

# THE JOURNAL OF PHILOSOPHY

VOLUME CVII, NUMBER 5  
MAY 2010

*page*

221 *Lessons from the Context Sensitivity of Causal Talk* Eric Swanson  
243 *Closure on Skepticism* Sherrilyn Roush

COMMENTS AND CRITICISM

257 *Why Responsible Belief Is Blameless Belief* Anthony Booth  
and Rik Peels

BOOK REVIEWS

266 Galen Strawson: Colin McGinn  
*Real Materialism and Other Essays*

272 NEW BOOKS

*Published by The Journal of Philosophy, Inc.*

## CLOSURE ON SKEPTICISM\*

It is received wisdom that the skeptic has a devastating line of argument in the following. You probably think, he says, that you know that you have hands. But if you knew that you had hands, then you would also know that you were not a brain in a vat, a brain suspended in fluid with electrodes feeding you perfectly coordinated impressions that are generated by a supercomputer, of a world that looks and moves just like this one. You would know you were not in this state if you knew you had hands, since having hands implies you are no brain in a vat. You obviously do not know you are not a brain in a vat, though—you have no evidence that would distinguish that state from the normal one you think you are in. Therefore, by *modus tollens*, you do not know you have hands. At least, the skeptic has a devastating argument, it is thought, if we grant him closure of knowledge under known implication, which many of us are inclined to do: roughly, if you know  $p$ , and you know that  $p$  implies  $q$ , then you know  $q$ .<sup>1</sup>

To say that this is an intuitively compelling argument is an understatement; the project of finding a reply that is not table-thumping, or obfuscating, or special pleading has exercised philosophers for some time. The steps of the argument have been scoured in detail to find cracks that will yield under pressure. Some of these efforts have been intriguing, and illuminating, and some, I think, even provide dialectical victories that shift the burden of proof back to the skeptic. For all this, though, as I will argue, we have missed a very simple point: though the skeptical argument above is valid, it has a false premise, namely, the claim that the thing we seem obviously to know *implies* the thing we seem on inspection obviously not to know. I will argue that this part of the argument cannot be repaired in a way that preserves the skeptical threat. Thus, if the skeptic wants to convince us to worry about our ordinary knowledge, he will have to come up with a completely different argument.

Closure of knowledge under known implication (hereafter “closure”), is necessary for the skeptical argument presented above but obviously not sufficient. For the closure principle to apply to our case, we would have to know that having hands implies that one is not a

\* Thanks to John MacFarlane and Paolo Mancosu for helpful discussion.

<sup>1</sup> For an up-to-date discussion of this argument, see John Greco, “External World Skepticism,” *Philosophy Compass*, 11, 4 (July 2007): 625–49.

brain in a vat. We cannot know that, as epistemologists are already aware, because the implication does not hold and false claims cannot be known. The implication does not hold because one could be a brain in a vat, so far as that is described above, with hands. The hands would be attached seamlessly to the brain, hence yours in an undeniable sense. These stipulations describe a scenario no less plausible than the original one of a brain in a vat. The scenario ruins the implication the skeptic needs because a handed brain in a vat is a counterexample to the claim that having a hand implies you are not a brain in a vat.



The Addams Family—Season 1  
 © 1964 Orion Pictures Corporation  
 All Rights Reserved.  
 Courtesy of MGM CLIP+STILL

Epistemologists are aware that the implication claim first stated by the skeptic does not hold, due to the possibility just described, so the implication claim typically gets propped up in the obvious way, by saying that having hands implies one is not a *handless* brain in a vat. Sometimes one puts a tone on the emphasized word to convey the judgment that this detail is tiresome. One then moves along in development of the skeptical line to get to the more interesting issues, confident that the patch has done no harm to the argument because implication has been achieved. However, it is not enough that there be an implication. It must be an implication from something we think we do know to something we pretty clearly do not, in order to set us

up for a *modus tollens*. What is wrong with this particular patch is that weakening the conclusion to “I am not a handless brain in a vat” trivializes it for this purpose. If we assume I know that I have a hand, then we should not have the slightest hesitation to credit me with knowledge that I am not a handless brain in a vat.

No appeal to the closure principle is needed to support this conclusion. The claim is independently obvious because that you are not a handless brain in a vat is just not much to know. If we know that someone has hands then it follows that she is not a handless person with high blood pressure, or a handless victim of child abuse, but this would not give us any assurance that she need not go to a doctor for these conditions. To a person who already knows she has hands these claims say nothing at all about how far she might or might not be susceptible to heart disease or suicide. For this reason they are statements that it is trivially easy to know if you know that you have hands. If I know that I have hands, then in virtue of that I know I am not a handless anything. The implication is achieved in the skeptical argument, but only by letting the issue of brains in vats swing free of it.

The problem with my claim, one might think, is that it assumes that whether or not one has a hand is independent of whether or not one is a brain in a vat. The blood pressure example would look very different if not having a hand was correlated with having high blood pressure and you knew it. Then, indeed, finding you have a hand would give you a reason not to worry about your blood pressure. In our case, one might say, not having a hand is part of what we *meant* by being a brain in a vat. It is not an extra piece of news. The word “handless” gets added to the conclusion of the skeptic’s argument only in order to make this explicit, so that one can see how clear the implication is. This idea is also a good explanation of our tone of tiresomeness—it should be obvious that a brain in a vat, in the sense we had in mind, has no hands. The implication holds, and the conclusion is not trivial.

If this is what we meant, then, I submit, it is not what we wanted to mean, or should have meant, given our collective state of puzzlement and distress over this skeptical argument, for the conclusion imagined is still trivial. Having a hand does make you distinct from the brain in a vat of imagination that has no limbs, but it does so in only one respect. It tells us nothing about whether you resemble it or not in any other respect. Let a brain in a vat be a thing that by definition has no hands. Having a hand still allows you to be a thing that is like a brain in a vat in every respect except that it has a hand seamlessly attached to it. The question now is how significant it is to find out that you are not a brain in a vat, when you still could be the same thing but for a hand attached; you still could be systematically deceived about just

about everything. The possession of these hands does not imply the thing we seem obviously not to know, which is that we are not subject to systematic deception. It is the latter concept that insures the intuition that we do not know the conclusion of the skeptical argument. Either the conclusion of the skeptic's argument is weak enough to be implied by the premise that I have a hand, but not strong enough to seem hard to know, or the conclusion of his argument is strong enough to appear obviously unknown to us, but not weak enough to be implied by my having a hand.

If I am right, then why have we been under the impression all this time that the adjusted conclusion "I am not a handless brain in a vat" is nontrivial? One reason is that philosophers are like all human beings in being susceptible to associational "thinking," that is, in drawing conclusions that have not been stated, purely on the basis of the proximity of words to one another. All people are sometimes victims, for example, of the devices of highly trained advertising agencies that do psychological research on how we are moved by associations. There was an ad recently that said, above a vivid picture of a train, "Legally, we can't say you can throw it under a train," of the TOUGHBOOK laptop computer. The ad did not assert that you can throw it under a train (and have it survive), but because precisely that clause was inscribed—see the original sentence—an exaggerated impression was created, in just about everyone I would venture, of just how tough the TOUGHBOOK is. Similarly, the words of our adjusted conclusion are "I am not a...brain in a vat," and this created a strong impression that this sentence without the ellipses had been asserted, or at least that some information was conveyed about this matter. Philosophers are not immune to such unconscious mistakes; we are all apt to make them when our conscious attention is directed elsewhere.

A second reason that the sentence "I am not a handless brain in a vat" seemed to carry the content that I am not a brain in a vat is conversational implicature. Suppose a man says that he enjoys talking to me. I ask him whether he has a wife and he replies "I don't have a wife I can *talk* to," where the word "talk" is not only emphasized but raised in pitch. The content of his reply contains no information about whether he has a wife. However, the emphasis conveys very clearly that he does. What is relevant about this case is that the content of the sentence is perfectly consistent with the message that he does have a wife, despite the fact that the sentence contains the phrase "I don't have a wife." Similarly, the content of the sentence "I am not a handless brain in a vat" is perfectly consistent with my being a brain in a vat. This is why it is even possible to make a strong suggestion that I *am* a brain in a vat, by saying "I'm not a *handless* brain in a vat," if

the word “handless” is emphasized and higher in pitch. To say the sentence “I’m not a handless brain in a vat” with a high-pitched emphasis on “handless” would reveal the triviality of the claim with respect to the matter of whether one is a brain in a vat (on the assumption one does know one has a hand), but I never hear epistemologists say the sentence that way.

The word “handless” is sometimes introduced with an emphasis that lowers the pitch on this word (to convey that tone of tiresomeness), but this hides the fact that no information has been conveyed that I am not a brain in a vat, just as “I don’t have a wife I can talk to,” may well fail to set off the wife alarm if the word “talk” is not raised in pitch. Admittedly, epistemologists also sometimes say the conclusion of our argument straight. In that case one is likely presuming that the word “handless” merely brings out an assumption already in what we meant by the phrase “brain in a vat,” and politely leaving out the tone of tiresomeness. As I argued above, tone or no tone, the conclusion that follows is thereby trivialized, and easy to know if you know you have a hand.

Another plausible reason for the mistake is an equivocation on the term “brain in a vat.” One could mean by this phrase a literal, specific image of a brain with no limbs or funny stuff, or one could mean this image as a kind of stand-in for any of a host of scenarios in which one is systematically deceived. Knowing you have a hand is plenty good enough to rule out the first, and miserably inadequate for ruling out the second, even if the host of scenarios is a set of small variations on a single theme. Our confidence in the implication has come from the first reading of “brain in a vat,” and our confidence that the conclusion is something we do not know comes from the second. By equivocation we conclude that something we obviously do not know is implied by something we obviously do.

The initial patch I have described is of course not the only recourse the skeptic has. He could find a different way to weaken the conclusion, in which case the task is still to avoid making it trivially knowable. I will canvas another way of using this conclusion-weakening strategy below. The other obvious approach is to strengthen the premises. In this strategy we would keep the conclusion the same—I am not a brain in a vat in the originally intended sense—and add premises to make sure that what we think we obviously know does imply this conclusion it seems we clearly do not know. This turns out to be harder than it may seem, for even if we added claims that we have feet, and likewise for other body parts, things we know just as obviously as we know about our hands, the possibility of systematic deception does not go away. We could imagine an entity like a brain in a vat in every respect except that it had hands, feet, and so on, attached.

The number of attachments is not the issue in how much it takes to rule out systematic deception. The poor captured people who are used as batteries by the Matrix of movie fame have kept their entire bodies, but their brains are being fed impressions of a colorful world nothing like the dank storage facility in which their pods are suspended. This scenario would be as disturbing as the image of ourselves as “mere” brains in vats, and as obviously difficult to rule out. What makes something a brain in a vat in the relevant sense is that you are not related to the real world in the way you appear to yourself to be, and you have no indication of that; thus the world your hands and feet exist in is nothing like the world of your impressions. I will call this scenario in which you are systematically deceived one where you are a *brain in a vat* to indicate that this feature is essential to the scenario, while failing to have limbs, for example, is not. The denial of this envattedness, which it seems independently obvious we do not know, needs to follow from things we think we clearly do know; knowing that one has ever so many hands and feet does not rule out the disturbing and indiscriminable, and hence essential, feature of the *brain-in-a-vat* hypothesis.

What would rule out the skeptical hypothesis? The *brain-in-a-vat* hypothesis implies something about the vat-brain person’s relation to the world: her lack of discriminating evidence about not only the world but her situation in it. The skeptic’s premise must rule out all logically possible ways of realizing this. Once we understand this another repair strategy suggests itself. Merely that I possess hands is not enough, but perhaps this is because that claim does not say that the hand is connected up to my impressions, and intentions to move, in the normal way that it is when I have evidence and a nondeceptive set of impressions of the world. It seems that the claim of a hand that is normal in the relevant way—which we can as innocently agree we have knowledge of, when the skeptic asks, as we can agree about the previous claim—will do the trick of implying the claim that I am not a *brain in a vat*, since the *brain-in-a-vat* hypothesis says I am systematically deceived about everything, and this says there is something about which I am not deceived.

Call the first type of hand that is unconnected to my impressions a “floppy” hand, and the second a “hooked-up” hand. The problem is that a hooked-up hand is also far too little to rule out the *brain-in-a-vat* hypothesis. Though the hands of the poor people in the actual Matrix are floppy in the sense just introduced, we can easily imagine them having hooked-up hands, as long as we also enlarge the pod to allow their free movement. Their movement would require some movement of the arms, but they have those too and we can imagine them

hooked up. The impressions they have of their hands and arms, both sensory and motor, would come from the hands and arms, whereas their impressions of everything else would come from the supercomputer stimulations. The real and the fake would have to be coordinated with each other, the fake impressions of objects responding just as real objects would, to the interventions of the real hands. But there is nothing impossible about this.

An instance of the idea would be a video game: your control of the joystick is real, but what it is controlling is representations of things that are not real, and what it is controlling is a world that the player can increasingly come to inhabit as if it is real. Suppose such a player becomes fully entranced, without any longer having a sense of the set-up or movement of the rest of his body. Then he is systematically deceived. He will not come out of that world by any prompt within the game-world, but only by a screen that pops up saying he has run out of money, or by a bout of thirst, or intervention from a parent. We can imagine a case in which none of those external cues are available. It is clear that having hooked-up hands does not imply that one is not a *brain in a vat* any more than having a collection of floppy hands and feet did. When we want to know that we are not systematically deceived we expect more than that there is one little thing from which we are not hopelessly unconnected. Thus the “not” in the phrase “I am not a *brain in a vat*” does not function intuitively as it seems it should logically, issuing in a weak claim because it is denying a strong claim. Intuitively, “I am not a *brain in a vat*” means that most ordinary things are pretty much as they seem. If it does not mean this, then it is a claim that is too easy to know for the skeptic’s purposes—just wave your hand. The reason the denial of the real *brain-in-a-vat* hypothesis is so strong is that “I am a *brain in a vat*” is a disjunction of lots and lots of quite similar ways one could be uncorrectably deceived about just about everything: a brain disconnected from everything except your little toe, deceived about everything except the existence of the floor, and so on. Each of those is easy to know in virtue of its being easy to know you have a little toe, and that there is a floor, but no one or two of them denies the sort of systematic deception the *brain-in-a-vat* hypothesis is about.

One might wonder if the problem is that we have not taken into account enough body parts. Having a greater number of floppy body parts did not help, but maybe it will if the numerous parts are hooked up. However, we can draw out the scenario just discussed with any number of body parts we like by imagining the interface between the joystick and the hand growing into an interface between the entire body and a control surface. Now my whole body is doing every



motion I think it is doing, and I am feeling whatever is impinging on the surface of my body. However, none of my impressions bears any indicative relation to the way the world is. The body is pushing and pulling around a real interface, but the interface is pushing around false representations (from my point of view), or objects that do not match my impressions (from an objective point of view), or nothing at all. I could be a whole-body-hooked-up *brain in a vat*.

Perhaps, then, it is not about me and my body, and ruling out the *brain-in-a-vat* scenario requires adding to the premises some things that I apparently obviously know about the *world*. Take the table of skeptical lore. This will not do either, since that premise typically states only that a table exists, and we already know that mere existence allows the possibility of floppiness—where I have no appropriate connection to the table. What if we suppose that I am hooked up to the worldly object, the table, in some appropriate way, say causally. Suppose also that my visual impressions of the table are perfectly coordinated with my other impressions of the world, whether those are fraudulent or true. But this does no good. I could be resting my arms on a table while the rest of my being is perfectly engrossed in a video game on the screen in front of me. The table impressions are properly produced: it is not just that when I have the leaning feeling in my arms it is because my arms are leaning, as we had already with the hooked-up arms, but also that when the arms are truly leaning, and I am having the impression of their leaning on a table, they are leaning on a real table. Apart from the hooked-up table, though, the entire world of my impressions is a fraud. Make the screen bigger and bigger until it surrounds me; hook me up to a feeding tube; make sure the game world never ceases to be interesting and has a backup generator; imprison anyone who might care to save me; and I am a *brain in a vat*.

One might think that the problem is that we are only considering hooking me up in the normal way to *one* object. The world has many objects, and if we suppose me hooked up to many, many of them, then we are imagining a scenario in which I surely cannot be deceived about very much in my physical surroundings. Is that not enough? Given that we are assuming from the previous steps that I know I am hooked up to my whole body, too, does this scenario not, for all intents and purposes, rule out the possibility that I am a *brain in a vat*? We can look at this approach in two different ways, as either a strengthening of the premises or a weakening of the conclusion. We will see that neither strategy helps the skeptic.

Our strategy now will be to put into the premises enough claims about body parts and objects that a robust denial of the *brain-in-a-vat*

thesis will be implied. In strengthening the premises this way we want to include enough hooked-up objects to insure that I am pretty much in touch with the world around me. Throw in the table, the chairs, the kitchen sink, the lamps and couches, the truck I see outside the window, the sunshine, the floor and ceiling, the walls. Are we there yet? Does all of this imply we are not systematically deceived in the appropriate sense? One problem is that there are a whole lot of things left off of this list. Does the friend you think you just talked to on the phone exist? Is there really a building supporting the room you are sitting in when you are not looking at the building? Why think that closed closet door does not open into outer space? Assume that you do know all of those things you list. They do not imply what the skeptic needs because the list you make, however long, will always leave out an infinite number of important aspects of the world. On the other hand, the things you will manage to list do not appear to imply anything that it would be surprising to think you know on the assumption that you know *them*. For example, it will not be doubtful that you know your hand is not a fake hand. To assist the skeptic we constructed strengthened premises with the object of making them imply a denial of systematic deception, so we had to assume that the hand you know about is a hooked-up hand, a really hooked-up hand.

The list of things that would need to be claimed in a set of premises implying the denial that one is a *brain in a vat* is of course infinite. However, given infinite time one could verify each claim on the list, the way one does with the claim that one has a hand, by directly inspecting them seriatim. The problem is that the knowledge so produced that the closet door does not open to outer space expires when I move away to inspect the lamps in the living room. Can we not have that knowledge in a different way? Not if we are trying to help the skeptic, whose target is those of our beliefs that we think we most obviously do know. We need to make the claims that go into the skeptical argument's premises very, very hard to believe I do not know, the way that it is hard to believe that I do not know I have a hand, since I can feel it and wave it in front of myself. Much if not all of our confidence that we know we have hands is this direct verification. This cannot be done with all of the claims we need in the premises, even if we cut off the list to a large finite set, because we cannot sufficiently directly verify them all at the same time, even roughly, which is what we need to do in order to assert our knowledge of them as premises of a single argument.

The things that we can claim simultaneously obviously to know do not appear to be strong enough to imply anything that we obviously do not know, and so surely not that we are not systematically

deceived. But one might think there is an obvious solution to all of this. You can express all of that information, that there is a table, chairs, sunshine, a building supporting me, whatever you see, simply by making a generalization that includes all of those examples without listing them individually. The generalization captures everything we need in one expression, perhaps making it possible to verify it all at the same time. What would the generalization look like? In order to capture all the things that I should be properly connected to if I am going to rule out being systematically deceived I must say that, modulo correctible errors—false beliefs which observations potentially could correct—things are pretty much as they appear to me to be not just at this moment, but also according to the general assumptions that the perceptual process typically has me making, such as that objects do not disappear in virtue of my turning away, and so on. Thus, that there is a building holding up my office counts as part of how things appear to me to be in this sense of “appear.”

But now we have come full circle. In order to get premises strong enough to imply the conclusion that I am not a *brain in a vat*, we have had to add so much information to the premises, and in such a generalized form, that if we know those premises, then there can be no surprise that we also know we are not *brains in vats*, for what is left for us to be systematically deceived about? We may be wrong about many things, even ordinary things, but only in the normal way of being wrong, not uncorrectably so. We have closed the implicational gap, but only by inflating the premises to the point of recognition. Alternatively, we might think that the strengthened premises of the argument are far too much for us to know in any obvious way. In that case they also give us no reason to think we know we are not *brains in vats*, but that does not give us a *modus tollens* since there is no assumed obvious knowledge for it to undermine.

It still may seem that we have something to worry about, in that we have exposed that we may not know that we are not *brains in vats*. Sure, we know we have hands, but what we see now is that even if we assume closure this does not mean we know we are not *brains in vats*, because that does not follow. We do, surely, go around implicitly believing we are not so thoroughly deceived, though, so if we cannot defend that claim there still seems to be a problem. Part of the reason for this worry is not yet having fully taken on board the claim of this paper. Lack of knowledge that you are not a *brain in a vat* undermines your claim to knowledge only of those things inconsistent with your being a *brain in a vat*. A given list of beliefs about things around us being thus and so, and even our being rightly hooked up for knowing

that they are thus and so, is obviously not inconsistent with being a *brain in a vat*. This may seem like a bad thing—all of the things we are most confident we know will never get us to the reassuring knowledge we are not otherwise systematically deceived about many, many things. But it is just as much a good thing: we do not need to know we are not *brains in vats* in order to know however long a list of the familiar things we think we know. For all the skeptic has done, we can take the skeptic's first premise—you know that you have hands—and go home with it. We can take our feet home too, and keep assuming we know the closet door does not open into outer space. Nothing in his subsequent argument touches what we are permitted to think we know of such things.

The kicker, one might think, is in those assumptions that perception has us automatically making, such as that objects remain when I am not looking at them. Such claims are generalizations and so not claims I can directly verify in the way discussed above, yet we believe them and think we know them. However, granting that we think we know such generalizations, and granting that we cannot verify them directly, this still does not pose a problem. The skeptic has not shown that direct verification is necessary for knowledge. We think of this standard because the skeptic focused on an example, the claim that we have hands, where we fulfill it, and he focused on this because direct verification seems of all things overwhelmingly sufficient for knowledge, and he needed a premise we seem very obviously to know. We did not need to assume that direct verification is necessary for knowledge in order to take the skeptic seriously in the first place, and his argument leaves the question whether we know these generalizations just as it was found.

The effect of the argument of this paper somewhat resembles the outcome of views of knowledge that deny closure. In both you have a split decision where it is possible for you to know you have hands without knowing you are not a *brain in a vat*. But here the reason for the split is that it is possible to be a *brain in a vat* even if you have hands. The difference is in whether we deny that knowing you have a hand, and knowing that your having a hand implies that you are not a *brain in a vat*, implies that you know you are not a brain in a vat (closure); or deny that "I have a hand," and claims relevantly like it, imply "I am not a *brain in a vat*." There is no need to deny closure in order to defeat the skeptic in the way advocated here. There is no need to deny any general principle about knowledge, as far as I can see. Here, we got generality over the moves the skeptic might make to repair his situation by explaining the trade-off he will always face in trying to identify both a logical implication and

a huge intuitive knowledge gap. The skeptic needs a conclusion strong enough to be obviously unknown by us, and weak enough to follow from something we obviously know. His problem is that the closer we get to an implication, the farther we get from this intuitive combination.

The argument of this paper clearly does not appeal to a denial of closure, but one might think it tends to suggest the opposite, closure, and even, perhaps, to depend on it. This is because it is sufficient for a counterexample to closure if we find a case where we obviously know something, obviously know that it implies something else, and obviously do not know the something else. If I am right that the skeptic cannot find the kind of example he needs then it looks like a counterexample to closure cannot be found either. This is not quite right. There may be counterexamples to closure that do not aid the *brain-in-a-vat* skeptic. For example, one might think that knowing that one has a hand does not give one knowledge that it is not a fake hand although the first implies the second. However, this will not help a *brain-in-a-vat* skeptic for if this is a failure of closure then one does, or can, know one has a hand despite not knowing that it is not a fake hand, and thus that one is not a *brain in a vat*. Thus, one does not get to do a modus tollens to undermine the claim to knowledge of a hand.<sup>2</sup> *Actual* counterexamples to closure do not undermine my argument. However, all of this is very confusing. The approach to defeating skepticism that denies closure<sup>3</sup> assumes the skeptic needs closure in order to make his argument go. The argument of this paper suggests that what the skeptic needs is a counterexample to closure. How could both of these be true of his one argument?

<sup>2</sup>The same goes for the potential counterexample from "That is a zebra" to "That is not a cleverly disguised mule." If knowledge is not closed then not knowing whether there is a subterfuge does not undermine your knowledge that it is a zebra. Of course, one might think of these hand and zebra examples as helpful to the skeptic because if one maintains closure then the apparent fact that we do not know the conclusions comes back to undermine our knowledge of the premises. My response to this is to ask exactly what kind of hand it was you thought you knew about when you claimed that you knew you had a hand. It seems that would have been a real hand, in which case why exactly do you not know it is not fake? The other possibility is that you were claiming you knew you had an at least floppy hand, in the sense above, in which case there is no reason to expect you would know it was not fake, and so no modus tollens. A similar point can be made about zebras and mules.

<sup>3</sup>See Fred Dretske's "Epistemic Operators," this JOURNAL, LXVII, 24 (December 1970): 1007–23, "Conclusive Reasons," *Australasian Journal of Philosophy*, XLIX, 1 (May 1971): 1–22, and *Knowledge and the Flow of Information* (Cambridge: MIT, 1981), and Robert Nozick's *Philosophical Explanations* (Cambridge: Harvard, 1981).

Closure and closure failure are both at work because the skeptic's argument is a *reductio ad absurdum*, and the way down is different from the way up. The way down appeals to an implication claim, and two intuitions:

- |   |  |
|---|--|
| (1) I know "I have hands."  | Intuition                                    |
| (2) "I have hands" implies "I am not a <i>brain in a vat</i> ."         | "Logic"                                      |
| (3) A normal person and a <i>brain in a vat</i> have the same evidence. | definition, stipulation, or...? <sup>4</sup> |
| (4) I (obviously) do not know "I am not a <i>brain in a vat</i> ."      | (3) plus (independent) intuition             |

If all of these statements are true, then this is a counterexample to closure, in which case the skeptic fails to undermine our knowledge of our hands. This is not what the skeptic wants, but he does need to make these four statements all *look* true. He needs to produce an apparent counterexample to closure on the way down, but one that does not actually disturb your conviction that knowledge is closed, so that your only option is to do a *modus tollens* and lose confidence in your ordinary beliefs. Another way out would be Moore's dogmatic one, of course, insisting that because one knows one has a hand, and because knowledge is closed, therefore one does know that one is not a *brain in a vat*, but the skeptic hopes you find that laughable, and if my argument is right then Moore's argument also has a false implication claim.

The option of denying closure admits that (1)–(4) are all quite convincing and denies the skeptic the move from "I do not know 'I am not a *brain in a vat*'" to "I do not know 'I have hands'." In other words, it denies the way up. What is distinctive about the argument here is that I am denying the skeptic the way *down*, via an argument that his apparent counterexamples to closure are illusory, and intuitions to the contrary are due to sloppiness about implication. Once we see what we need for his implication claims we see by inspection, independently, that the premises we clearly know or do not know line up only with conclusions we clearly know or do not know, respectively; the skeptic has not created a problem or a reason to deny closure, because he has a problem defending (1), (2), and (4) simultaneously.

<sup>4</sup>In *Knowledge and Its Limits* (New York: Oxford, 2000), Timothy Williamson points out that the skeptic's argument is not serious if he plans simply to stipulate that the brain in a vat and I could have the same evidence. "Same evidence" must be defined, and the claim that it is possible defended, which Williamson argues cannot be done. Whether or not this challenge can be met, I grant the skeptic's assumption here for the sake of argument.

The argument here proceeds without Moorean-style dogmatism, for I do not claim that we do know we have hands or that we do know that we are not *brains in vats*, only that the skeptic has given us no reason to think we do not. I also differ from Moore in refusing the claim that knowledge of something momentous—that I am in large part undeceived about my body and the world—follows from knowledge of something skimpy, such as that I have one or two real hands. The knowledge does not follow because the thing itself does not follow. Too little attention had been paid to how much is necessary for an implication claim and what the contents of the imagined premises and conclusions were or must be if the goals of the argument were to be achieved.

We can defeat the skeptic without denying closure because in his initial foray he needs to convince us of an apparent violation of it but only comes up with a case where the principle does not apply or else is not intuitively violated. Anything you know as well as that you have a hand will carry so little information that it will not imply you are systematically undeceived about much of the world. Anything that is so informative as to imply this is either something we do not plausibly—and certainly do not obviously—know, or else something the knowing of which would also make us obviously know we are largely systematically undeceived. The kind of example the skeptic needs is a will-o'-the-wisp.

None of this means that we need to worry that we are knowledge-poor. It means that the skeptic's argument has not shown anything about our knowledge. Rather, he has engaged, with our assistance, in an iterated shell game. (Ten dollars if you can tell me where the knowledge went!) Most people think, *contra* G. E. Moore, that you cannot get out of radical skepticism by waving your hands. What we have seen here is that you cannot get *to* a radically skeptical challenge by hand-waving either. These are both true for the same reason: you should not expect knowledge that you have a hand to give you knowledge of a world, not because of closure failure but because there being a hand does not imply there is a world, much less one that is like we think it is. Even the Romantics, who told us that we can see the world in a grain of sand, or the universe in a drop of water, did not think we could expect to do so by logical implication.

SHERRILYN ROUSH

University of California, Berkeley