexample to (WES). Still, (WES) faces another problem. As Timothy Williamson suggests, it is insufficiently general.²⁰ That is, it fails to account for all cases of knowledge. Here is Williamson's example: Let my present situation be one in which it appears to me that I am sitting in front of a computer screen in my office. Now imagine a case, β , in which I am a brain-in-a-vat but in which it appears to me that I am climbing a mountain. Williamson claims that I know in my present situation that I'm not in β , for "things do not appear to me at all as they would in β ."²¹ I do *not*, however, sensitively believe that I'm not in β : I would believe that I was not in β even if I were in β .²² Thus, if (WES) is to account for my knowing that I'm not in β , (EXP) must be satisfied. Williamson concentrates on the following instance of (EXP):

- (i $_{\beta}$) Where my being in β implies
 - (M) it now appears to me as if I'm climbing a mountain, and
- (ii $_{\beta}$) we take it that I know that $\sim M$,
- (iii $_{\beta}$) my being in β fails to explain how I came falsely to believe that $\sim M$.

Although (i $_{\beta}$) and (ii $_{\beta}$) are (obviously) true, (iii $_{\beta}$) is not. For the truth of (iii $_{\beta}$) requires the existence of the unexplained false belief that $\sim M$.²³ But, in β , I come *truly* to believe that M : in β , I come to hold the belief that it now appears to me as if I'm climbing a mountain, and that belief is true. Thus, absent from β is the false belief that $\sim M$. Since the false belief that should fail to be explained by my being in β is nowhere in β to be found, (iii $_{\beta}$) is false, and so (EXP) is not satisfied. It follows that (WES) does not account for my knowing that I'm not in β .

Unlike (WES), however, (WES*) accounts for my knowing that I'm not in β . For my being in β fails to explain how I would come falsely to believe, via a particular belief-forming method, that it does not appear to me that I'm climbing a mountain. Recall that ' β ' is equivalent

to the conjunction ‘I am a brain-in-a-vat and it appears to me that I am climbing a mountain’, which means that ‘I’m not in β ’ is equivalent to the disjunction ‘Either I am not a brain-in-a-vat, or it does not appear to me that I am climbing a mountain’. On our proposal, it is crucial that we identify the method that produces my belief that I’m not in β . It seems, moreover, that my belief that I’m not in β is ultimately based on *introspection* – I come to believe via introspection that it does not appear to me that I am climbing a mountain, and then I come to believe that I’m not in β on the basis of an inference from this introspective belief. What’s more, it seems that I come to *know* that I’m not in β in this way. Happily, unlike (WES), (WES*) can account for this piece of knowledge. For, even though my belief that I’m not in β is insensitive, my being in β *does* fail to explain how I would come falsely to believe via introspection that it does not appear to me that I’m climbing a mountain. Even in β , introspection will lead me *truly* to believe that it *does* appear to me that I’m climbing a mountain. (WES*) therefore allows me to know via introspection that I’m not in β , and our proposal yields the proper result in this case.

Moreover, Williamson’s example helps us to see that one of (WES*)’s crucial restrictions is in place for good reason, namely, its restriction that the conditions for S’s knowing that P make reference only to the method that S actually uses in forming her belief that P. Once again, ‘I’m not in β ’ is equivalent to the disjunction ‘Either I am not a brain-in-a-vat, or it does not appear to me that I am climbing a mountain’. And it’s the insensitivity of my belief in the *former* disjunct, a belief that is produced by perception, that leads to the insensitivity of my belief that I’m not in β . It therefore seems that perception is responsible for the insensitivity of my belief that I’m not in β . Yet, as we’ve seen, a different belief-forming method – *introspection* – produces my belief that it does not appear to me that I’m climbing a mountain. Perception’s failure, though, does nothing at all to impugn introspection’s ability to produce

beliefs that count as knowledge. This is confirmed in the present case. For, in the nearest β world, introspection does just fine with respect to whether it appears to me that I'm climbing a mountain. In that world, it appears to me that I'm climbing a mountain, and introspection leads me, quite appropriately and truly, so to believe. Thus, in spite of perception's breakdown in this case, we still have every reason to believe that introspection can produce beliefs that will count as knowledge. Only breakdowns in introspection give us reason to question introspection's ability to produce such beliefs. In this way, then, this case helps us to see that there is good reason for (WES*)'s restriction on methods.

(WES*) helps with other difficulties as well. In his discussion of (WES), Williamson calls on Alvin Goldman's dachshund case.²⁴ In this case, I see a dachshund and believe that there is a dachshund before me. In the nearest counterfactual situation in which there is not a dachshund before me, there is instead another kind of dog before me, but I do not mistake this dog for a dachshund. So, my belief that there is a dachshund before me is *sensitive*. However, from my belief that there's a dachshund before me, and from my knowledge that all dachshunds are dogs, I competently infer that there is a dog before me. This belief, however, is *not* sensitive since in the nearest counterfactual situation in which there is no dog before me, there is a *wolf* before me, and I mistake that wolf for a dog. In this case, then, (WES) suggests that I know that there is a dachshund before me, but that I *fail* to know that there is a dog before me.

On our proposal, however, we do not encounter this problem, for I meet (WES*)'s conditions with respect to my belief that there is a dog before me. Even though my belief is insensitive, I can know that there is a dog before me since (i) *there is no dog before me* implies *there is no dachshund before me*, (ii) we take it that I know that there is a dachshund before me, and (iii) the counterfactual situation in which there is a wolf rather than a dog before me *does* fail

to explain how I might come falsely to believe that there is a dachshund before me. For, in the situation in which there is a wolf before me, though I mistake the wolf for a dog, I do not mistake it for a dachshund.

Perhaps the dachshund case suggests other difficulties, however. Williamson maintains that the case helps to show that conditions on knowledge that appeal to methods of belief formation must *finely individuate* those methods.²⁵ That is, rather than coarsely individuating a belief-forming method – say, as vision – the dachshund case suggests that we must individuate methods more finely – for example, as vision-in-a-particular-set-of-circumstances-*C*, or as inference-from-a-particular-set-of-visual-experiences-*E*.²⁶ Consider a methodized version of Nozick's sensitivity condition:

(NSC_M) If S knows via method *m* that P, then if P were false and S were to use *m* to arrive at a belief as to whether P, then S would not believe via *m* that P.²⁷

Williamson maintains that (NSC_M) is false when belief-forming methods are coarsely individuated. For, when I see a dachshund before me, I know via vision that there is a dog before me, but if there were a *wolf* before me (and no dog) and I were to use vision to arrive at a belief as to whether there is a dog before me, I would nevertheless believe via vision that there is a dog before me. Thus, in order to preserve (NSC_M), it seems that its advocates must resort to finely individuated belief-forming methods.

Our proposal, however, suggests a strategy for dealing with this sort of objection, for (WES*) seems to handle the dachshund case even when belief-forming methods are coarsely individuated. This is one very clear advantage that our proposal has over Nozick's methodized sensitivity condition. Suppose that vision, coarsely individuated as such, is responsible both for

my belief that there is a dog before me in the actual case (in which there is a dachshund before me) and for my belief that there is a dog before me in the counterfactual case (in which there is a wolf before me). Even so, if there were a wolf before me, vision would not lead me to mistake that wolf for a *dachshund*. Thus, the counterfactual situation fails to explain how I would come falsely to believe via vision that there is a dachshund before me. (WES*) therefore allows me to know via vision that there is a dog before me. Since this is true, the dachshund case fails to show that proponents of (WES*) must resort to finely individuated belief-forming methods. One benefit of our proposal, then, is that it might help counterfactualists to sidestep objections that proceed from the idea that they must finely individuate belief-forming methods. Of course, we have more work to do before we can say how, or even whether, our proposal can forestall such objections. The details must come later, however, since we haven't the space to deal with these issues here. Nevertheless, it certainly seems that our proposal holds promise in this regard.

4. CONCLUSION

We have now seen that in spite of the existence of serious and decisive objections to (WES), there is a rationale for adopting conditions on knowledge like those it expresses, which include conditions like those expressed in (EXP). In spite of this, however, certain instances of (EXP) are false, making (WES) susceptible to counterexample. We therefore suggest a further restriction, namely, a restriction to just those belief-forming methods that we actually use in coming to hold certain beliefs. This restriction suggests a revision of (WES) that both incorporates (EXP) and avoids the serious and decisive objections to (WES). More importantly, it teaches the valuable lesson that any adequate account of knowledge should make reference only to the belief-forming methods that we actually employ in coming to hold certain beliefs.²⁸

NOTES

¹ R. Nozick, *Philosophical Explanations* (Cambridge, MA: Harvard University Press, 1981).

² For a nice collection of critical and diagnostic papers, see S. Luper-Foy, ed., *The Possibility of Knowledge: Nozick and His Critics* (Totowa, NJ: Rowman and Littlefield, 1987).

³ See K. DeRose, 'Solving the Skeptical Problem', in K. DeRose and T. A. Warfield (eds), *Skepticism: A Contemporary Reader* (Oxford University Press, 1999), pp. 183-219. DeRose (1999) employs his proposed condition in providing a contextualist response to skeptical arguments.

⁴ See Nozick (1981), p. 173. (For Gettier's counterexamples to the justified-true-belief account of knowledge, see E. L. Gettier, 'Is Justified True Belief Knowledge?', *Analysis* 23: 121-123.) According to Nozick, the sensitivity condition helps with other troubling cases as well. Suppose that Henry is driving through a region populated by *papier-mâché* barn facsimiles. Even if Henry spies a bona fide barn—and so even if he believes, justifiedly and truly, that this is a barn—he doesn't *know* that it's a barn because he would believe that it's a barn even if it had been a *papier-mâché* facsimile (see Nozick (1981), pp. 174-175).

⁵ See DeRose (1999), pp. 196-7.

⁶ See DeRose (1999), pp. 196-7.

⁷ DeRose (1999), p. 197.

⁸ That is, we take it that S knows *in the actual world* that $\sim Q$. To see that this is what DeRose has in mind, see the passage quoted above (from DeRose (1999), p. 197) and the passage in which he discusses the intelligent dog hypothesis (DeRose (1999), pp. 196-7).

⁹ We do not endorse any particular account of epistemic relevance, and nothing in what follows will depend on the truth of any particular such account.

¹⁰ See DeRose (1999), pp. 196-7.

¹¹ There are, perhaps, more plausible, albeit more complicated, formulations of the closure principle. Nevertheless, since nothing in our argument turns on our choosing any particular formulation, we work here with a less complicated one.

¹² This point is fully general. Thus, if I come to believe that Q through an inference, in accordance with the relevant instance of (CLO), that involves my *m*-produced belief that P (where *m* is a belief-forming method), then my belief that Q will have been produced by an inference from an *m*-produced belief. One might think, however, that I can know *through deduction alone* that Q. Yet this is not the case. I know through deduction only if I have knowledge of all the premises involved in the deduction. (Gettier taught us this lesson: lucky true beliefs (which do not, of course, count as knowledge) that serve as premises in an argument for the true conclusion that Q do *not* allow us to know that Q.) Since knowledge of a conclusion requires knowledge of each of the argument's premises, the method by which a conclusion is known will include at least the method(s) by which the premises are known (as well as, probably, deduction).

¹³ The form of sonar that I use is undetectable by me because it generates no phenomenology and has no qualitative feel whatsoever. This counterexample was suggested by Timothy Williamson's criticism of DeRose's proposal. See T. Williamson, *Knowledge and its Limits* (Oxford: Oxford University Press, 2000), especially p. 159.

¹⁴ DeRose suggests, in the passage quoted above, that (EXP) will be true – and hence that (WES) will be true (in cases in which our beliefs are insensitive) – when there is *just one* Q that makes each of (EXP)’s three components true. Thus, if we are to provide a successful counterexample to (WES), we must provide a case in which there is *no* Q that will make each of (EXP)’s components true.

¹⁵ Even if we somehow still manage to take it that we know that this sort of sonar is *not* available to us, hence allowing the implication to satisfy component (ii) of (EXP), the implication will nevertheless fail to satisfy component (iii). For if I possess this sort of sonar but nevertheless believe that I lack it, then the method’s *undetectability* will explain why I’ve come falsely to believe that I lack the method.

¹⁶ Furthermore, in the situation described in C, I falsely believe that I’m not in that situation; that is, I falsely believe that C is false. And I believe that C is false either on the basis of my undetectable and unreliable form of sonar – perhaps I believe that C is false because my sonar seems to inform me that I have hands – or on the basis of an inference from sonar-induced beliefs – perhaps I believe that C is false on the basis of an inference from my sonar-induced belief that I have hands.

¹⁷ Nozick introduces methods into his tracking theory in order to deal with a case in which a grandmother knows that her grandson is well. If he had he *not* been well, however, she nevertheless would have believed that he *was* well. For, in that case, the boy’s parents would not have allowed her to see him and, to spare her upset, they would have told her that he was well. This does not mean, however, that the grandmother does not know that her grandson is well when she sees him (see Nozick (1981), p. 179). Nozick’s invoking belief-forming methods in

this way helps him to deal with problematic cases like the grandmother case. Yet Nozick neglects to provide any more thoroughgoing philosophical motivation for introducing methods. We, on the other hand, do provide such a motivation. In that respect, then, our introduction of belief-forming methods stands on a firmer philosophical foundation than does Nozick's. Moreover, even though Nozick does assign a role to methods (see Nozick (1981), pp. 179-185), he never recognizes the need to introduce a clause like (EXP), one that assigns a certain *explanatory* role to methods. Thus, he never recognizes the need to move to something like our (WES*). This is another respect in which our proposal markedly improves on Nozick's, allowing our proposal to account for richer and more complex cases. Furthermore, our account has at least one other advantage over Nozick's—it treats troublesome cases in a manner that allows us to avoid the principal objection to Nozick's account, namely, that it's committed to a denial of (CLO). For these reasons, our more sophisticated sensitivity account, which includes a condition like (EXP), represents an improvement over Nozick's original account.

¹⁸ We do not here tackle the question whether (EXP)'s introduction also helps us to avoid skeptical troubles that stem from a consideration of, for example, brain-in-a-vat hypotheses. So far as we can tell, (EXP)'s introduction leaves both the skeptic and the non-skeptic with resources. This is another welcome advantage of our proposal. So, the skeptic might argue that (EXP)'s introduction allows us to maintain that I *don't* know that I'm not a brain-in-a-vat. For not only do I insensitively believe that I'm not a brain-in-a-vat, but my being a brain-in-a-vat also explains how I would come falsely to believe the denial of certain propositions that are entailed by my being a brain-in-a-vat. For example, my being a brain-in-a-vat who is electronically stimulated in the appropriate ways explains how I would come falsely to believe

that I have hands. The *non*-skeptic, on the other hand, might argue that the belief-forming methods of cognizers like us do not include those of brains-in-vats. If this is the case, then my being a brain-in-a-vat fails to explain how I would come mistakenly to believe, on the basis of methods *that I actually use*, that I have hands. (EXP)'s introduction gives us the room, therefore, to maintain that I *do* know that I'm not a brain-in-a-vat. (For further details regarding this non-skeptical strategy, see *<an author's paper>*.)

¹⁹ C* entails a necessary falsehood because the method in question is picked out by using 'actual'.

²⁰ See Williamson (2000), p. 158.

²¹ Williamson (2000), p. 158.

²² Williamson (2000), p. 158, says, "No matter what my situation, I cannot sensitively believe *p*," where *p* is the proposition that I am not in β .

²³ As Williamson seems to recognize, there are readings of (iii) $_{\beta}$ on which its truth does not require the existence of the unexplained false belief that $\sim M$. Indeed, one might contend that (iii) $_{\beta}$ is true precisely *because* the false belief that $\sim M$ is absent from β , for whenever an explanandum (such as my false belief that $\sim M$) is absent, nothing (including my being in β) can explain its being present. One way to defend (EXP) would be to argue for such a reading of (iii) $_{\beta}$. Rather than proceeding in this way, however, we defend (EXP) in a different way.

²⁴ See Williamson (2000), pp. 153, 159. The case appears originally in A. Goldman, 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73: 771-791.

²⁵ See Williamson (2000), pp. 153-4.

²⁶ It might not be a problem that conditions on knowledge must finely individuate belief-forming methods. Whether it is a problem depends on various debates about the generality objection. According to that objection, which is usually leveled against process reliabilists but which is equally relevant to the sensitivity theory under consideration here, finely individuating methods leads to undesirable consequences. For example, if we individuate methods too finely, the beliefs formed via those methods can't but be sensitive. Consider, for example, a case in which my belief that there's a computer screen before me is produced by the finely individuated method *vision-in-this-very-set-of-circumstances*. There's simply no question here that if there were no computer screen in front of me and I were to use *vision-in-this-very-set-of-circumstances* to arrive at a belief as to whether there's a computer screen in front of me, then I would not believe via *vision-in-this-very-set-of-circumstances* that there's a computer screen in front of me. One complaint, therefore, is that sensitivity accounts do no real work – they are useless – if they are forced to individuate methods finely.

²⁷ See Nozick (1981), p. 179, and Williamson (2000), p. 153.

²⁸ For their helpful comments on earlier drafts of this paper, we wish to thank Albert Casullo, Ram Neta, Duncan Pritchard, Guy Rohrbaugh, Jonathan Schaffer, Elka Shortsleeve, an audience at the 2004 Inland Northwest Philosophy Conference, and two anonymous referees for this journal.