

Skeptical Success

Troy Cross

The Project

You intend to write a textbook on the theory of knowledge. Chapter One is slated to cover skepticism. You need some colorful skeptical scenarios, sure to vex your audience. Later, your anti-skeptical arguments will ease their pain, but for now, the goal is to inflict it.

"You think you know you have hands," you begin, "but maybe, *in fact*, you don't have any hands!"

You stop. Isn't that a skeptical scenario? It entails that you don't know something that you now take yourself to know, something you take yourself *obviously* to know. Yet somehow your hypothesis doesn't seem to live up to Descartes' dreaming argument (13) or his evil genius idea (15), Russell's five-minute hypothesis (2008, 104), Goodman's grue possibility (72), Goldman's phony barn country (1976, 772-773), or even the Wachowski brothers' Matrix movies (1999).¹

You set yourself once more to the task. "You think you know you have hands," you begin, "but maybe you don't even *believe* you have hands! Belief is just as much a requirement for knowledge as truth, so if you don't even believe you have hands, you can't know you do."

This feels even worse. Your dream of Cartesian fame is fading rapidly.

One more try. "You think you know you have hands," you write, "and maybe you do have hands, but as it happens, you are correct just as a matter of luck! Lucky true belief is not knowledge, so you don't know you have hands."

A bit more subtle, a slight improvement over the others. Still, your scenario doesn't begin to approach the classics.

"What's missing?" you ask yourself. "It can't be that Descartes' scenarios are antecedently judged to be more *likely* than mine. Quite the opposite. I know probability theory. If I am a handless victim of an evil genius, then I am still handless. That's guaranteed. But I may be handless for reasons other than being the genius's victim. I may be handless and in the matrix. I may be handless because of an accident on my way to work and still in shock, in denial about my loss. No evil genius in sight! Likewise, if I am dreaming, then my belief, if true, is true by virtue of luck. But dreaming a truth is not the only way to be luckily correct. Maybe I am looking at the equivalent of a stopped clock that just happens to read the correct time (Russell, 1994, 113). Maybe I am in phony *whatever* country. In sum, if Descartes' scenarios obtain, so do mine, but the reverse does not hold. Heads, we both win; tails, I win and Descartes loses. If someone's

¹The phony barn scenario, though popularized by Alvin Goldman, is standardly attributed to Carl Ginet. In the paper I'll make use of the brain-in-a-vat scenario. Its origin is unknown, though like Keith Lehrer's (1971) 'Googol' example, it was probably a way of making the evil genius case compatible with materialism.

scenarios are to be judged more likely, it should be *mine* that are so judged."

In spite of this airtight demonstration that your scenarios necessarily equal and possibly even *best* Descartes', you have to admit they exert a negligible force on the psyche.

Ah! Perhaps it is the generality, the reach, of the classic scenarios that explains their intuitive appeal. Some skeptical scenarios seem parochial enough -- mules cleverly disguised as zebras (Dretske, 1971, 1016), white walls with red lights shining on them (Dretske, 1971, 1015) -- but the really big ones are always sweeping, foundation-shaking. If an evil genius is in control, then one is ignorant about virtually everything that one takes oneself to know. So substitute "almost all of your beliefs" for "you have hands" in the above. Is it improved? Not much.

What's missing? What makes some skeptical scenarios successful and others unsuccessful? That is the question I intend to ask. And by "success" I do not mean to prejudge the question of whether successful skeptical scenarios ultimately rob us of knowledge. By "success" I mean only that the scenarios work their magic on the intended audience, giving rise to an intuition that is palpable and extremely difficult to shed--an intuition that one does not know, or even cannot know, that the skeptical scenario presented does not obtain. Skeptic and Moorean, contextualist and invariantist, internalist and externalist, all owe an account of what distinguishes successful from unsuccessful scenarios. All owe an account of why some, but not all, skeptical scenarios exert that peculiar, disconcerting force whose proper management is the subject matter of so much professional epistemology.

As we have seen, the imputed lack of true belief, or the lack of true justified belief, or of true not-merely-lucky belief, or for that matter, just the lack of knowledge alone does not a successful skeptical scenario make. The traditional question in epistemology, Plato's question in the *Theaetetus*, is: what is knowledge? What, when added to true belief, yields knowledge, distinguishes knowing from veridical opining? The starting point of the present paper is, rather, the question of what must be added to (or subtracted from, as it were) the lack of knowledge to yield skeptical success. What distinguishes skeptical success from mere cases of imputed ignorance? Answering this question will, I think, shed a surprising amount of light on the traditional question.

Successful skeptical scenarios, I shall argue, are marked by their explanatory prowess. A successful skeptical scenario not only *entails* that you don't know some things you take yourself to know, but also *explains* your taking yourself to know those things even though you don't. It does seem, on its face, that what is lacking in the failed attempts above is a back story about how you're supposed to have gone wrong. How is it that you in fact know less than you think you do? We're not told.

What of the "airtight" argument that the simple skeptical scenarios envisioned must be judged at least as likely as Descartes'? The argument is indeed sound, but what it shows is only that skeptical success is not a matter of probability. A less probable scenario can be more successful at inducing skeptical unease than a more probable scenario. If skeptical success is explanatory success, this should not be a surprise. Explanatory adequacy and probability are often at odds.

How many philosophers does it take to diagnose a broken light bulb? The lights go out unexpectedly. Many hypotheses come to mind about why you are sitting in the dark. The light

bulb broke. The power lines are down. A rat chewed through a wire in the walls. Someone flipped a switch in another room. More probable than all of these is that electrons are no longer coursing through the filament in the bulb. It is, after all, true on all of the other hypotheses entertained. Suppose you believe that the bulb is broken, whereas I believe only that that electrons are no longer flowing through the filament. Then if you're right, I'm right, but if you're wrong I might still be right. Nevertheless, your belief is, in some sense, the better one. While more likely to be false, it is also a better story, one we can actually use to restore illumination. He who believes only the more probable remains in the dark, resting comfortably in the assurance that he is less likely to be mistaken.

Our doxastic engine is tuned so that we are strongly drawn to explanatory hypotheses. We favor them over non-explanatory hypotheses, somehow even when they are *less* probable. The most dramatic illustration of this strange attraction is what is known as the "conjunction fallacy".

Tversky and Kahneman (1974) posed the following question to test subjects:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. Which is more likely?

1. Linda is a bank teller.
2. Linda is a bank teller and is active in the feminist movement.

A vast majority of respondents say 2 is more likely, contrary to probability theory.

Are we really *that* irrational? The conclusion is hard to resist. The experiment has been replicated, while shifting many of the variables involved. For instance, you might think subjects take 1 to imply that Linda is not a feminist, but the question has been posed substituting "Linda is a bank teller, whether or not she is a feminist" for 1, with the same result (Kahneman & Tversky, 1983). How troubling!

Why the attraction to 2 over 1? Because 2 *explains* the data and 1 does not. Obviously. And clearly, we place a high epistemic value on believing good explanations. The reason for this, I think, is the two-fold goal of belief. As William James pointed out, our goal in belief formation is not simply to *avoid error*, but also to *believe truths* (2003). The two goals are, to some extent, opposed to one another. Other things being equal, a miserly believer will avoid error, but miss out on important truths. A profligate believer will grasp many truths, but also fall more frequently into error. Proper belief formation charts a middle course between miserliness and profligacy, attempting to gain the most true belief available while incurring the minimal risk of error. Good explanations are potentially important truths, ones that might be fruitful as guides to other truths and to the avoidance of errors; they are thus worth some risk of going wrong.

My own, admittedly unscientific, diagnosis of Tversky and Kahneman's results is that subjects are confusing the general question of "What is better to believe on the basis of this data?" with "Regardless of which of these hypotheses is better to believe given the data, which is less likely to be false?" Perhaps even stated that way, subjects would continue to confuse the questions. Separating out these two virtues of a hypothesis is something philosophers routinely do. But for ordinary folk, there is just, I think, an overriding, overall judgment of the epistemic "goodness" of the hypothesis, which is a *combined score* on the Jamesian "avoiding error" and "gaining truth"

tests.

Even for philosophers, the following question is a bit annoying. Given your experience, which of the following is more probable:

1*. I have hands.

2*. It seems to me that I have hands.

Only reluctantly do we admit that 2* is more probable, and only because we explicitly think through the possibilities. It seems that in admitting 2* is more probable, we are giving ground to the skeptic; we are acknowledging that his belief is in some important way *better* than ours. And we are loathe to admit that.

Indeed, one kind of skeptic simply plays on probability theory, urging us to believe that which is more probable, that which is less likely to be false. That sort of skeptic exploits our desire to avoid error. Let's call such a skeptic a "probability skeptic". The probability skeptic is easily banished. We simply note that avoiding error is not all that we are trying to accomplish in forming beliefs. We drive him away with aphorisms. "Nothing ventured, nothing gained," we say. "Better to have loved and lost..."

What is interesting, I think, is that this is not skepticism as we know it in epistemology. Rather, successful skeptical scenarios are designed to exploit our attraction to explanatory hypotheses, our desire to form beliefs in accordance with the *gain true belief* side of the Jamesian equation. This attraction to explanatory hypotheses is precisely what allows us to ignore the probability skeptic. Our anti-skeptical impulse is thus harnessed and pressed into the service of skepticism.

Just as, in an anti-skeptical embrace of truth, we latch onto non-skeptical hypotheses that are explanatory but improbable, we also, in the same spirit, cannot bring ourselves to brush aside explanatory skeptical hypotheses, while dismissing certain logically weaker things. To state the observation once again, the advantage of the hypothesis that I have hands over the hypothesis that it seems to me that I have hands is precisely the advantage of the hypothesis that I am a handless victim of the evil genius over the hypothesis that I lack hands. Yet, in the first case, that advantage works in favor of my knowing I have hands, whereas, in the second case, it works in favor of my *not* knowing I have hands.

That is the irony of skeptical success.

From Sensitivity to Explanation

Robert Nozick's so-called "sensitivity" condition requires that to know *p*, it must be the case that if *p* were false, you would not believe it, i.e., that belief be *sensitive* to the truth of what is believed (176). Sensitivity was offered as a necessary condition for knowledge, and typically, as part of a sufficient condition as well. I prefer to examine insensitivity as a requirement of skeptical success. When we try to explain what is missing in the scenarios above, and what is present in the classic skeptical scenarios, it is quite natural to say that even if you were in one of the classical scenarios, you would still believe you weren't, whereas, if you didn't, in fact, have any hands, you would *not* still believe that you did. You would believe something like: "Ah! I don't have any hands, just bloody stumps where hands used to be!" Though knowledge is absent in all of the

cases discussed, insensitivity is something all of the classical scenarios display, and something all of my simple, failed attempts at skeptical scenarios lack. Whatever its prospects as part of a theory of knowledge, perhaps it is the missing ingredient that, when added to unknown belief, yields skeptical success.

Sensitivity has died the death of all the other 1980's epistemic theories: it is buried under multitudes upon multitudes of clever, clever counterexamples. Sensitivity thus lies next to the causal theory of knowing (Goldman, 1967), the nomic account of knowing (Armstrong), the "no undefeated defeaters" theory (Klein, 1971; Lehrer & Paxson, 1969), and all of the responses to Edmund Gettier's famous puzzle (1963).

But sensitivity is the dead horse that epistemologists cannot stop beating. Nozick himself, one of sensitivity's original proponents, kicked things off, as we shall see, by pre-emptively raising some counterexamples. According to epistemology lore, soon thereafter Saul Kripke refuted the view. Alvin Goldman (1983) and Jonathan Vogel piled on (1987). And as if that weren't enough, just recently, philosophical heavyweight Timothy Williamson subjected sensitivity to one of its most impressive thrashings yet (2002). Still the anti-sensitivity literature grows. Perhaps, like new proofs of the Pythagorean theorem, epistemologists find this worthy of their time not because they want to establish that sensitivity fails, but to do so in an original way.

I am going to jump into the fray, working my way through the counterexamples and various fixes. I will do so for two reasons. First, the counterexamples to sensitivity as a necessary condition for knowing are also counterexamples to insensitivity as the missing ingredient for skeptical success. Second, my suggestion that skeptical success is explanatory success, though motivated by a direct intuition and the Jamesian reasoning above, dodges every arrow slung against sensitivity--which is an impressive display of arrow dodging.

As I said, sensitivity, as originally presented, has been pretty much abandoned. The field has moved on to related notions like safety (Sosa) -- sensitivity's contrapositive -- or sensitivity with closure (Roush), or has simply given up on modal epistemology (Vogel, 2007). But I think there is something deeply intuitive, deeply telling about the sensitivity requirement. It is a natural answer, a good answer to the question of why I hesitate to say I know I am not dreaming. I hesitate because in my dreams, I do not think I am dreaming. And it is an equally good answer to the question of what is wrong with the simplistic skeptical hypotheses with which I began this paper.

To say sensitivity is a good answer to these questions, however, is not to say it is the best one. I aim to defend the idea 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.² I will end up favoring the explanatory account I have gestured

²I said the post-Gettier theories lie buried under a heap of clever, clever counterexamples. That is not quite right. Some of the theories collapsed under the weight of their own caveats and qualifiers, formed in response to the same clever, clever counterexamples. William Lycan's exemplar here is Marshall Swain, whose impenetrable developments of the defeasibility and causal accounts may or may not be free of counter-example, for all anyone knows (Lycan, 149). A small sampling follows:

S has nonbasic knowledge that p iff (i) p is true; (ii) S believes that p; (iii) S's justification renders p evident for S;... (iv*) [w]here 'e' designates the portion of S's total evidence E that is immediately relevant to the justification of p, either (A) there is a nondefective causal chain from P to BSe; or (B) there is some event or state of affairs Q such that (i) there is a

towards above, but along the way discuss other more conservative principles that some may find more attractive than my final proposal.

In keeping with the literature, I'll discuss the counterexamples as cases of known but insensitive belief, and I'll talk about sensitivity as a condition on knowing. But bear in mind, these counterexamples of known but insensitive belief can be transformed into cases of unsuccessful skeptical scenarios. If p is intuitively judged to be known yet insensitive, then p is also insensitive and yet a skeptical failure.

Once through the gauntlet of cases, I will show that the explanatory account of skeptical success not only survives, but thrives. It explains, for instance, why we say that we know in induction cases but not lottery cases. And finally, I will briefly draw attention to the consequences of the explanatory theory of skeptical success for theories of knowledge, or for the semantics of "knowledge".

Counterexamples and Fixes

Nozick himself gave us the following case of known but insensitive belief (179):

Grandmother. Grandmother sees that grandson is well; If grandson had not been well, the family would have hidden him away and she would still have believed he was well. Yet, she knows he is well: she just saw him standing there and in perfect health!

Nozick solves the Grandmother case by modifying the rule of sensitivity to appeal to methods of belief formation (179). Grandmother formed her belief by visually inspecting her grandson. If she had visually inspected him and he was not well, then, using that same method, she would not have believed he was well. Sensitivity restored.

But Nozick's fix was only temporary. Goldman soon raised the following case (1983, 84):

Dachshund. You're looking at a Dachshund and form the belief: "There's a dog." But if there hadn't been a dog there, there would have been a hyena instead, which you would have taken to be a dog. Looking at the little Dachshund, you know there's a dog in front of you, but would still have believed this even if it were false and you were looking at (visually inspecting) the fearsome hyena.

Nozick's "methods" can be applied to resolve the Dachshund case, but problems ensue. Suppose

nondefective causal chain from Q to BSe ; and (ii) there is a nondefective causal chain from Q to P ; or (C) there is some event or state of affairs H such that (i) there is a nondefective causal chain from H to BSe ; and (ii) H is a nondefective pseudo-overdeterminant of P . [Where a causal chain $X \rightarrow Y$ is 'defective' with respect to S 's justification for p based on evidence e iff: Either (I) (a) there is some event or state of affairs U in $X \rightarrow Y$ such that S would be justified in believing that U did not occur and (b) it is essential to S 's justifiably believing that p on the basis of the evidence e that S would be justified in believing that U did not occur; or (II) there is some significant alternative C^* to $X \rightarrow Y$ with respect to S justifiably believing that p on the basis of e . [Where C^* is a 'significant alternative' to $X \rightarrow Y$ with respect to S justifiably believing that p on the basis of e if (a) it is objectively likely that C^* should have occurred rather than $X \rightarrow Y$; and (b) if C^* had occurred instead of $X \rightarrow Y$, then there would have been an event or state of affairs U in C^* such that S would not be justified in believing that p if S were justified in believing that U occurred.] (Swain 1972: 292; 1978: 110-11, 115-16).

that the method is typed narrowly not as mere visual inspection, but as seeing-a-Dachshund-sort-of-thing-and-taking-it-to-be-a-dog. Now you can say you know there's a dog, because if you used that method, you wouldn't form a false belief.

The problem with such narrow individuation of methods is that it is difficult to see how you use the same narrowly individuated method in the Grandmother case when looking at grandson well and when looking at him ill. One thing you know about grandson is that he isn't covered with unsightly tumors. But if he were covered with unsightly tumors he would look very different and you wouldn't be able to employ the same narrowly-defined method to form beliefs about him in both "tumored" and "un-tumored" states.

Nozick must avoid overly narrow methods because a perfectly narrow method virtually guarantees perfect insensitivity. If methods differed with any internal difference, then the only way to use the same method would be to have exactly the same internally characterized experience. And this would practically guarantee that while using that method the subject would reach the same belief he in fact reaches regardless of its truth.

Matters only got worse for sensitivity. Jonathan Vogel noticed that a pretty central kind of knowledge doesn't seem to be sensitive: inductive knowledge (1987).

Ice. You leave some ice in the backyard out in the sun. An hour later, you come to believe the ice has melted. If it didn't melt, if the heat by chance did not find its way into the ice, you would still believe it had melted. But you know the ice melted.

Trash chute (Sosa, 1999). You drop your trash down the garbage chute in your apartment high rise, coming to believe soon thereafter that it hit the bottom. If it had not hit the bottom, but perhaps caught on a snag, you would still have believed it did. But you know your trash hit the bottom.

Heartbreaker (Vogel, 2007). Sixty golfers are playing in a tournament tomorrow. The hardest par three on the course is called "The Heartbreaker". You know not all sixty golfers will shoot a hole-in-one on that hole. But if they were going to, you'd still believe they weren't.

And then, if getting induction wrong wasn't bad enough, there was something even more embarrassing.

Not-false-belief. Vogel (1987) observed that while sensitivity allows lots of ordinary knowledge, for any piece of such knowledge--say, that I have hands--suppose I form the belief that I do not falsely believe it, e.g., I do not falsely believe I have hands: $\neg(Bp \ \& \ \neg p)$. If I falsely believed I had hands, I would still think I didn't falsely believe that I had hands -- so Moore's paradox guarantees. $(Bp \ \& \ \neg p) \rightarrow \neg(Bp \ \& \ \neg p)$. And yet it seems as easy to know that I don't falsely believe I have hands as to know that I have hands. Call this case the not-false-belief case (hereafter NFB).

Let's set aside the induction cases for now and focus on the NFB case. DeRose notices that this case differs from classic skeptical cases in an important respect (1995, 23). That I falsely believe I have hands implies that I do not know something I take myself to know-- namely, that I have

hands--but it does not *explain* how I came to believe falsely that I have hands. Classic skeptical scenarios also imply that I do not know things I take myself to know, such as the proposition that I have hands, but they provide an explanation for how I came to believe those things falsely: for example, by saying that the evil genius misleads me.

Suppose we say, based on DeRose's suggestion, that the insensitivity of p inclines us against ascribing knowledge that p only in those cases where if not- p implies that we do not know things we ordinarily take ourselves to know, then not- p also explains how we came to falsely believe p . Now the insensitivity of NFB does not prevent us from saying we know it. Let q be the proposition that I do not falsely believe that I have hands. q is insensitive, but not- q is also one of those funny propositions that implies that I do not know something I ordinarily take myself to know and yet does not explain how I came to believe it. So q 's insensitivity does not in this case dispose us to deny that we have knowledge of it. q 's non-explanatoriness is an antidote to its insensitivity.

Williamson (158), following an example of Stephen Schiffer's (331), offers the following as a counterexample to any such suggestion:

BIV-mountain climber. You believe you're not a BIV who thinks he is now climbing a mountain. "Maybe I'm a BIV," you say to yourself, "but not one who thinks he's climbing a mountain. If I'm a BIV, I'm a BIV who thinks he's reading a lousy epistemology paper!" The belief that you are not a BIV who thinks he's climbing a mountain is not only insensitive but also explanatory. If your belief were false, and you were a BIV who thinks he's climbing a mountain, there would be a ready explanation for why you would think you're not a BIV who thinks he's climbing a mountain. BIVs doing such-and-such are always thinking they're not BIVs doing such-and-such. It's something about being a BIV everyone understands.³

Williamson, following out a suggestion by DeRose, offers a repair, which I will quote directly:

(W1) Necessarily, if S knows p then, for some proposition q : q entails p , S sensitively believes q , and $\neg p$ does not explain how S could falsely believe q (159).

Williamson proposes, moreover, that we might weaken the entailment requirement to something else and that we might require that S believe p *because* S believes q (159). Let us say q is a ground for S 's believing p just in case S believes p on the basis of q , q entails p , and q is true. Then (W1) is the claim that if S knows that p , then S sensitively believes a ground, q , for p and not- p does not explain how S could falsely believe q .

Here is how (W1) addresses the strange counterexample. You believe p , that you are not a BIV-mountain climber, on the basis of q , that it seems to you as though you are not climbing a

³DeRose is actually mischaracterized by Williamson at this point. I want to follow out Williamson's suggested fix, however, because it leads us in a helpful direction. I will later return to a discussion of DeRose's *actual* proposal, which, properly understood, is close to my own and avoids the BIV-mountain-climber case altogether.

mountain (Williamson, 159). Your belief that it *seems* thus to you is sensitive. If q were false, you would not believe that q . q entails p . And not- p does not explain your belief that q . That you are a BIV who thinks he's a mountain climber does not explain why you believe that it seems to you as though you are not climbing a mountain. So, while p is insensitive, and while p 's falsity would explain your (counterfactual) belief that p , you still know that p . You still have a sensitive ground for your belief that isn't explained by the hypothesis.

Williamson does not revisit all of the other counterexamples to sensitivity, but his suggestion on behalf of sensitivity works surprisingly well on some of the earlier cases discussed.

Grandmother. Grandmother believes grandson is well on the basis of her belief that she sees that he is well. That she sees grandson well entails that he is well, that she sees he is well is sensitive, and that he is not well does not explain how she could believe that she sees that he is well. All this without any mention of methods.

Dachshund. You believe there is a dog in front of you because you believe there is a dachshund in front of you. If there weren't a Dachshund in front of you, you wouldn't think there was. That there is a Dachshund in front of you entails that there is a dog in front of you. That there is no dog in front of you does not explain your belief that there is a Dachshund in front of you. So you can know there is a dog in front of you by seeing the Dachshund, even if a hyena lurks behind the Dachshund (Williamson, 159).

NFB. I know that p : I don't falsely believe I have hands. On what basis? On the grounds that q : I believe I have hands and I have hands. q is sensitive, because arguably, if it were false, both conjuncts would be false, and I would therefore not falsely believe the conjunction. q entails p . p 's falsity does not explain my belief that q . Of course, p 's falsity predicts that I will believe q . But it does not explain how this false belief came about.⁴

Induction. The induction cases are not clearly solved. Suppose I believe the ice I left in the backyard this morning is now melted. Suppose that belief is insensitive, i.e., it is not the case that if it were false I would not have believed it. What is the sensitive belief that entails and grounds this belief? It's unclear that there is one. And without a relevant sensitive belief we cannot even apply the explanatory condition.

This is significant progress; all cases considered except for induction are resolved nicely. Perhaps there is some independent way of solving that particular problem. But Williamson is not interested in reflecting on the successes of his proposal, because he has his own counterexamples, which are cases of systematic insensitivity:

⁴Williamson himself gives no indication that the NFB cases are to be handled this way but seems to think of his proposal as an extension of DeRose's suggestion that even if p is insensitive, not- p must explain how I came to believe what I take myself to know in order to rule out knowledge that p . However, p is insensitive in the NFB case, while, according to Williamson's analysis, there must be a sensitive belief, q , that entails p in order for S to know that p , so Williamson's modification does not allow that the relevant $q = p$ in the NFB case, as did DeRose's original formulation. Luckily, there is a relevant q in the conjunction of p and I believe that p .

*Revised Dachshund.*⁵ Through a tiny aperture, you see a dachshund. Behind, and totally obscured by the dachshund is a basset hound. Behind and totally obscured by the basset hound is a hyena. You form the belief that you're looking at a dog, and also that you're looking at a dachshund. Now, you confuse dachshunds with basset hounds. They look like the same breed to you: dachshund! If you weren't looking at a dachshund, you'd still believe you were, because you'd be looking at the basset hound. And if you weren't looking at a dog, you'd be looking at the hyena, and you'd still think it was a dog. So, let p be your belief that there's a dog in front of you. Not only is it insensitive, but there is also no relevant belief q that entails p that is itself sensitive. Intuitively, you don't know you are looking at a dachshund, because you'd mistake a dachshund for a basset hound. Yet, when looking at a dachshund, you do know you're looking at a dog; you would not mistake a dachshund for a non-dog or a non-dog for a dachshund.

The Pole. You're looking at a two-foot pole and form the belief p , that the pole is less than twenty feet tall. You underestimate heights just a bit. If it were just greater than twenty feet, let's suppose you would still think it was less than twenty feet. (All we need, strictly speaking, is that it is not the case that you would not believe that it was less than twenty feet.) So, your belief is insensitive. And there is no other height belief q , entailing p , on the basis of which you believe p , and which is sensitive. Say you believe the pole is less than four feet, and that's why you think it's less than twenty feet. Because you are *systematic* underestimator, your belief that the pole is less than four feet is insensitive. If it were just over four feet, you'd think it was just under. Thus, we will search in vain for a sensitive ground for your belief that the two-foot pole is less than twenty feet tall.

Williamson's cases of systematic insensitivity prompt him finally to abandon even his amended version of sensitivity. But this is premature. These are cases of slight insensitivity, and there may be a way to cordon them off. DeRose suggests that while the underestimator may think a pole greater than twenty feet tall is under twenty, he will hold this belief to a very low degree; if forced to guess, he'll say the pole is under twenty feet, but he is not going to bet on it, whereas when looking at a two foot pole, he will hold the same belief to a high degree (Williamson, 161, *ftnt.* 10). Thus, his degree of belief is sensitive to the truth of the belief, even if his credence does not drop below .5 exactly when the pole exceeds twenty feet.

Williamson's reply is unconvincing. He asks us to imagine a being for whom belief does not come in degrees, but is instead an all-or-nothing matter (or at least he has rigid degrees of belief, perhaps not one and zero, but .1 and .9) (Williamson 161). Such a being looking at a two-foot pole could still know that the pole was less than twenty feet, Williamson thinks, even though the being would believe to the same degree of confidence of a pole just more than twenty feet that it is less than twenty feet.

DeRose isn't sure whether such beings could have much knowledge at all, at least knowledge of the "x is less than y ft." variety.⁶ It is, in fact, difficult to imagine what such beings would be like.

⁵Though inspired by Williamson, some details of the case are my own.

⁶DeRose's reply is from conversation and as characterized by Williamson (161, *ftnt.*)

Their betting behavior and belief updating would be truly bizarre. Imagine that a two-foot pole slowly grows taller and taller. As the pole grows taller, these creatures continue to bet that it is less than twenty feet tall and will take the very same odds until exactly one moment when the pole is slightly taller than twenty feet. Then, there is a discontinuous change in credence and they bet the other way, at some particular odds that remain fixed regardless of how tall the pole grows beyond twenty feet.

You don't want such a creature working as your parking attendant. Suppose you tell him the hood of your Bentley, which slopes out of the driver's view, is exactly four feet long. In parking your car, he'll underestimate the distance to the wall, smashing confidently into it. From exactly four feet away from the wall, he will be just as confident that he is more than four feet away as he was when he was a mile away. (If his degrees of belief are .7 and .3, he will, bizarrely, be somewhat wary of smashing into the wall when he is still a mile away from it.)

Who knows what to make of such creatures?⁷ If the objection from systematic insensitivity rests on what we think they know, it should not be given much weight. This may be a legitimate case of spoils to the victor. An otherwise workable theory earns its right to decide obscure cases like this. Maybe what the cases show is only the extent to which ordinary empirical knowledge depends on our capacity for changing our credence smoothly, by degrees. But it would be better to characterize these cases in accordance with intuition and I can imagine them in a sympathetic way.

Suppose our underestimator doesn't just look at the pole and form a belief about whether it is less than twenty feet, but rather looks at it and thinks, "Oh, it's one of those poles, the ones less than three feet. And therefore, it's less than twenty." Then he believes, full-on, that it is less than twenty feet. If the pole had been twenty feet, he would have formed his belief in a different and defective way. He would have thought, "Oh, it's one of those poles, the ones nearly twenty feet." And then he would have formed the very same full-on belief that the pole is less than twenty feet. It is not blazingly obvious that the defective nature of the latter process necessarily condemns the first. So how do we capture the cases?

Here is an idea: relative sensitivity. Suppose that rather than requiring that one's grounds for belief be sensitive, we require that they be sensitive to the truth of that belief for which they are grounds. That is, if p is believed on the basis of q , require that if p were false, you wouldn't believe q .⁸ Let's take Williamson's proposal for sensitivity and modify it as follows:

⁷The creature may have a discontinuous credence function only in part of the spectrum of cases, but similar considerations will apply.

⁸Note: this proposal differs significantly from Dretske's in 'Conclusive Reasons' (1971, 1). For Dretske, the truth of a conclusive reason, R , had to be counterfactually dependent on the truth of the proposition P , for which it was a reason (i.e., R would not be the case if P were not the case). In (W2) it is the *belief* in the ground that must be sensitive to the truth of that for which it is a ground. The relation between grounding and grounded propositions is entailment, though may be weakened. On Dretske's proposal, but not mine, it may be that if the grounded proposition were false, I would still believe all of the propositions that ground it. Because Dretske is focused on the content of the grounding belief rather than the believing, he must characterize the grounds as epistemically meritorious in some way; he says they are either 'known' or an 'experimental state'

(W2) Necessarily, if S knows p then, for some true proposition q: q entails p, S believes p on the basis of q, S's belief that q is sensitive to the truth of p, and $\neg p$ does not explain how S could falsely believe q.

Now, back to the pole case. Suppose S believes p, that the pole is less than twenty feet, on the basis of q, its being less than three feet. q entails p. S would not believe q if p were false, i.e., if the pole were greater than twenty feet tall, S wouldn't believe it was less than three feet tall. And that the pole is greater than twenty feet does not explain how S came to think falsely that it is less than three. Thus, (W2) does not stand in the way of S's knowing that the pole is less than twenty feet.

Suppose S is an underestimator looking at a 19'11" pole and that he underestimates by 6". Now there is no true proposition he believes about the pole that entails that it is less than 20' and also would not be believed if the pole were over 20'. He believes the pole is at 19'11" or less. But, as a 6" underestimator, he might very well hold this if the pole were over 20'. He also believes the pole is 19'5" or less. This, he would not believe if the pole were over 20'. It is not relatively sensitive in the right way.

Revised Dachshund case. Let p = there's a dog. Let q = there's a dachshund. q entails p. S believes p on the basis of q. S's belief that q is sensitive to the truth of p: if there weren't a dog, S wouldn't believe there was a dachshund. And that there is no dog does not explain how S came to believe falsely that there was a dachshund.

Just as a reminder, let's see that you cannot know you're not a victim of a traditional skeptical scenario.

BIV scenario. Let p = I'm not a handless BIV. Let q = I have hands. Suppose S believes p on the basis of q. q entails p. But S's belief that q is not sensitive to the truth of p. If p were false, S would go on believing q. Also, not-p explains S's belief that q. That S is a handless BIV explains why he believes that he isn't.

So, (W2) corrects for systematic insensitivity while preserving the core intuition that we do not know we're not BIVs. Does it get the intuitively right results for all of the other cases, excluding the induction cases? Not quite. With (W1), we were able to dismiss NFB cases by appealing to the sensitivity of "p and I believe p", which grounds "I don't falsely believe p." But on (W2), the sensitivity of a ground for p is not enough. To disqualify p as a skeptical success, a ground must be sensitive to p itself. But if I did believe p falsely, I would still believe (p and that I believe p). So, if we want to address the NFB and induction cases, we will need to modify (W2) further.

Perhaps there is some other way of addressing those cases. I've argued elsewhere that NFB cases may simply be accepted by proponents of sensitivity (2007, 92-93). The sentence "I do not falsely believe p" can express, or at least convey, a number of sensitively believed things:

(1971, 13). The first qualification is circular. I need only say that the grounds are beliefs and that they are sensitive to that for which they are grounds, avoiding the circularity in the knowledge requirement.

It can convey the proposition that I truly believe p.

It can convey the proposition that I do not believe p, a false proposition. (*I* am not one of those delusional folks who falsely believes that p!)⁹

It can convey, *de re*, of a certain proposition I believe, namely, p, that it isn't false.

All of these things are such that if they were false, I would not believe them. They're all sensitive. And all of these are easily confused with the single insensitive proposition in the neighborhood: not-(not-p and I believe p). Perhaps that one insensitive proposition really is not known, but seems to be only because it is embedded in a web of sensitive belief, conveyable by the same natural-language expression.

I have also argued that the induction cases are anything but clear, because we do not know how to resolve the relevant counterfactuals (2007, 94-95). If my trash had not hit the bottom, would I still believe that it had? I'm not sure what to say. It may be that if the trash had not hit bottom, it would have caught high up in the chute and I would have heard it. Or maybe I would have seen signs warning that the chute is full, signs directing me to hold off temporarily on trash disposal. Or maybe I would have received a call from the apartment supervisor drawing my attention to the defective chute. And so on.

If we explicitly rule out such possibilities, if we say that what I know is not just that my trash reached the bottom, but also that it is not the case that my trash failed to reach the bottom in some way as to leave me no indicator that it failed to reach the bottom, then sensitivity fans should own up to my ignorance. I do not and cannot know such a thing. But that does not stand in the way of ordinary inductive knowledge, which is not obviously insensitive.

Since NFB cases cannot be addressed directly by (W2), we may drop the explanatory rider, whose only function was to rule them out, taking us back to a clean, modal relation. We're left with what I call Relative Sensitivity:

(RS) Necessarily, if S knows p then, for some true proposition q: q entails p, S believes p on the basis of q, and S's belief that q is sensitive to the truth of p.

We have here a simple successor to sensitivity, one that solves for systematic insensitivity, but which, admittedly, won't satisfy everyone on NFB and induction cases. I think it can, in fact, be stated even more simply. I have said nothing so far about "believing on the basis of" and I will not say much, because I want to leave (RS) as flexible as I possibly can. But counterfactuals will be a rough guide. If S believes p on the basis of q, then if S didn't believe q, S wouldn't believe p. This dependence needn't be causal, because p and q needn't be distinct. I want to leave it open that $p=q$. Nor should it be read to imply that there is any conscious inference. Likewise, as Williamson suggested, the relation holding between q and p might be weakened to something less than entailment; I leave that an open question. Now, let's say that if p is believed on the basis of q and q entails p and q is true, that q constitutes *grounds* for p. (Remember: everything trivially grounds itself.) Now we can state Relative Sensitivity intuitively as:

⁹I owe this reading to Elliot Paul.

(RS) Necessarily, if S knows p, then S's belief in some ground for p is sensitive to p.

So stated, (RS) is exactly what you want to say explains the skeptic's power. We can't say we know we're not BIVs because if we were BIVs, we would go on believing everything that grounds our belief that we aren't. Some readers may stop at this stage, content that the explanatory power of sensitivity has survived its subjection to the epistemologist's extraordinary counterexamplifying engine. I think (RS) can be improved, however, to better answer the worries about inductive knowledge.

The problem with inductive knowledge is that sensitivity requires that if p were to be false, S would not believe p, where the negation takes narrow scope. Put in terms of possible-world semantics, if p is sensitive, then in all of the nearest not-p worlds, S does not believe p. So far, I have said that in the induction cases, it is unclear whether if p were false, S would not believe it. It is unclear whether if the trash had failed to hit bottom, I would still have believed it had. But "unclear" is not good enough to ensure inductive knowledge. Even if I probably would not have believed it hit bottom in such a case, I may still fail the sensitivity condition. The condition requires that in all of the nearest possible not-p worlds, I refrain from believing p. Just one not-p world tied for nearest, in which I continue to believe p, will rob me of knowledge.

Now, it is difficult to know how to interpret the subjunctive conditional in these induction cases. It can be natural both to say that I would have known it if my trash hadn't hit bottom, because there would likely have been a warning of some kind, or a noise on the way down, and also natural to say that I wouldn't know if my trash hadn't hit bottom, because I did not, in fact, hear a noise or see a warning sign. Contextual cues can resolve the subjunctive either way. Or they can resolve the subjunctive as indeterminate, because I might have gone on believing and I might not have. It's difficult to know what to do with these intuitions.¹⁰

¹⁰It's worth noting that the easy ways of saying my inductive beliefs are sensitive use backtracking counterfactuals. For instance, the non-actual sign warning me of accumulated trash would have to have been installed prior to my dropping of the trash in order for it to effect my belief. David Lewis has argued that backtrackers require a "special" resolution of the counterfactual, and that the "standard" resolution of the counterfactual does not allow backtracking (457). So perhaps inductive knowledge is definitely insensitive after all.

But this needn't frighten away the sensitivity proponent, for two reasons. First, Lewis's case is not airtight. Even the standard resolution of counterfactuals requires some backtracking. For example, when evaluating "If I were President, I would stop this war," we fill in some kind of story – a series of assassinations and appointments perhaps – according to which I became President, so the transition from the actual past to the counterfactual present is smooth (Bennett 202-220). Perhaps in this smoothing, I would get indicators of what has happened to my trash (or my ice, or what have you). Second, even if Lewis is correct, the fact that we do invoke backtrackers in ordinary speech allows us to co-opt them for the purpose of our analysis. If backtrackers are required for sensitivity to work, and backtrackers are sometimes invoked in ordinary speech, then let us allow that sensitivity (or relative sensitivity) employs these conditionals in such a way that allows for backtrackers. The "special" resolution, even if non-standard, is not off limits to the theorist.

Still, the rule of sensitivity seems very strong. Perhaps too strong. The existence of one bad world where you believe p even though it's false, amongst all of the closest not- p worlds, isn't obviously enough to make a successful skeptical scenario. Classic skeptical scenarios have me going wrong in *all* of the nearest worlds where p is false, not just one of them. It seems that any decent scenario should have me going wrong in at least a sizeable portion of not- p worlds. It's far from clear that common sense crowns every scenario in which I might believe p if p were false with skeptical success.

Suppose the actual world harbors a lottery-loving evil genius. The genius watches over me constantly, waiting for me to lose my hands. Should I lose my hands, he would hold a billion-ticket lottery – a very speedy one – to determine whether or not to deceive me into thinking I still had hands. I would be assigned one ticket. If my ticket were a winner, he would make me think I still have hands. If my ticket were a loser, I would stare in horror at the remaining stumps while he does nothing. Now, I do not in fact lose my hands. There is no lottery and no deception. Even if I had lost my hands, I almost certainly would have lost the lottery too, and bemoaned my loss of hands. Yet, according to (RS), if the actual world contains such a genius, I do not now know I have hands. There is one close world at which I lack hands but believe that I have them. That world is as close as any of the others where I lose my hands and I lose the lottery, because there is nothing special about my ticket. (Note: In the actual world, since I don't lose my hands, I don't even have a ticket!)

We can easily remedy the excessive strength of sensitivity by giving the negation wide scope in the interpretation of sensitivity. We may read "If p were false, S would not believe p " as "It's not the case that if p were false, S would believe p ." This sounds intuitive enough to me: if S knows that p , then if p were false, it's false that S would go on believing p anyway. Call this weak sensitivity. S 's belief that p is weakly sensitive iff it is false that if p were false, S would believe p . Applied to Relative Sensitivity, we get Weak Relative Sensitivity:

(WRS) Necessarily, if S knows that p , then S has grounds for p , S 's belief in which is weakly sensitive to p .

Weak Relative Sensitivity allows for inductive knowledge. There is no good reason to think that in all nearby worlds where my trash does not hit bottom, I still believe it does. At the very least, I might find out about it. Yet, I still do not know I'm not a handless BIV, or that my lottery ticket is a loser, because in all of the nearest worlds where these beliefs are false, I go on holding them.

We may have reached a stopping point. (WRS) gets all of the cases right except NFB cases, and those are misleading anyway. It's still simple. It still captures the core intuition behind sensitivity. But suppose we want to avoid an error theory for NFB cases and we do not like my pragmatic strategy for explaining their appeal. Or, suppose we think that weak insensitivity still disqualifies us from knowledge. (WRS) is as far as I have been able to push a purely conditionals-based approach. But we *can do better*--and without launching into realms of technocratic obscurity. We can do so with a notion DeRose and Williamson have already invoked and we earlier abandoned because its only real purpose was to solve the NFB cases: the notion of *explanation*. Recall (W2), our relativization of Williamson's proposal on behalf of DeRose:

(W2) Necessarily, if S knows p then, for some true proposition q : q entails p , S believes p on the basis of q , S 's belief that q is sensitive to the truth of p , and $\neg p$ does not explain how

S could falsely believe q.

We can simplify (W2) because the non-explanatory requirement and the relative sensitivity requirement are closely related. In particular: if the relative sensitivity requirement is met, the non-explanatory requirement will also be met. If belief in q is sensitive to p, then not-p cannot explain S's belief that q; if not-p, then S wouldn't believe that q at all. Not-p cannot explain something it ensures would not happen. Accordingly, if a skeptical scenario is explanatory, insensitivity comes along for free.

This suggests that we try removing the relative sensitivity clause, as it is only *one way* of satisfying the non-explanation requirement. Relative Explanation:

(RE) Necessarily, if S knows p then, for some true proposition q: q entails p, S believes p on the basis of q, and the hypothesis that not-p does not explain how S could have come to believe that q.¹¹

Relative Explanation can be put in somewhat intuitive terms. What it says is that if you know p, then the hypothesis that not-p doesn't explain how you could have come to believe your (actual) grounds for p.

We may be able to improve on (RE) further still by dropping the grounds clause and dropping the strange "could have come to believe". What's left is the following. Explanation:

(E) Necessarily, if S knows that p then the hypothesis that not-p does not explain S's belief that p.

Equivalently: if the hypothesis that not-p explains S's belief that p, then S doesn't know that p.

Think about the case of the BIV who takes himself to be climbing a mountain. Williamson suggested on behalf of DeRose that we know we are not in the scenario because we have a sensitive ground for belief (its seeming to us as though we're not on a mountain) and the scenario fails to explain how we "could have come to believe" this sensitive ground.¹² (RE)

¹¹I insert "the hypothesis that.." because I mean "explanation" to be read in a non-factive way, and I find the factive reading of "explanation" far less tempting with the prefix.

¹²DeRose's actual qualification:

"The limitation of SCA's generalization that's suggested by these cases is this: 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."
(23)

Three observations. First, while Williamson credits DeRose with inspiration for (W2), DeRose's actual proposal evades the counterexample. One of the things you take yourself to know is that it doesn't seem to you that you're climbing a mountain. The hypothesis doesn't explain how you came to believe that without its being true! Second, this is a limitation of DeRose's subjunctive

ignores the sensitivity requirement for the grounding belief itself and deals with the case by noting that there is a belief in the ground that is not explained by the scenario. (E) dispatches with the case even more directly. It is most obvious in the first-person case. Does the hypothesis that I am, right now, a BIV who thinks he's climbing a mountain explain my belief that I am not such a BIV? Clearly not! Anything that explains (to me) my belief that I'm not a BIV who thinks he's a mountain climber must respect my phenomenology (more on which in a moment). And the hypothesis that I am a BIV who takes himself to be a mountain climber clearly violates my phenomenology.

Notice that the locution, "could have come to believe that q" invites a *subjunctive* test for explanatoriness, whereas (E) invites an *indicative* test. If you want to know whether a hypothesis explains how you *could have come to believe* that q, you ask yourself what would be true if the hypothesis *were* true, and whether your purported belief *would* thereby be explained. If the BIV-mountain climber hypothesis were true, there would be an explanation for why you would believe you are not such a BIV -- namely, the trickery of the scientists who would make it seem as though you really are climbing a mountain. The scenario explains, in fact, *all* of my counterfactual beliefs, the beliefs I would have if I were in the scenario. However, the BIV-mountain-climbing hypothesis would not successfully explain my actual belief that it does not seem to me that I am climbing a mountain, and this is one of the grounds for thinking I am not in the scenario. That is, it would not, even counterfactually, explain my *actual* grounds.

Returning to the indicative, if I suppose that it actually *is* the case that I am a BIV who thinks he is climbing a mountain, I get seriously confused. "What!? This is what it feels like to climb a mountain? Like writing a philosophy paper? How disappointing!" Anything that counts as an explanation of my beliefs must somehow accommodate my actual grounds for those beliefs. Else, it fails to be explanatory. And so we accomplish what Williamson accomplishes with the "grounds" clause, but without any explicit mention of grounds.

It is less obvious, but still plausible, in the third-person case that we can dispense with the grounds condition. Sungil is a friend writing a philosophy paper alongside me in the library. Does he know he is not a BIV who thinks he is a mountain climber? I think he does know that, and I think the hypothesis that he is so envatted does not explain his belief. It's true that I take Sungil's phenomenology to be pretty much like mine: very non-mountain-climber-ish, very library-ish. But

conditionals account, but as we have seen, once the explanatory limitation is introduced, it makes the subjunctive conditionals redundant, at least as a necessary condition for knowledge. Third, the limitation operates on the truth component. It requires that not-P entail that some proposition we take ourselves to know is in fact, false, and also explains how that proposition came to be falsely believed. It imputes and explains not only ignorance but also false belief in cases where one takes oneself to know. My earlier proposals (E) and (RE) do not mention knowledge, but only truth and belief. My later proposals (SSI) and (SSS) relate only ignorance -- of which false belief is one variety -- and what we take ourselves to know. That is, the earlier proposals require an explanation of false belief (any false belief, not just cases where we take ourselves to know) to disqualify a belief for knowledge. The later proposals are targeted at skeptical success, and require an explanation of ignorance where we take ourselves to know, and nothing more, i.e., do not require ignorance *by way of false belief*. All are very closely related to, and indeed indebted to, DeRose's proposal.

I admit that it is much less obvious than in the first-person case that I must hold his actual phenomenology (relatively) fixed to explain, to my own satisfaction, or to others', his belief that he is not envatted. I do think I must respect *his* grounds for taking himself to be not so envatted in any satisfactory explanation of one of his beliefs. And I think his grounds include the non-mountainous feel of his surroundings.

Williamson has argued, persuasively, that almost no condition is "luminous", i.e., almost no condition is such that if you are in it, you are invariably in a position to know you are in it (Chapter 4). Thus, Williamson would say that its seeming as if you are not climbing a mountain is, like almost every condition, one that you could be in without knowing you're in it. You could, after all, be in a borderline case of the state -- "This feels much like hill climbing, and hardly mountaineering proper!" you would say -- and not be in a position to know whether your case is or isn't mountainclimber-ish.

But the fact that almost nothing is luminous, and the fact that luminosity cannot therefore usefully divide the internal from the external, makes no difference to my project. I've said a skeptical hypothesis, if it is to meet with success, must respect my phenomenology, but I've not said what that entails. I haven't said, in particular, that to respect my phenomenology is to attribute to me absolutely no errors about my phenomenal states, or to put me in a position to know I am in a phenomenal state whenever I'm in it. A successful hypothesis can say that I am on one side of a phenomenal boundary when I believe myself to be on the other, e.g., it may say I am mistaking a light pain on the skin for a hard itch. But it cannot attribute anything more than minor phenomenological error or it ceases to be explanatory, e.g., if it says I am presently mistaken in thinking it does not seem to me, right now, as if I am climbing a mountain.

Here, I am not importing my tacit commitment to an internalist epistemology; I am just reporting what works and doesn't work as a skeptical hypothesis. I could just as easily have begun this paper with the proposed textbook case: *you think you know that it doesn't seem you're climbing a mountain right now, but maybe it does seem that way!* This is a more obviously failed skeptical scenario than any I offered, because it's not only a mystery how I went wrong in taking myself to know something, it's also staggeringly obvious that I didn't go wrong. Williamson's anti-luminosity argument, compelling as it is, just doesn't generate the same kind of skeptical intuitions as the dreaming argument (nor is it designed to do so).

When I am very cold, first thing in the morning, I'll have no patience for the skeptic who is telling me I may be just over the borderline of not-cold. "Shut up and get me a blanket!" I'll say. Midday, I will hear him out and worry that he is right, eventually agreeing that I am not cold. Moving in the other direction, as the sun sets and I move from not-cold to barely cold to very cold, I will first of all find the scenario true, then I'll think it's false but still explanatory, and then as I grow seriously cold again, will find it utterly non-explanatory. Of course, it will be a vague matter exactly where I begin, or cease, to take the scenario seriously. But why should that trouble me?

There is no analogue of this sort of thing for radical skeptical scenarios regarding the external world, other minds, the past, and so on. There is no relevant continuous dimension such that if I move a little further along it, I will be able to brush off the skeptic in the way I am now brushing off the mountain-climbing-BIV skeptic, anxiety-free. My guess is that this inability to generate skeptical unease, given the phenomenal state I'm now in, is a surer cognitive anchor than luminosity ever was for demarcating an interesting class of states. And it is better, too, than the

"respecting my phenomenology" requirement gestured towards above. We have more of a purchase on what generates skeptical intuitions while in certain phenomenal states and on what counts as an explanation of our beliefs, than we do on 'respecting my phenomenology'. Consequently, I am not attempting to reduce or analyze the idea of explanatory success. If I were aiming for reduction or analysis here, I would go in the opposite direction.

In any case, if you agree with me about what it takes to explain an actual belief, even in the third-person case, then we may proceed with (E). If you do not agree, then proceed with (RE). Either proves more powerful than previous suggestions in explaining away counterexamples to sensitivity.

Take the NFB case. Let p = I don't falsely believe I have hands. Consider the case where p grounds itself. Now, the hypothesis that not- p does not explain how I could falsely believe that p , nor how I do in fact falsely believe p . This works directly, without any detour through some other sensitive belief like the belief that I have hands and believe I have hands. But of course it works in the BIV-mountain climber case as well, since there, my belief is clearly based on some distinct grounds. That I am a BIV-mountain climber does not and would not explain my actual belief that my experience is to the contrary.

Grandmother's belief that grandson is well is not explained by his being unwell. If we use (RE) instead, then consider Grandmother's ground for believing grandson is well, *viz.*, her visual beliefs about his healthy appearance. Those beliefs are unexplained by the hypothesis that he is unwell.

Recall the Dachshund case: you are looking at a Dachshund and form the belief that there is a dog in front of you. But if your belief were false and you were not looking at a dog, you would be looking at a hyena, which you would take to be a dog. Intuitions differ about this case. And they differ, I think, because we are not sure what to say your grounds are for believing there is a dog in front of you. Equally, we are not exactly sure what would count as an explanation of your belief. But if we fill in the story appropriately, so that you first determine that it is one of those low, long-ish dogs, and then infer that it is a dog, then it seems the hypothesis that you are looking at a non-dog does not explain your actual belief, nor would it explain your grounds for belief.

Systematic insensitivity. My actual belief that the pole in front of me (which is two feet tall) is less than twenty feet tall is not explained by its being greater than twenty feet tall. This is so even if, were the pole taller than twenty feet, I would have the same belief I do. A satisfactory explanation of my belief must not contradict my grounds for belief, and I believe the pole is less than twenty feet, because I take it to be one of those little stubby poles. Likewise, if we employ (RE), the hypothesis that the pole is less than twenty feet would not explain my belief that it is less than four feet.

Even the induction cases can now be met in a deeply satisfying way. I have left my ice outside. That my ice has not melted does not *explain* my belief that it has melted. There are ways of filling out the case so that it does explain my belief. But, as it is, baldly stated, it does not explain my false belief. It is much like the hypothesis that I do not have hands. It doesn't explain why I falsely believe I have hands.

The application of (RE) is similar. Would the hypothesis that my ice did not melt explain my grounds for believing it did? It is, I think, indeterminate what would have happened if my ice had

not melted. The hypothesis, considered counterfactually, does nothing to rule out the possibility that the ice fails to melt because my spouse brings it into the house, puts it back in the freezer, and calls me to complain that, once again, I've left a bucket of perfectly good ice out to melt. And in this case, I would have believed that the ice did not melt. So the hypothesis that it didn't melt wouldn't necessarily *explain* my belief that it did, because it might be that I would not even have believed that it melted.

Now, let the relevant p = it's not the case that my ice has failed to melt purely as a matter of chance, without anyone noticing, without any traces of its non-melting that I've actually detected that are different from the actual ones. p is not only insensitive, but it also fails (E) and (RE). This 'pure chance' hypothesis in fact explains my false belief that the ice has melted, and if it *were* true, it *would* explain my false belief as well as my grounds for it.

But here again, I think the verdict is correct: the intuition is that I don't know. This is just a lottery belief. The hypothesis that my lottery ticket is a winner explains my belief that it is a loser. (It is background knowledge, readily available, that losing and winning tickets look just the same.) Likewise, if the "winning ticket hypothesis" *were* true, that *would* explain my grounds for thinking it's false. The fact that some cases will fall in between the lottery case and the standard case of inductive belief discussed -- that in some cases it will be unclear whether there is an explanation or not for the belief that p -- is to be embraced. Induction cases, as they are filled in with more and more details, will come closer and closer to lottery-level stories about *how* we could be going wrong. And at every point, our judgments of whether not- p explains the belief that p will swing in tandem with our judgments of whether S knows that p .

Skeptical Success Revisited

In the previous section, I followed convention in discussing sensitivity and revisions thereof as necessary conditions for knowledge. However, as noted in the introductory section, my primary aim is something different. It is to consider sensitivity, and then my proposals, as candidate answers to the question of what makes a skeptical scenario successful, i.e., what must be added to imputed ignorance to yield skeptical success.

First, notice that success, for a skeptical scenario, is the creation of a powerful intuition that one does not and perhaps cannot know something one takes oneself to know. It may be that a scenario succeeds, in this sense, but that the skeptical intuition generated is somehow overcome, perhaps through exposure to Moorean proofs. So it may be that versions of (E) or (RE) account for skeptical success without providing necessary conditions for knowledge. Second, even if successful scenarios are things one cannot know do not obtain, the project of spelling out necessary conditions for knowledge and for describing skeptical success may be subtly different. Recall (E): Necessarily, if S knows that p then the hypothesis that not- p does not explain S 's belief that p . Transformed into a criterion of skeptical success, (E) yields:

(SE) p is a successful skeptical scenario for S iff the hypothesis that p explains S 's belief that not- p .

Recall (RE): Necessarily, if S knows that p then, for some true proposition q : q entails p , S believes p on the basis of q , not- p does not explain how S could have come to believe that q . Transformed into a criterion of skeptical success, (RE) yields:

(SRE): p is a successful skeptical scenario for S iff the hypothesis that p explains, for each of S 's grounds, q_i , for not- p , how S could have come to believe q_i .

(SE) and (SRE) neatly classify not only all of the original bogus scenarios with which I began the paper, but also all of the cases of insensitive but known belief as explanatory failures, and therefore as skeptical failures. Let me put things in terms of (SE), which I find easier to apply: that I falsely believe I have hands doesn't explain my belief that I have hands; that grandson is not well doesn't explain Grandmother's belief that he is well; that the ice hasn't melted doesn't explain my belief that it has; that I am a brain in a vat who thinks he's climbing a mountain doesn't explain my belief that I am not; that there is not a twenty-foot pole in front of me doesn't explain my belief, given the way things appear to me now, that there is a twenty-foot pole in front of me, even if I am a slight but systematic under-estimator of heights. Any of the scenarios just listed would make for a miserable failure as an introduction to skepticism in an epistemology text, arguably because they fail to explain.

So far, so good. But do (SE) and (RSE) cover the full range of skeptical scenarios? Do they capture what it is that, when added to imputed ignorance, yields skeptical success? There are some prima facie reasons to think they do not. It could be argued that a subject has no belief whatsoever about whether a certain skeptical scenario obtains, never having thought about it. If I have never heard or thought about the Matrix, it is reasonable to say I have no belief about whether I'm in it. But the Matrix scenario entails that I do not know things that I take myself to know, e.g., that I move around freely with my own body. It is therefore a skeptical scenario, and at least moderately successful.

Of course, it is open to reply that I am *disposed* to believe that I am not in the Matrix, and so in this dispositional sense I do believe that I am not in the Matrix. But I imagine not all philosophers will be happy with a dispositionalist response, either because of their general views on belief or because they will fashion counterexamples to this particular proposal. Some will imagine skeptical scenarios that by their very nature can never be disbelieved although they can be considered (a "killer-evil-genius-scenario"). Some will imagine skeptical scenarios that, when considered, actually *induce* belief in their truth.¹³ Second, some skeptical scenarios may deny that we have beliefs: you are dreaming, and because you are dreaming, lack any beliefs, even the belief that you are not dreaming.¹⁴ This may follow from a certain philosophical theory of belief that requires it to be connected to behavior in a way that is lacking in sleep. Since, according to the scenario, you don't have any beliefs, the scenario can't explain your *belief* that it doesn't obtain. Thirdly, and this is a more systematic consideration, note that (SE) and (SRE) operate on the truth

¹³Nevermind that there will be some who insist that these are masked or finked dispositions. Others will find the division between masked dispositional belief and unmasked non-belief to be indefensible.

¹⁴I owe this objection to Zoltan Gendler Szabo. Zoltan mentioned Malcolm (1959) on dreams, but surely the Wittgensteinian idea is that the proper understanding of dreams makes the dreaming hypothesis a non-starter. Since I am now judging myself not to be dreaming, I'm now judging, and therefore, am not dreaming -- a clever, if hopeless response to the dreaming hypothesis. What a Wittgensteinian will never do is marry behaviorism about belief with skepticism, as the envisioned scenario requires.

condition, the truth of what is believed. But, as I began the paper, one of the other conditions, like the belief condition or the warrant condition, may fail according to a skeptical scenario.

If you take yourself to know something, then presumably, you take yourself to meet the necessary conditions for your having knowledge, and you take yourself to know those conditions are fulfilled as well. This pressures us towards thinking that any skeptical scenario denies the truth of some putative piece of knowledge, whatever condition it attacks -- truth, belief, or warrant. But again, if you haven't done an analysis of knowledge or haven't thought about the conditions, then, as in the above cases, it may seem strange to say that you believe those conditions obtain, or, more strongly, take yourself to know that they obtain, in every case in which you take yourself to know something.

A broader, more straightforward characterization, better suited for answering my question than the more traditional question of what is required for knowledge, and which may or may not be equivalent to (SE), depending on certain auxiliary assumptions about the nature of belief and the extent of what one takes oneself to know, would be this:

(SSI) p is a skeptical success iff p entails that S does not know something, q , that S takes herself to know and p explains how S takes herself to know q even though she doesn't.

I prefer (SSI) as it is, but it can also be given a subjunctive treatment that refers to grounds. The subjunctive version says:

(SSS) p is a skeptical success iff p entails that S does not know something, q , that S takes herself to know and p *would* explain how S could have her (actual) grounds for q , even though she doesn't know that q .

These formulations avoid the complications above. Even if one has no belief whatsoever about p , or for extraneous reasons is inclined to *accept* p when p is considered, p can still qualify as a skeptical success according to (SSI) and (SSS) so long as it denies that one knows things one takes oneself to know and explains how this is, or could be, the case. Even if the skeptical scenario according to which dreamers have no *beliefs* is coherent, dreamers still, in some sense, typically take themselves not to be dreaming. Sometimes, of course, one dreams that one is dreaming. (Note: the hypothesis that I am now in the kind of dream in which I take myself to be dreaming is a skeptical failure, for the same reason as the mountain-climbing BIV.) There *is* a difference between these two kinds of dreams, dreams in which I do and dreams in which I don't take myself to be dreaming. And if this difference doesn't qualify as a difference in belief, it still qualifies as a difference in how one takes oneself to be.

Suppose a hard-core externalist denies that there is any sense to be made of "taking oneself to know..." apart from believing oneself to know. In that case, I don't understand the so-called skeptical scenario on offer. How is it supposed to go?

*You're dreaming. You don't have any beliefs when you're dreaming, and you don't take anything to be a particular way either, and therefore you don't take yourself to know anything at all. Hence, the present scenario, while clearly a skeptical scenario, doesn't explain your taking yourself to know things even though you don't know them.
Counterexample!*

Is this skeptical success? Nevermind success. Is it even skeptical? It doesn't impute to me any misrepresentation of the world or any ignorance where I think I have knowledge. It doesn't impute any misrepresentation even about whether I misrepresent anything.

Fallout

The first two sections of this paper are mutually supporting. In the first section, I made a *prima facie* judgment that what makes a skeptical hypothesis successful is its explanatory prowess. I drew that conclusion from a direct intuition about what was missing in the simple considered scenarios (no back story!) and on an observation about the tension between probability and explanatoriness, a tension that expresses itself in parallel fashion in skeptical and anti-skeptical ways. In the second section, I examined sensitivity and the counterexamples to its necessity for knowledge. Skeptical success requires that we come to have a strong intuition of ignorance. And if insensitive beliefs are not known, then perhaps insensitivity could explain skeptical success. But if insensitive beliefs are sometimes known, then insensitivity, strictly speaking, is not what explains skeptical success. It may *indicate* skeptical success, but it cannot guarantee it. Fortunately, the explanatory proposals account for all the counterexamples, keeping the initial intuitive judgment intact.

Moreover, the intuitive appeal of sensitivity can be accounted for. If p is insensitively believed, then typically but not always, the hypothesis that not- p explains the belief of p . (It does in all of the classical skeptical scenarios anyway, and until recently, those were the only ones people were talking about.) Finally, explanatoriness, as a marker of skeptical success, accounts for our differing intuitions about lottery-based skepticism and denials of ordinary claims to inductive knowledge. The lottery setup provides a ready-made explanation for error, whereas the story of how false inductive belief arises is sometimes absent. If a case of inductive skepticism is properly filled in, if it becomes genuinely explanatory of false belief, then it also succeeds in bringing about in us the strong intuition that we do not know we are not in the skeptical scenario, just as the account predicts. In the last section, I transformed the accounts reached in the traditional discussion of the necessary conditions for knowledge into accounts of skeptical success, and I argued that this transformation is not trivial, but once accomplished, much more natural.

The views I've indicated here are woefully undeveloped. This is not by accident. The less I say about explanation and about grounds the less likely I am to say something false, and the more plausible the general point about explanation being the marker of skeptical success. But I should outline a few relevant, and I hope, uncontroversial features.

Degrees. Some explanations are better than others. Since my account simply invokes the difference between hypotheses that explain and hypotheses that do not explain, I assume some kind of threshold for how well something must explain to count as an explanation. I have said nothing about what that threshold is. One possibility is that the threshold is extremely high, and nothing counts as an explanation except an a priori proof from first principles. Another possibility is that the threshold is so low that any attempted, coherent answer to a why or how question that is *prima facie* compatible with the explicandum is an explanation, just not a good one. If we think an explanation is necessarily an a priori proof from first principles, there are really no successful skeptical scenarios. If we think an explanation is any answer that does not obviously contradict what it attempts to explain, then skeptical success is as easy as denying a case of knowledge.

Factivity. There is a reading of 'explanation' on which what is false is never an explanation. This reading would doom the indicative formulations I've offered and I do not intend it.

Audience-relativity. You begin your text: "You think your car is where you left it. But what if it isn't?" The paranoid reader might begin to worry about the whereabouts of her vehicle, concocting more complete scenarios about how it could have been stolen or borrowed. In a few moments, she brings herself around to the conclusion that she doesn't know if her car has been moved. But the dogmatic sort can brush off the suggestion. "If my car had been moved, I'd know about it," she'd say, "People ask me before they move my car!" And people do. Descartes did not rely on the reader. He filled out his own scenarios so that paranoid and the bold alike would suffer. But lesser hypotheses reveal the variability of success, and this should not come as a surprise if skeptical success is explanatory success. What counts as an explanation depends on one's background beliefs and tendencies towards belief.

If this all sounds contextualist-friendly, that's because it is. If the magic of skeptical scenarios lies in their ability to explain, then that magic will arise in some groups and not others, depending on their background beliefs and the assumptions in play. Should we take the intuitions generated by skeptical scenarios to guide the semantics of knowledge attribution, then that semantics will be just as shifty and context-dependent as explanation itself.¹⁵ ¹⁶

¹⁵This is by no means the first mention of explanation in epistemic theory. Goldman (1984), Rieber (1998), and Jenkins (2006) all offer explanatory accounts of knowledge. However, these all give sufficient as well as necessary conditions for knowledge. p is known, on these accounts, just in case the right kind of explanatory relation holds between p and the belief that p . Noted problems for such accounts include knowledge of necessary truths and knowledge of the future. I would be happy if some explanation-based, or explanation-related theory of knowledge could be defended (see Neta 2002, for example). The goal of the present paper, however, is different, more modest, and compatible with a much broader range of theories. Though I think we have mathematical knowledge and knowledge of the future, it is not obvious to me that mathematical truths explain our beliefs about them, nor that the future explains our beliefs about it. In general, it is not obvious that we are ignorant of everything about which our beliefs are not explained by the facts. It is obvious, however, that if a hypothesis explains my believing all of my grounds for thinking it's false, then I am disinclined to say that I know the hypothesis is false. That is, skepticism is driven not by a lack of explanatory connection between belief and fact, but the existence of an explanatory connection between the scenario and ignorance, typically between the scenario and false belief. This echoes the contrast between the probability skeptic and the traditional skeptic mentioned in the first section.

¹⁶This paper grew out of comments I gave on a paper by Jonathan Vogel and later a seminar I taught with Keith DeRose on knowledge and conditionals. Thanks to Jonathan, Keith, the participants in the seminar, and also the philosophy faculty at Yale where an earlier version of the paper was discussed in a lunchtime colloquium. I'm also indebted to Peter Klein and Ernest Sosa, who started me thinking about these topics, and to Amia Srinivasan and Tamar Kreps for comments on earlier drafts.

1. DM Armstrong, *Belief, Truth, and Knowledge* (Cambridge University Press).
2. Jonathan Francis Bennett, *A Philosophical Guide to Conditionals*, 2003.
3. J. Cornman, G. S. Pappas, and M. Swain, *Essays on Knowledge and Justification*, 1978.
4. Troy Cross, "Comments on Vogel," *Philosophical Studies* 134, no. 1 (May 23, 2007): 89-98, doi:10.1007/s11098-006-9014-7.
5. Keith DeRose, "Solving the Skeptical Problem," *The Philosophical Review* 104, no. 1 (January 1995): 1-52.
6. Rene Descartes, *The Philosophical Writings of Descartes: Volume 2* (Cambridge University Press, 1985).
7. Fred Dretske, "Epistemic Operators," *The Journal of Philosophy* 67, no. 24 (December 24, 1970): 1007-1023.
8. Fred Dretske, "Conclusive reasons," *Australasian Journal of Philosophy* 49, no. 1 (1971): 1-22.
9. Edmund L. Gettier, "Is Justified True Belief Knowledge?," *Analysis* 23, no. 6 (June 1963): 121-123.
10. Alvin I. Goldman, "A Causal Theory of Knowing," *The Journal of Philosophy* 64, no. 12 (June 22, 1967): 357-372.
11. Alvin I. Goldman, "Discrimination and Perceptual Knowledge," *The Journal of Philosophy* 73, no. 20 (November 18, 1976): 771-791.
12. Alvin I. Goldman, "Review: [untitled]," *The Philosophical Review* 92, no. 1 (January 1983): 81-88.
13. A.H. Goldman, "An Explanatory Analysis of Knowledge," *American Philosophical Quarterly Pittsburgh, Pa* 21, no. 1 (1984): 101-108.
14. Nelson Goodman and Hilary Putnam, *Fact, Fiction, and Forecast, Fourth Edition*, 4th ed. (Harvard University Press, 2006).
15. Stephen Hetherington, *Epistemology Futures* (Oxford University Press, USA, 2006).
16. William James, *The Will to Believe and Other Essays in Popular Philosophy* (Dover Publications Inc., 2003).
17. C. S. Jenkins, "Knowledge and Explanation," *Canadian Journal of Philosophy* 36, no. 2 (2006): 137.
18. D. Kahneman and A. Tversky, "Extensional versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment.," *Psychological Review* 90, no. 1983 (1983): 293-315.
19. Peter D. Klein, "A Proposed Definition of Propositional Knowledge," *The Journal of Philosophy* 68, no. 16 (August 19, 1971): 471-482.
20. Peter D. Klein, "Misleading 'Misleading Defeaters'," *The Journal of Philosophy* 76, no. 7 (July 1979): 382-386.
21. K. Lehrer, "Why Not Scepticism?," in *The Philosophical Forum*, vol. 2, 1971, 283-298.
22. Keith Lehrer and Thomas Paxson, "Knowledge: Undefeated Justified True Belief," *The Journal of Philosophy* 66, no. 8 (April 24, 1969): 225-237.
23. David Lewis, "Counterfactual Dependence and Time's Arrow," *Noûs* 13, no. 4 (November 1979): 455-476.
24. Steven Luper, *The Possibility of Knowledge: Nozick and His Critics* (Rowman & Littlefield Publishers, 1987).
25. W. Lycan, "On the Gettier Problem'," *Epistemology Futures*: 148-168.
26. Norman Malcolm, *Dreaming* (Routledge & Paul, 1959).
27. Ram Neta, "S Knows That P," *Noûs* 36, no. 4 (2002): 663-681.
28. Robert Nozick, *Philosophical Explanations*, New edition. (Harvard University Press, 1981).

29. G. S. Pappas, "Marshall Swain, eds. 1978," *Essays on Knowledge and*.
30. Steven Rieber, "Skepticism and Contrastive Explanation," *Noûs* 32, no. 2 (June 1998): 189-204.
31. Sherrilyn Roush, *Tracking Truth: Knowledge, Evidence, and Science* (OUP Oxford, 2007).
32. Bertrand Russell, *Human Knowledge: Its Scope and Limits*, 2nd ed. (Routledge, 1994).
33. Bertrand Russell, *The Analysis of the Mind* (NuVision Publications, 2008).
34. Stephen Schiffer, "Contextualist Solutions to Scepticism," *Proceedings of the Aristotelian Society* 96, New Series (1996): 317-333.
35. Ernest Sosa, "How to Defeat Opposition to Moore," *Noûs* 33 (1999): 141-153.
36. Marshall Swain, "Epistemic Defeasibility," *American Philosophical Quarterly* 11 (1974): 15-25.
37. Amos Tversky and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," *Science* 185, no. 4157, New Series (September 27, 1974): 1124-1131.
38. J. Vogel, "Tracking, Closure, and Inductive Knowledge," in *The Possibility of Knowledge*, 1987.
39. Jonathan Vogel, "Subjunctivitis," *Philosophical Studies* 134, no. 1 (May 23, 2007): 73-88, doi:10.1007/s11098-006-9013-8.
40. Andy Wachowski and Larry Wachowski, *The Matrix* (Warner Brothers, 1999).
41. Timothy Williamson, *Knowledge and Its Limits* (Oxford University Press, USA, 2002).