

Closure on Skepticismⁱ

The Journal of Philosophy 107 (5): 243-256, 2010.

Sherrilyn Roush

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 weren't in this state if you knew you had hands, since having hands implies you are no brain in a vat. You obviously don't know you're not a brain in a vat, though—you have no evidence that would distinguish that state from the normal one you think you're in. Therefore, by modus tollens, you don't 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 .ⁱⁱ

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 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.

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 due to the weakness of the claim that one is not a handless brain in a vat. 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 I have hands, then in virtue of that I know I’m 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. To put my point neutrally, 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. This pattern continues. If you shake your foot and say that the kind of brain in a vat you have in mind has no feet, we can grant your point and reply that you still could be a thing that is like a brain in a vat in every respect except that it has a foot. These hands and feet do not tell us what we wanted to know if we wanted to be reassured that we are not systematically deceived, which is typically taken to be what we need. Why else be distressed? If you know you have a hand or foot then it is entirely unsurprising that you know that you are not a handless or footless brain in a vat, since in that case the latter is trivial to know.

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 non-trivial? 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 the areas of our strong 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. (Why else would the advertisers have done this?) 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 the matter. Philosophers are not immune to such unconscious mistakes; we are all apt to make them when our conscious attention is directed elsewhere.

The other 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. I never hear epistemologists say the sentence “I’m not a handless brain in a vat” with a high-pitched emphasis on “handless;” to do that 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).

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. When a qualification is added to the content of a sentence, in our cases “that I can talk to” or “handless,” we have the choice, via tone and pitch, to emphasize it, hide it, or say it straight. Epistemologists sometimes hide the trivializing effect of the “handless” qualification with the lowered pitch on “handless,” but they also sometimes say it 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.” As I argued above the conclusion that is claimed to be known is thereby trivialized.

Another plausible reason for the mistake is an equivocation on the term “brain in a vat.” One could mean by this phrase a literal, very 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. Having 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 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, one might think. 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, although, as we have seen, the task would be to avoid making it trivial. 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. For example, we could add, in addition to the claim that we have hands, the claim that we have feet, and so on for other parts of our bodies. It is no less plausible that we know these things than that we know that we have hands, in the innocent first step of the skeptic’s argument. This is far too modest a beefing up of the premise, though, as indicated above. It is hopeless, since the hands and feet do not rule out that we are brains in vats with hands and feet. We could imagine an entity like a brain in a vat in every respect except that it had hands and feet, etc., attached, and the possibility of systematic deception does not go away.

How many “attachments” can a brain have before it ceases to be a brain in a vat and becomes a normally epistemically functioning person? The number of attachments is not the issue, of course. 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. To answer the question what makes something a brain in a vat in the important sense, we have to hew closely to what is disturbing about not knowing that you are not one of these things, which is closely related to what is disturbing about being one of these things: your evidence about the external world is systematically corrupt—which derives from the fact 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. You may have hands and feet. Having hands and feet does not rule out the possibility that you are systemically deceived, that is, that 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. This claim that it seems independently obvious we do not know needs to follow from





things we think we clearly do know; otherwise our knowing the second cannot, even with closure, imply that we know the first, and our not knowing the first will not threaten our claim to knowledge of the second. Having ever so many hands and feet does not rule out the difficult to discriminate and disturbing, and hence essential, feature of the *brain-in-a-vat* hypothesis.

What *would* rule out the skeptical hypothesis? The *brain in a vat* hypothesis has the special feature that not only is it a claim about how the world is—such as that there is a human brain suspended in a vat, with all those hookups to electrical stimulations—but also implies something about the vat-brain-person’s relation to the world, and thereby to her lack of discriminating evidence about her situation vis-à-vis the world. This is the essential part of the hypothesis, so this is what needs to be ruled out, and in presenting the skeptic’s argument, people have, typically, pretty grossly underestimated what is required to do so. And the premise must rule it out, since the skeptic’s claim is one of implication, which means that the premise, e.g., some version of “I have hands,” must rule out all *logically* possible ways for me to be a *brain in a vat*.ⁱⁱⁱ That I have hands is not enough for this. But perhaps the reason it is not enough is that we have considered a premise that asserts nothing more than that there is a hand attached to the brain. It does not say that the hand is connected up to my impressions, and intentions to move, in the normal way that I have when I have evidence and a non-deceptive set of impressions of the world. Maybe this kind of claim—which we can as innocently agree we have knowledge of when the skeptic asks as we agree about the previous claim—will do the trick of implying the claim that I am not a *brain in a vat*.

Call the first type of hand that is unconnected to my impressions a “floppy” hand, and the second a “hooked-up” hand. The poor people in the Matrix have real, physical hands, but in the sense just introduced the hands are floppy. However, 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 movement of the arms, but they have those too. 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 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 and all-surrounding. 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's 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 is 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.

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. No luck here, though. We can concoct 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). 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 *brains 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. Thwarted again, 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 I am perfectly engrossed in a video game on the screen in front of me. The table impressions are properly produced, which requires 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 video 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 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. Both aspects will emerge in what follows, but 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 the full denial of the *brain in a vat* thesis will be implied. In strengthening the premises in this way we want to include enough hooked-up objects to insure that I am not systematically deceived about 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? One problem is that there are a whole lot of things left out 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. They do not imply the denial of the *brain in a vat* hypothesis (systematic deception) because the list you make will always leave out important aspects of the world no matter how many you have already put in. Moreover, the things you have listed do not appear to imply anything that it would be surprising to think you know on the assumption that you know *them*. So even if the skeptic is not attached to the *brain in a vat* hypothesis but is willing to run the argument with any conclusion we obviously do not know derived from a premise we obviously do, this will not achieve what he needs.

The list is of course infinite of the things that would need to be part of a set of premises that implies the denial that one is a *brain in a vat*. However, given infinite time one could verify each claim in the list, the way one does with the claim that one has a hand, by directly inspecting each of these things 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't we have that knowledge in a different way? Not if we are trying to help out the skeptic, whose target is those of our beliefs that we think we most obviously do know. We need to make the premises that go into the skeptical argument claims that it is 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 the list off to a large finite set, because we cannot sufficiently directly verify them all at the same time. To play a role as premises to an argument, the claims have to be asserted, and in this argument apparently known in a special way, all in the same time interval.

The things that we obviously, simultaneously, 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, by simply making a generalization that includes all of those examples without listing them individually. That captures everything we need in one expression, perhaps making it possible to verify it all at the same time, and then we are done. 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, the

claim must be something very general, not referring merely to a list, however long, of specific objects. That is, I must say that, modulo correctible errors—false beliefs where there are nevertheless potential observations I could make to correct them—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, etc. 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 have closed the implicational gap, but only by inflating the premises to the point of recognition. Alternatively, we might think that strengthening the premises of the argument in this extreme way means we *don't* know them. In that case there is also no reason to think we know we are not brains in vats, but the modus tollens we could do on that does not have any assumed knowledge to call into doubt.

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. So, knowing I have a hand does not mean I should know that I am not a *brain in a vat*. Don't I still know, somehow, that I am not systematically deceived? Doesn't not knowing this still somehow interfere with the assumptions of my daily life? 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 systematically deceived. 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. We can take the skeptic's first premise: you know 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 a different set of things we believe, and think we know, 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. 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 where we fulfill it: the claim that we have hands, 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 we found it.

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 hands. The difference is in whether we deny that knowing p, and knowing that p implies q, implies that you know q (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. We need not attribute to the skeptic a false claim about knowledge, since he has a much more basic false factual premise about an implication. The contents “I have a hand” and “I am not a *brain in a vat*,” simply are not suitable for applying principles that will get him to his goal. Here we get 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. His problem is that the closer we get to an implication, the farther we get from a combination of a premise we think we know and a conclusion we think we do not know.

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. The argument I have made says that the closer he gets to intuitively known premises, the farther he gets from intuitively unknown conclusions, in cases where the implication does truly hold. But all of this is very confusing. The approach to defeating skepticism that denies closure (e.g. Dretske 1970, 1971, 1981; Nozick 1981) assumes the skeptic needs closure in order to make his argument go. The argument here suggests that what the skeptic needs is a *counterexample* to closure. How could both of these be true of his one argument?

Closure and closure failure are both at work because the skeptic’s argument is a reduction 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 brain in a vat.” | “Logic” |
| 3) A normal person and a brain in a vat have the same evidence | definition, stipulation, or? ^{iv} |
| 4) I (obviously) don’t know “I am not a brain in a vat.” | 3) plus (independent) intuition |

If all of these statements are true, then this is a counterexample to closure. The skeptic needs to produce an apparent counterexample to closure on the way down, in order to convince you that your only

options are to deny closure or, assuming closure, do a modus tollens and lose confidence in your ordinary beliefs. Another plausible 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 has a false implication claim). The option of denying closure denies the skeptic the move from "I don't know 'I am not a brain in a vat'" to "I don't know 'I have hands,'" in other words, denies the way up.

What is distinctive about my argument here is that I am denying the skeptic the way *down*, via an argument that, as it happens, also *defends* closure. I am doing so without Moorean dogmatism, for I do not claim that we do know we have hands or we do know that we are not brains in vats, only that the skeptic has given us no reason at all to think we do not. This is because the implication claims on the way down that have historically been taken to be intuitive are either false or not achievable consistently with the wide intuitive gap that the skeptic needs between the plausibility of our knowing the premise and the plausibility of our knowing the conclusion. This I have argued possible implication-patch by possible implication-patch, rather than via an appeal to or direct argument for closure. And I have argued it by focusing attention on the neglected aspect of how much it really takes to defend a claim of implication. That eliminative argument suggests an idea for further consideration in the wider discussion of closure, that when we think we have found apparent counterexamples to closure, it may be because we are not paying close enough attention to what it takes to defend an implication claim and to the consequences for intuitions about knowledge of the premises and the conclusion of doing the patching. A further apparent consequence of the argument of this paper is that our confidence in our ordinary knowledge is safe from the brains-in vats skeptic *whether or not* knowledge is closed under known implication.

We can defeat the skeptic without denying closure because in his initial foray he needs to be convincing us of an apparent violation of it, but can only come up with a case where the principle simply does not apply or else is not intuitively violated. Anything you know as well as that you have a hand is going to carry so little information that it is not going to imply that you are not systematically deceived about most 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 obviously make us also know we are not systematically deceived. 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 can't 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, and for the same reason: knowledge that you have a hand cannot give you knowledge of a world, because having a hand does not imply there is a world. 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 that by logical implication.

References

Dretske, Fred (1970), "Epistemic Operators," *Journal of Philosophy* **67**: 1007-23.
----- (1971), "Conclusive Reasons," *Australasian Journal of Philosophy* **49**: 1-22.
----- (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
Greco, John (2007), "External World Skepticism," *Philosophy Compass* **2**: 625-649.
Nozick, Robert (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

ⁱ Thanks to John MacFarlane for helpful discussion.

ⁱⁱ For an up-to-date discussion of this argument, see Greco 2007.

ⁱⁱⁱ It does seem possible to re-run the skeptic's argument using an inductive inference, though I am not aware that this has been done. In this version, we would say that "I have hands" makes it probable that I am not a brain in a vat, and improbable that I am. However, we can look at this in terms of the same dilemma I am developing above. It does not seem that the mere fact that I have hands makes it improbable that I am a *brain in a vat*. It's not just that it is logically possible for me to be a *brain in a vat* with hands. It is that this seems at best marginally less probable than that I am a *brain in a vat* was in the first place. If I didn't know I was not a *brain in a vat* before, adding the hand does very little to change the probability. The skeptic is stymied on this side because whatever probability my having a hand confers on the claim that I'm not a *brain in a vat* that does not give me any more knowledge than it is plausible to think I have.

^{iv} Williamson (2000) points out that the skeptic's argument is not serious if he plans simply to stipulate that the brain in 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. I am inclined to think this challenge can be met, and in any case I grant it here for the sake of the argument.