

*Sensitivity meets explanation: an improved
counterfactual condition on knowledge*

Peter Murphy and Tim Black

Robert Nozick (1981) famously suggests 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 A principal objection to this condition is that it forces us to deny a very plausible epistemic closure principle, according to which

(Closure) If S knows that P, and if S knows that P implies Q, then S knows that Q.²

Ernest Sosa, Saul Kripke, and Timothy Williamson have each leveled this sort of objection, and each has done so by presenting one or more cases in which the sensitivity condition is supposed to force us to deny very obvious instances of Closure. They then urge us to abandon the prospects for a sensitivity condition.

But a bit of caution is in order here. We shouldn't give up on a sensitivity condition if, as we will argue, some such condition is defensible. We take our cue from a suggestion Keith DeRose makes in his seminal paper, "Solving the skeptical problem." There, DeRose points out that:

We don't so 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.³

The revised condition that DeRose gestures toward allows S to know P despite insensitively believing P. How? By sensitively believing some other proposition, Q, from which S can deduce P, where this holds at

¹ See also Dretske (1971). For a nice collection of critical and diagnostic papers on Nozick's views about the condition, see Luper-Foy (1987c).

² There are 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.

³ DeRose 1995, 197.

least as long as the following is true: the not-P situation that gets in the way of S's having a sensitive belief that P fails to explain how S might come to hold the false belief that Q. To say that the relevant not-P situation fails to explain how S *might* come to hold the false belief that Q is, we take it, to say that it fails to explain how S might come to hold the false belief that Q in a situation in which Q is false but which is in every other respect exactly like the actual situation. Given this, we can express DeRose's suggestion more clearly in the following way, as a condition that we will call *explanationist counterfactualism*, or (EC):

- (EC) S knows that P only if *either* S sensitively believes that P *or*, where P is implied by some Q (that is distinct from P), S knows that Q, and -P fails to explain how S might come to hold the false belief that Q.

We contend that when (EC)'s conditions are met, the not-P situation is not a threat to S's knowing that P, given, of course, that S is competent at deduction, and that S is in a position to know P by deduction. As we might put it, S can know that P because she can modus ponens her way to P from her sensitive belief that Q.⁴ Such knowledge that P is not threatened by not-P situations that fail to explain how she might come to hold the false belief that Q.

Our aim here is to show that (EC) is not threatened by objections like those leveled by Sosa, Kripke, and Williamson. In showing this, we will also challenge some attempts to show that counterfactualist accounts of knowledge are fated to be incompatible with Closure. In an effort to preserve Closure, we may not have to abandon the prospects for a counterfactual condition that involves sensitivity. We proceed by first providing a brief review of the reasons to find a sensitivity condition attractive. Then we turn to our main task – a discussion of the leading cases in which Nozick's ur-condition forces us absurdly to deny Closure. We show how the revised condition handles each case, and how it does so in a way that is both illuminating and non-ad hoc. We end with some reasons for thinking that (EC) has the additional virtue of helping to better clarify and to focus the debate over skepticism.

I WHY SENSITIVITY?

Nozick himself abandons the most basic form of sensitivity, according to which S knows that P only if S would not believe that P if P were false,

⁴ We are tentatively inclined to think that to know P, it is not enough that S *can* modus ponens her way to P from her sensitive belief that Q. This ability suffices only to put S in a position to know

and adopts a condition that makes reference to methods of belief formation: S knows that P only if

(SEN) S does not believe that P in the nearest non-actual worlds in which P is false and in which she uses the same method that she uses in the actual world in coming to form a belief as to whether P.⁵

There are considerations, both theoretical and case-based, that favor building an account of knowledge around a sensitivity condition like (SEN). In this section, we briefly review a few of these considerations.

First, the theoretical considerations. When we inquire into the nature of knowledge, it seems that S must appropriately respond to her environment if she is to know something about it, say, that P. If S is so to respond, she must come to hold the true belief that P. Yet this cannot be the end of the story. For, as we have learned from Gettier,⁶ if S's belief that P is luckily or coincidentally acquired, then it does not count as knowledge, not even if it is true and suitably formed. How, then, do we keep luckily acquired beliefs from counting as knowledge? We must demand more of S than that she respond appropriately to her environment by suitably forming the true belief that P. One idea is to demand that S respond appropriately in environments in which it is *not* the case that P. Nozick puts the point like this:

A belief that *p* is knowledge that *p* only if it somehow varies with the truth of *p* ... An account that says nothing about what is believed in any situation when *p* is false cannot give us any mode of varying with the fact.⁷

But since S's present environment is, *ex hypothesi*, an environment in which P is true, our additional demand must take the form of a contrary-to-fact condition: *if P had not been true, S would have responded differently to her environment and, in particular, would not have believed that P.*

As this suggests, it's true not only that these conditions – truth and belief – must be met if one is to know, but also that those conditions themselves be related. In particular, S's knowing that P calls for an intimate connection between whether P is true and whether S believes that P.

P. S will know P only when S *does* modus ponens her way to P from Q (in addition to her belief's meeting any other conditions on knowledge besides sensitivity).

⁵ See Nozick 1981, 179. Nozick's appeal to methods has been criticized by Luper-Foy 1987a, 3, and Williamson 2000, 152–56. We do not respond to these criticisms here, as it is not our goal in this paper to defend sensitivity against all charges.

⁶ See Gettier 1963.

⁷ Nozick 1981, 208.

Any adequate epistemological account should have something to say about the relationship between truth and belief, and a sensitivity account does so: it straightforwardly captures the idea that if one is to know, one's beliefs must vary with the truth.

Let's now consider a few cases. We'll see first that a sensitivity account handles cases that count against reliabilism. 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 (since his vision is reliable) and truly, that this is a barn – it seems that Henry doesn't *know* that it's a barn. A sensitivity account can explain why Henry fails to know: ff this were *not* a barn, Henry would nevertheless have believed that it *is* a barn. That is, Henry fails to know that this is a barn because he insensitively believes that it's a barn.⁸

We can strengthen the case for sensitivity by appealing to additional cases. In addition to handling Gettier's original cases against the justified-true-belief account of knowledge, a sensitivity account can handle a nice array of cases that sprang up in the wake of Gettier's paper. For example, in Brian Skyrms' Sure-Fire match case, which concerns inductive knowledge and knowledge of the future, a pyromaniac justifiedly believes that the next Sure-Fire match he strikes will ignite.⁹ However, the match is defective – it ignites, but only because it was hit by a coincidental burst of Q-radiation. The pyromaniac has a justified true belief that the match will ignite, but he does not *know* that it will. Again, sensitivity handles the case: if the match had not ignited – if there had been no coincidental burst of Q-radiation – the pyromaniac would nevertheless have believed that the next match he strikes will ignite. Again, since the pyromaniac's belief does not suitably vary with the truth, it fails to count as knowledge.

These cases, and others, suggest the rather modest point that we should explore the prospects for providing an account of knowledge that involves sensitivity. Further motivation comes from theoretical considerations which suggest that on an account of knowledge built around sensitivity, the only beliefs that count as knowledge are those that suitably vary with the truth.

⁸ In fact, Goldman himself added a sensitivity condition to his analysis of knowledge. He idiosyncratically referred to it as the local reliability condition (and referred to the more common reliabilist condition as the global reliability condition); see Goldman 1986, 44–47.

⁹ See Skyrms 1967.

II SOSA'S FALSE BELIEF CASE

In spite of its plausibility, however, there are reasonable objections to (SEN). This brings us to the first case we wish to consider. Ernest Sosa points out that whenever S knows that P, no matter what P is, the sensitivity condition prevents her from knowing that *S is not wrong in believing that P*.¹⁰ For, in the nearest counterfactual situations in which S is wrong in believing that P, she nonetheless believes that she is not wrong in believing that P. Consider, for instance, my belief that

(H) I have hands.

Skeptical worries aside, this belief amounts to knowledge, and it meets the condition expressed in (SEN). Vision, the method that I use in forming my belief that H, responds appropriately to my environment. In the relevant worlds in which I don't have hands, I was born without hands or I had an accident that cost me my hands. In those worlds, vision does not lead me to believe that I have hands. But now consider my belief that

(F) It is not the case that my belief that I have hands is false.

This belief also amounts to knowledge: we certainly don't think that my belief that I have hands is a false belief, and my belief that (F) is both justified and subject to no Gettier problems. Yet my belief that F does *not* meet the condition on knowledge expressed in (SEN). For, if I were in an environment in which it is true both that I believe that I have hands and that that belief is false, I would nevertheless believe that my belief that I have hands is *not* false – since I believe in this environment that I have hands, I will surely believe in addition that my belief that I have hands is *not* false.¹¹

We agree with Sosa: this result is simply not plausible. That is, it's simply implausible to suppose both that I know that H and that I fail in these circumstances to know that F. It seems quite clear in this case that I should be able to extend my knowledge from H to F. For I know that H, I know with relative ease that H entails F, and I competently perform the simple deductive inference from H to F. Still, despite my impeccable epistemic and deductive performances, defenders of (SEN) are committed to saying that I do *not* know that F. This means that (SEN) leads us, quite absurdly, to regard as false the operative instance of Closure: if I know

¹⁰ Sosa 2002. DeRose discusses this case in DeRose 1995.

¹¹ See DeRose 1995, 196–97.

that H, and if I know that H implies F, then I know that F. We need some account that does not force us absurdly to deny this claim.

(EC) is such an account. At the heart of this account is a requirement that involves *explanation*. This requirement, which we call (EXP), is expressed in the second disjunct of (EC)'s consequent. Articulating it more precisely, (EXP) has three components:

- (EXP) (i) Where P is implied by some Q (that is distinct from P),
 (ii) S knows that Q,¹² and
 (iii) \neg P fails to explain how S might come to hold the false belief that Q.¹³

The introduction of (EXP) allows us to hand down the proper verdict in Sosa's case. That is, it allows us to say that I know that F. For (i) F – *it is not the case that my belief that I have hands is false* – is implied by H – *I have hands*,¹⁴ (ii) I know that I have hands, and (iii) F 's simply being *false* – that is, its simply being the case that my belief that I have hands is false – fails to explain how I might come to hold this false belief. (i) and (ii) are obvious. As for (iii), to simply describe a situation as one in which my belief that I have hands is false is to so severely under-describe it that there is a failure to explain how I might come to hold this false belief. That is, the description in $\neg F$ is too sparse to explain how I might come to hold the false belief that H in a non-actual world which is much like the actual world, but differs from it only in H's being false. Typically explanations cite causes, but this description of the situation in which I do not have hands says absolutely nothing about how, or about via which methods, I might come to hold the false belief that I have hands.¹⁵

Further reflection on Sosa's case reveals that (EC) is not ad hoc. Note that so sparsely describing a situation as one in which my belief that I have hands is false but in which I nevertheless believe that I

¹² That is, S knows *in the actual world* that Q.

¹³ As we are about to elaborate, we take $\neg P$ in (EXP)'s clause (iii) to be a *description* of the closest not-P world.

¹⁴ This implication is true at least in contexts in which I believe that I have hands. We stipulate that this is such a context.

¹⁵ Moreover, in addition to the fact that raising the simple and bare possibility that my belief that I have hands is false does nothing either to keep me from knowing that I have hands or to keep me from knowing that H entails F, raising that possibility does nothing to impugn my ability to perform simple deductions. For these reasons, raising this possibility gives us no reason to think that I cannot come to know that F on the basis of a competent deduction from my knowledge that H.

don't hold the false belief that I have hands – call this description of that situation FALSE BELIEF ABOUT HANDS – does nothing to explain how I might have gone wrong in a situation much like my actual one when I formed a belief as to whether I have hands.¹⁶ The explanandum, notice, is not *that* I might have gone wrong, but rather *how* I might have gone wrong. The focus is the history of the belief, its genesis. What would it take to explain how I might have gone wrong in a situation much like my actual one? Since I might have gone wrong in the actual world if and only if there is a world that is close to the actual world in which I do go wrong, a description of a not-P world explains how I might have gone wrong in the actual situation just in case it explains how I do go wrong in a close possible world. FALSE BELIEF ABOUT HANDS, however, gives us no reason to think there is a close world in which I'm wrong in thinking that I have hands, and so it fails to provide the needed explanation. Furthermore, we should stick with the assessment that I know in the actual situation that I have hands *precisely because* the FALSE BELIEF ABOUT HANDS description does nothing to explain how I might come to hold the false belief that I have hands. Since this failure of explanation allows us to continue to judge that I know that I have hands, we have good reason to include an explanation requirement like (EXP) in an account of knowledge.¹⁷ That is, the requirement in (EXP) is not ad hoc.

One of our assumptions about how to understand (EXP) has now become evident. We assume that explanation is, in the first place, a relation between *descriptions* of situations (or events). Rather than pedantically highlighting this in the rest of the chapter, we will simply offer descriptions of situations with the understanding that it is these descriptions which are key. A second assumption, one that also falls out of a general view about explanation, is that conversational context plays a key role in filling in background facts that are relevant to assessing explanations. This will be especially apparent when we get to cases where a good explanation is in place and thus (iii) in (EXP) is not satisfied. The next case, we argue, is like this.

¹⁶ Our emphasis is on the need to explain how the process that I use in forming a belief as to whether I have hands might lead me to get things wrong, rather than right. We leave it open that other things might need to be explained as well if the condition is to be met, for example, why I formed a belief as to whether I have hands rather than a belief as to whether I have feet.

¹⁷ In Black and Murphy 2007, we argue that (SEN) and (EXP) are related, and that (EXP) has not merely been cobbled onto (SEN) in order to take care of certain troublesome cases.

III KRIPKE'S RED BARN CASE

It is reported that Saul Kripke offers a case like the following:¹⁸ Henry, still out on his drive, believes (truly) that there is a red barn before him. In the nearest counterfactual situation in which there is not a red barn before him, there is a blue barn before him, and in that situation Henry would not believe that there is a red barn before him. He therefore meets the sensitivity condition with respect to the proposition that *there is a red barn before him*. Trouble arises, however, because in the nearest counterfactual situation in which there is no barn at all before him, there is a barn façade there, and Henry mistakenly believes that there is a barn before him. This means that, according to sensitivity, Henry does not count as knowing that there is a barn before him. Once again we get two results that are strange bedfellows: Henry knows that there is a red barn before him, but he fails to know that there is a barn before him. This strikes us as odd because it is quite clear that Henry can easily infer, via conjunction elimination, from his knowledge that there is a red barn before him to the proposition that there is a barn before him.

We agree that we must avoid the absurd denial of Closure to which Kripke calls our attention. We first focus on determining whether Henry fails to know that there is a barn before him. Does the scenario involving the barn façade explain how Henry might go wrong in believing that there is a red barn before him? Clearly, it does, since the façade scenario explains how Henry might come mistakenly to believe that there is a *barn* before him. It follows, then, that Henry does not meet (EC)'s conditions with respect to the proposition that there is a barn before him.¹⁹ This brings out a strengthening dimension of our proposal – a dimension that *prevents* S from knowing *any* of the propositions involved in certain instances of Closure.

We need to reconsider whether Henry meets (EC)'s conditions with respect to the proposition that there is a red barn before him. Since this proposition is equivalent to the conjunction that there is a barn before him and there is something red before him,²⁰ we propose that we apply (EC) separately to each conjunct. When we do this, we see that Henry

¹⁸ The report is given in, among other places, Shatz 1987, 265n.18. See Kripke 2011.

¹⁹ We are unable to think of a candidate proposition that entails that there is a barn in front of Henry, is a proposition Henry sensitively believes, and meets (iii) of (EXP). If this makes our judgment that Henry does not know that there is a barn in front of him inductively based, we are confident that it is nonetheless well supported.

²⁰ He might also need a third conjunct: *and the red thing before him is a barn*.

fails to know that there is a red barn before him, for he fails to meet both conditions in (EC) with respect to one of the conjuncts, namely, that there is a barn before him. He fails to meet the sensitivity condition, since in the façade scenario, in which there is no barn before him, Henry nevertheless believes that there is a barn before him. He also fails to meet (EXP). This is because for any proposition that might be a candidate premise from which Henry could deduce that there is a barn before him, the façade scenario will explain how Henry might come to hold a false belief in that premise. For this reason, (iii) of (EXP) will not be met. Specifically, this is so because for whatever aspect of that premise that entails that there is a barn before Henry – different candidate premises will involve different aspects which entail this – the façade scenario will successfully explain how Henry might go wrong with respect to that aspect of the candidate premise. We leave it the reader to check that this is so by considering various candidate premises: for example, there is a red barn before Henry, there is a wooden barn before Henry, there is a beautiful barn before Henry, and so forth. Thus, according to (EC), Henry fails to know that there is a red barn before him. This plus the additional claim that knowing a conjunction requires knowing each of its conjuncts means that we are not committed to Kripke's absurd denial of Closure.

Part of our proposal is that we should apply (EC) separately to each individual conjunct of a believed conjunction. So, to know a conjunction, one must know each of its conjuncts, and therefore one must meet either the sensitivity condition or (EXP), the explanationist condition we're highlighting, with respect to each conjunct. Proponents of sensitivity need to require this if they are going to combat a general recipe for counterexamples that is otherwise available to their opponents. The general recipe begins with a sensitive belief, say Sally's sensitive belief that

(H) Sally has hands.

It then conjoins H with X, where X (i) is any true proposition that the subject believes and (ii) is such that the nearest $\neg X$ world is farther from the actual world than is the nearest $\neg H$ world. In any such case, Sally will meet the sensitivity condition with respect to the conjunction *H and X*. For the nearest counterfactual situation in which *H and X* is false is identical to the nearest counterfactual situation in which H is false, and Sally doesn't believe that H in that situation. A sensitivity account seems to suggest, then, that *for any such X*, Sally knows that *H and X*. But of course there is nothing here about whether Sally is epistemically competent with respect to X. Meeting the unadorned sensitivity condition

with respect to *H and X* is simply not sufficient for knowing that *H and X*, for even if one meets that condition, one needs to be epistemically competent with respect to *X* (as well as with respect to *H*). If one *fails* to meet the sensitivity condition with respect to *X*, one fails to know that *H and X*. No theory should allow Sally to know that *H and X* simply by meeting the sensitivity condition with respect to that conjunctive proposition, since one can do so by only meeting the sensitivity condition with respect to *H*. Thus, an adequate sensitivity condition on knowledge will maintain that I know that a conjunction is true only if I meet the prescribed epistemic conditions with respect to each of the conjunction's conjuncts.

IV WILLIAMSON'S DACHSHUND CASE

In discussing a condition like (EC), Timothy Williamson calls on Alvin Goldman's dachshund case.²¹ In this case, Oscar sees a dachshund and believes that there is a dachshund before him. Moreover, in the nearest counterfactual situation in which there is *not* a dachshund before him, there is instead another kind of dog before him, one that he does not mistake for a dachshund. Oscar therefore meets the counterfactual condition with respect to his belief that there is a dachshund before him. Knowing, then, that all dachshunds are dogs, he competently infers that there is a dog before him. And it is quite reasonable to suppose that, in this way, Oscar comes to *know* that there is a dog before him. Nevertheless, he fails to meet the sensitivity condition with respect to his belief that there is a dog before him. For suppose that in the nearest counterfactual situation in which there is no dog before him, there is a *wolf* before him, which he mistakes for a dog. In this case, sensitivity allows him to know that there is a dachshund before him and competently to deduce from this that there is a dog before him, but it does *not* allow him to know that there is a dog before him. Again, sensitivity theorists are forced into an absurd denial of the following obvious instance of Closure: if Oscar knows that there is a dachshund before him and if he knows that *there is a dachshund before him* implies *there is a dog before him*, then he knows that there is a dog before him. Proponents of a rudimentary sensitivity theory seem forced to say that Oscar knows that there is a dachshund before him, but that he fails to know something that he easily and competently deduces from this, namely, that there is a dog before him.

²¹ See Williamson 2000, 153, 159. The case originally appears in Goldman 1976.

(EC) handles this case quite well. To see how, return to the relevant closure claim: if *S* knows both that there is a dachshund before her and that all dachshunds are dogs, then she knows that there is a dog before her. (EC) has us consider a counterfactual situation in which there is a wolf before Oscar, one that he mistakes for a dog. Call this WOLF. Now, Oscar has a tendency to make a particular sort of mistake in WOLF – in such situations he believes that there is a dog before him. But the dogs for which he mistakes the wolf are dogs of certain sorts, and we have no reason to believe that *dachshunds* are among the dogs for which he mistakes the wolf. For, while some dogs are quite like wolves, dachshunds are not. So even though Oscar has trouble distinguishing *wolves* from (certain kinds of) dogs, he might nevertheless be very good – perhaps even infallible – at recognizing *dachshunds* both as dachshunds and as things that are pretty obviously dogs. So Oscar's failure in WOLF, where he mistakes a wolf for a dog, does nothing to explain how he might come mistakenly to believe that there is a dachshund before him (or how he might go wrong in deductively reasoning from the fact that there is a dachshund before him to the fact that there is a dog before him). Thus, Oscar satisfies (EC) with respect to the belief that there is a dog before him. For (i) *there is a dog before him* is implied by *there is a dachshund before him*, (ii) Oscar knows that there is a dachshund before him, and (iii) WOLF *fails* to explain how Oscar might come to hold the false belief that there is a dachshund before him. According to (EC), Oscar can know in this case both that there is a dachshund before him and that there is a dog before him.

We learn an important lesson when we examine (EC)'s response to the dachshund case. WOLF fails to explain how Oscar might come to hold the false belief that there is a dachshund before him. But what is responsible for WOLF's explanatory failure? Here is one line of reasoning that suggests an answer to this question and from which we can draw a more general conclusion about which situations fail to provide the explanation called for in (EC)'s explanation condition. (EC)'s response to the dachshund case brings out the fact that a difference in *percepts* can contribute to explanatory difficulties. When we talk here of a difference in percepts, we have in mind Alvin Goldman's notion of perceptual equivalence, according to which *qualitatively identical percepts* are percepts that are *exactly* similar, while *perceptual equivalents* are percepts that are at least *very* similar.²² A situation in which there is a dachshund before Oscar – call

²² See Goldman 1976.

this DACH – is one in which Oscar has percepts that are dramatically and importantly different from those he has in WOLF. Moreover, the percepts he enjoys in DACH would not lead him to make the mistake that he makes in WOLF. Because of the character of his percepts in DACH, Oscar would not believe that there is a dog before him. This suggests that a difference in percepts from one situation, α , to another nearby situation, w , can keep w from providing the explanation called for in (EC)'s explanation condition. This suggestion is captured in this principle:

- (PE) Some situation, w , explains how S might come to hold the false perceptual belief that Q only if w is a situation that is perceptually equivalent to α .³³

In the case at hand, w is WOLF and α is DACH. According to (PE), then, since WOLF is perceptually distinct from DACH, it fails to explain how Oscar might come to hold the false belief in DACH that there is a dachshund before him. We may therefore maintain that Oscar knows in DACH both that there's a dachshund before him and, via deductive inference, that there's a dog before him. In this way, (EC) takes care of the dachshund case.

We have now seen, through an examination of several prominent and challenging cases, that (EC) – and hence an epistemological account that involves sensitivity – need not absurdly deny Closure. (EC) therefore stands up against the charge that Nozick's sensitivity condition implies absurd denials of instances of Closure. That charge should not deter us from adopting our more developed counterfactualist account.

V IMPLICATIONS FOR THE SKEPTICISM DEBATE

Besides handling the cases put forth by Sosa, Kripke, and Williamson, our proposal has another virtue: it clarifies and helps bring into focus the recent literature on skepticism. This literature has concentrated on skeptical arguments that employ Closure, arguments like this:

- (P₁) If I know that I have hands and I know that my having hands entails my not being a brain-in-a-vat, then I know that I'm not a brain-in-a-vat.
 (P₂) I don't know that I'm not a brain-in-a-vat.
 (P₃) I know that my having hands entails my not being a brain-in-a-vat.
 (C) Thus, I don't know that I have hands.

³³ As with (EXP)'s (iii), it is really a particular *description* of w that is in play in (PE).

Many prominent epistemologists²⁴ have sought to avoid the skeptical conclusion of this argument by denying Closure, an instance of which appears in the above argument as (P1). Typically, they support their denials of Closure by appealing to a counterfactualist account of knowledge.

However, for all we have seen so far, (EC), a far more plausible counterfactualist account, may be compatible with instances of Closure like (P1). One is therefore advised not to first adopt some plausible counterfactualism and then just simply claim that one's favored counterfactual condition demands that we deny (P1). In this section, we use the lens of (EC) to explore two other options: denying (P2), and accepting the argument as it is.²⁵ Both the former, Moorean response, and the latter, skeptical response leave the instance of Closure at (P1) intact. Both of these responses agree that it is true *either* that I know both that I have hands and that I'm not a brain-in-a-vat, as Moore alleged, *or*, as the skeptic alleges, that I know neither of those things. Does (EC) shed any light on the debate between the Moorean and the skeptic?

Whether it does depends on whether the counterfactual situation in which I am a brain-in-a-vat can explain my making a mistake about whether I have hands. Both the Moorean and the skeptic can appeal to compelling considerations here – certain considerations demand that we take seriously the legitimate threat of skepticism, but other considerations suggest a promising strategy for responding to this threat in such a way as to maintain that I *do* know that I'm not a brain-in-a-vat. We recommend that these considerations be the future focus of the skeptical debate.

In brief, here are the considerations. Begin with this question: does the counterfactual situation in which I am a brain-in-a-vat explain how I might make a mistake about whether I have hands? The skeptic claims that it *does*. My being a brain-in-a-vat, the skeptic might argue, will surely explain how I would form the false belief that I have hands. After all, brains-in-vats are handless creatures who are electrochemically stimulated in a way that generates perceptual experiences that are exactly similar to

²⁴ Including, among others, Dretske 1970; Goldman 1986; McGinn 1984; and Nozick 1981.

²⁵ There is a third option as well: We can adopt a version of epistemological contextualism. We feel that (EC) is compatible both with contextualism and with its denial. For one thing, some contextualists are motivated, at least in part, by the desire to retain both closure and counterfactualism (see Cohen 1988, 1999, and DeRose 1995). Yet (EC) allows us to retain both closure and counterfactualism *without adopting contextualism*. Still, the main elements of (EC) – sensitivity and explanation – are amenable to a contextualist treatment. It has been argued that the notion of explanation is context-sensitive (see Lipton 1990), and as DeRose's work shows, contextualist accounts can successfully incorporate the notion of sensitivity. This suggests that (EC), too, would be amenable to a contextualist treatment. We regret that we don't have the space here to pursue this issue further.

those that we are now having in what we take to be normal circumstances. Putting this in terms of perceptual equivalence, the skeptic might maintain that my percepts in w , the situation in which I am a brain-in-a-vat, and my percepts here in α are not simply *perceptual equivalents*, but that they are in fact *qualitatively identical percepts*. This fact gives the skeptic the freedom, at least to the extent to which I satisfy the necessary condition expressed in (PE), to maintain that w explains how I would come to hold the false belief in α that Q (where Q can be any – or almost any – positive claim about the external world). Given this, since I am now having perceptual experiences as of my having hands, it is reasonable to believe, the skeptic will argue, that my being a brain-in-a-vat explains perfectly adequately how I would come mistakenly to believe that I have hands.

Things are not as clear as this, however. For the Moorean can argue that a satisfactory explanation must appeal to belief-forming methods that we actually use. That is, the Moorean might argue that the belief-forming methods that are actually available to cognizers like us are not available to brains-in-vats. Vision, for example, is a belief-forming method on which we rely, but no such method is employed by brains-in-vats. The Moorean might conclude, therefore, that my being a brain-in-a-vat *fails* to explain how I might come mistakenly to believe, on the basis of methods that I *actually* use, that I have hands.²⁶ This would clear the way for us to say that I know both that I have hands and that I'm not a brain-in-a-vat.

The point of tracing this debate is not to make a claim about which of its participants – the Moorean or the skeptic – is right. Our point is rather to show that (EC) pays further dividends: it can help us focus on correctly specifying the explanatory relation that helps to determine whether one knows both of the propositions that are involved in a given instance of Closure, or whether one fails to know either of those propositions. This in turn might help us eventually to settle the score between skepticism and Mooreanism. Of course, we have limited ourselves to just a few remarks about whether considerations about belief-forming methods or about perceptual equivalence play a role – perhaps even a crucial role – in providing an appropriate and adequate response to the above skeptical argument. One would need to do more in order to determine whether (EC)'s explanatory condition includes demands that concern belief-forming methods and perceptual equivalence, and it might very well be that we need to identify and examine other demands of that

²⁶ For an argument of this sort, see Black 2002.

condition. But once one does so, one would then be in a good position to render a verdict in favor of either skepticism or Mooreanism. If we meet the demands of (EC)'s explanatory condition, we should deny (P₂), and Mooreanism prevails; but if we fail to meet those demands, we should accept the skeptical argument as it is, and skepticism prevails. Here (EC) pays real dividends in helping us focus our attention on key elements in the skeptical debate, elements that might help us finally to put an end to that debate.