CARTESIAN SKEPTICISM, CONTENT EXTERNALISM, AND SELF-KNOWLEDGE

Anthony Brueckner*

SÍNTESE – Há um argumento cético clássico derivado das *Meditações sobre a filosofia primeira*. Este artigo oferece uma formulação contemporânea padrão do argumento, pretendendo mostrar que ninguém sabe qualquer coisa sobre o mundo extramental. A obra de Hilary Putnam na filosofia da linguagem e da mente parece fornecer uma resposta a uma versão atualizada do argumento cético cartesiano. Em sua maior parte, este artigo é dedicado a uma análise e crítica das meditações anti-céticas de Putnam.

PALAVRAS-CHAVE – Descartes. Putnam. Ceticismo. Cérebros em cubas. Externalismo de conteúdo.

ABSTRACT – There is a classical skeptical argument that derives from Descartes's *Meditations on first Philosophy*. This paper offers a standard contemporary formulation of the argument, which purports to show that no one knows anything about the world that exists outside our minds. The work of Hilary Putnam in the philosophy of language and mind seems to afford an answer to an updated version of the Cartesian skeptical argument. The bulk of this paper is devoted to an analysis and critique of Putnam's anti-skeptical meditations.

KEY WORDS – Descartes, Putnam, Skepticism, Brains in vats, Content externalism.

1   The Cartesian Skeptical Argument

Let us start by considering an updated version of the skeptical concerns that Descartes raised in the first Meditation.[1] Imagine a living brain floating in a vat of nutrient fluids while attached to a supercomputer. The computer is programmed to provide electrical stimuli to the brain that are just like those a normal brain receives due to input to the sense-organs connected to the normal brain. As a result, the brain in a vat is caused to have experiences just like those of a normal human. The brain in a vat, we are to imagine, is the subject of conscious mental life just like that of a normal human: experience, thought, belief, desire, intention, and so on. Upon having various experiences, the brain in a vat thinks, 'The waves are breaking nicely here on Sands Beach in Santa Barbara'. Imagine that the brain

---

*   Departament of Philosophy. University of California, Santa Barbara.
[1]   See *Meditations on First Philosophy*, trans. and ed. by J. Cottingham (Cambridge: Cambridge University Press, 1996).

in a vat is a mental twin of you: whenever you have an experience of seeing sand and surf, the brain in a vat has an exactly similar experience. When you think, 'The sand is hot, and the surf is big', the brain thinks, 'The sand is hot, and the surf is big'. According to the Cartesian philosopher, the brain in a vat is massively mistaken about his world. He believes that he is sitting on a beach, for example, even though he is not, and cannot, sit anywhere.

You might ask: in the imagined vat scenario, what is the explanation for the envatted brain's predicament? On one version of the thought experiment, we are to imagine a normal human who is kidnapped by evil neuroscientists who relieve him of his brain in order to test the deceptive powers of their supercomputer's brain-manipulating program. On a different version of the thought experiment (due to Hilary Putnam), we are to imagine that the brain has always been envatted; there are no evil neuroscientists; there are no trees or beaches in the vat world; there is nothing but the brain in its vat in a room with the supercomputer; and this arrangement has come about as a result of some cosmic accident. We will call this special Putnamian brain in a vat a *BIV*.[2] Let *SK* be the proposition that I am a BIV.

Now let us depart the arid vat world just considered and dally in an English garden. I say to my friend, 'That is a goldfinch in the tree'. My friend asks, 'Old boy, do you *know* that it is a goldfinch?' I reply, 'Yes, I know that it is a goldfinch because it has those markings'. He retorts, 'Do you realize that goldcrests also have those markings? How do you know, then, that it is not a goldcrest? How can you rule out that possibility on the basis of your evidence?'. I then admit, 'I do not know that it is not a goldcrest--I cannot rule out the goldcrest possibility. But I know that it is a goldfinch'. At this point, my friend throws up his hands and says, 'Look, old chum, if you *do* know that it is a goldfinch, and you realize that if it is a goldfinch, then it cannot be a goldcrest, then you must also know that it is *not* a goldcrest. Since you have admitted that your evidence does *not* enable you to know that it is *not* a goldcrest, you surely must in all fairness retract your claim to know that it is a goldfinch'.[3]

In this exchange, my friend's exasperation at my pigheadedness stems from his acceptance of the principle that *knowledge is closed under known logical implication*. Letting 'K(S,P)' stand for 'S knows that P', and letting

'→' stand for logical implication, we can formulate our *Closure Principle* as follows:

For all S,P,Q: {K(S,P) & K(S,[P → Q])} → K(S,Q)

---

[2]   See *Reason, Truth and History* (Cambridge: Cambridge University Press, 1991). In Descartes' first Meditation, the skeptical scenario involves my mental twin and a God-like Evil Genius in a world containing no physical objects. The immaterial Evil Genius is the cause of my experiences, playing a role in the thought experiment that parallels that of the supercomputer in the BIV scenario.

[3]   This dialogue is inspired by J. L. Austin's "Other Minds", in his *Philosophical Papers* (Oxford: Oxford University Press, 1961).

Let us now return to the proposition SK = I am a BIV. Suppose SK were true. By construction of the BIV scenario, my experiences would be no different in an SK world than they would be in a normal world. So the experiences I in fact have do not discriminate between an SK world, in which I am a BIV, and a normal world, in which I am embodied and have experiences caused by beaches, surf, and sand. These considerations lead the Cartesian skeptic to maintain that my sensory evidence does not justify me in believing ~SK = I am *not* a BIV. Since knowledge requires justification, it follows that I do not know ~SK.

We can now construct a Cartesian skeptical argument to show that I do not know, say, H = I have hands. Let 'AB' denote me.

(1) {K(AB,H) & K(AB,[H → ~SK])} → K(~SK)

(2) K(AB,[H → ~SK])

(3) ~K(AB, ~SK)

So, (4) ~K(AB,H)

Premise (1) is an instance of the Closure Principle. (2) simply expresses that fact that I know the logical implication in question. (3) is based upon the skeptic's claims about the indiscriminability of SK worlds and normal worlds. (4) follows from (1)-(3) by Modus Tollens and DeMorgan's Laws. The skeptic seems to have proved that I do not know that I have hands! He will invite me to run through a parallel argument for my claim to know that I see surf and sand. And so on.

2   **Semantic Externalism and Content Externalism**

Let us now consider a response to the skeptical argument that employs some ideas that derive from the work of Hilary Putnam.[4] Imagine that a Martian looks at a paint blob that causes him to form a mental image that exactly resembles my mental image of a palm tree. Imagine that the Martian has had no causal contact with trees and no causal contact with anyone who has had causal contact with trees. *My* image is an image *of* a palm tree; it *represents* and *refers to* palm trees. But even though the Martian's image exactly resembles mine, his image does not have the foregoing *semantic* features of aboutness, representation, and reference. Putnam's intuition is that there is a causal constraint on such semantic features that is not satisfied by the Martian's image.

We are assuming that the BIV has a language that he uses in thinking, believing, wondering, and so on. The BIV never speaks or writes, of course, but he has a rich mental life that involves use of his language. The BIV's language is

---

4 See again *Reason, Truth and History*. See also my "Brains in a Vat", *Journal of Philosophy* 83 (1986), 148-67, "Semantic Answers to Skepticism", reprinted in *Skepticism: a Contemporary Reader*, ed. by Keith DeRose and Ted A. Warfield (Oxford: Oxford University Press, 1999), and "Trees, Computer Program Features, and Skeptical Hypotheses", in *The Skeptics: Contemporary Essays*, ed. by Steven Luper (Aldershot: Ashgate, 2003).

superficially just like English (recall that we imagined him to be my mental twin, though, as we will shortly see, this cannot be quite right). For example, whenever I think a thought using the sentence 'That is a palm tree', my BIV twin thinks a thought using that sentence. But what do the BIV's words 'palm tree' refer to? Given the foregoing considerations concerning the Martian's image, Putnam maintains that the BIV's 'palm tree' does not refer to palm trees, just as the Martian's image does not refer to or represent palm trees. This is because the causal constraint on reference is not satisfied vis-a-vis the BIV's words 'palm tree' and palm trees. That is, there are no palm trees in the vat world to serve as referents for his words 'palm tree'; the BIV never has any causal contact – direct or indirect – with palm trees, and thus his words do not refer to palm trees. What do those words refer to, then? One answer that is suggested by Putnam is as follows. Find the entities in the vat-world that play a causal role vis-a-vis the BIV's uses of 'palm tree' that is analogous to the causal role that palm trees play vis-a-vis a normal English speaker's uses of 'palm tree'. The best candidates for such entities are the recurring computer program features which cause electrical stimuli in the brain, which stimuli in turn produce experiences just like those which are produced in normal humans as a result of seeing palm trees. Call these *PT program features*.

So on this view, the BIV's words 'palm tree' differ in their reference from a normal human's words 'palm tree'. So there must also be a difference in meaning between the BIV's words and a normal human's words, given that the meaning of a term determines its reference. These semantic differences will also induce a difference between the truth conditions of a BIV's sentence 'A palm tree is near' and those of a normal human's corresponding sentence. The normal human's sentence 'A palm tree is near' is true if and only if a palm tree is near to the human, whereas the BIV's sentence 'A palm tree is near' is true if and only if PT program features are running. The semantic differences between the BIV's language and that of a normal human are ultimately induced by the differences between the two creatures' external, causal environments. Let us call this Putnamian view *semantic externalism*.

Semantic externalism engenders some interesting results concerning the mental states of the BIV. Since the BIV's words 'palm tree' do not refer to palm trees and differ in meaning from a normal human's corresponding words, and since his sentence 'A palm tree is near' does not have the same truth conditions as a normal human's corresponding sentence, we cannot say that when the BIV thinks a thought using that sentence, he is thinking a thought with the *same content* as a normal human's thought *that a palm tree is near*. The semantic differences we have noted induce differences in the contents of the mental states had by the BIV and a normal human. The BIV cannot think that a palm tree is near, nor can he think that there are big waves at Sands Beach. The BIV's corresponding thoughts are about the entities in his external, causal environment which systematically produce his experiences, viz. recurring computer program features such as the PT program features. So the contents of the BIV's thoughts, beliefs, desires, and wonderings depend upon his external, causal environment. Let us call this view *content externalism*.

Now we see that the BIV cannot be the mental twin of a normal human. If God were to somehow "look inside" the minds of the BIV and his normal human counterpart, then He could not distinguish the two creatures' mental states purely on the basis of their inner phenomenology. But the states would nevertheless differ in content in virtue of the differences between the creatures' external, causal environments. Now we also see that one of our key initial Cartesian claims about the BIV thought experiment was not correct. That is, we characterized the BIV as being massively mistaken about his world. But that claim presupposed that when the BIV has a belief that he expresses using his sentence 'A palm tree is near', the content of the belief in question is *that a palm tree is near*. If the BIV had that belief, then he would be mistaken, since no palm tree is near to him. But according to semantic externalism, the BIV has no such belief. Instead, he has a belief about computer program features – a belief which is *not* mistaken, given that the appropriate program features are running. According to the content externalist, the Cartesian philosopher is mistaken in the following assumption: in the BIV thought experiment, we can *hold fixed* the mental life of a normal human while varying, in thought, his external environment (normal world to vat world), with the result that the BIV's beliefs are those of a normal human and hence massively mistaken.

3   The Self-Knowledge Quiz and the Transcendental Argument

Let us pause for a bit and cleanse our minds of all concerns regarding Cartesian skepticism, semantic externalism, and content externalism. Let us imagine that I am looking out my window with a puzzled expression. My friend asks me, 'What are you thinking?' I say, 'Oh, I'm thinking that some palm trees have palm rats. I had forgotten, but now I see a rat in that palm tree'. My friend says, 'But do you *know* that you are thinking that some palm trees have palm rats? Maybe you are thinking that quantum entanglement is amusing'. I reply, 'Are you crazy? Of course I know that I am thinking that some palm trees have palm rats. Maybe that thing is not a palm rat, and maybe there are no such things as palm rats. Who knows – maybe there are no such things as palm trees at all. Maybe some elaborate trick has been played on me all these years. But at least I know what I am thinking right now, and it's this: that some palm trees have palm rats'.

In this dialogue, I am expressing the natural view that I have easy, direct, unproblematic access to my own current mental states. Unlike my knowledge of others' mental states, my knowledge of my own mental states is almost never evidentially based. Further, in the vast majority of cases, when I believe that I am in a certain mental state, my belief is correct. Such accuracy is obviously lacking in my beliefs about the mental states of others.

Let us now return to the Cartesian skeptical argument formulated earlier. Its third premise is

(3)   ~K(AB, ~SK)

viz. I do not know that I am not a BIV. We will now construct an argument for the conclusion that ~SK – that I am not a BIV. If the argument is successful, then I can claim to *know* ~SK on the basis of knowing the premises of the argument and knowing that they logically imply the conclusion, ~SK. By this method, we will block the Cartesian skeptical argument by showing that its premise (3) is false. In order for this strategy to work, we must make a case for the claim that my knowledge of the constructed anti-skeptical argument's premises does not somehow rest upon or presuppose knowledge of propositions about the external world. As we will see, it is a rather delicate matter to decide whether our constructed argument satisfies this criterion of adequacy. We are trying to construct a sort of *transcendental argument* for ~SK: an argument whose premises are (i) thin enough to be allowed in play by the skeptic(since they do not somehow presuppose knowledge about the external world), but, at the same time, (ii) strong enough to establish the conclusion ~SK.

The first premise of our transcendental argument is justified by considerations that flow from content externalism. In general, if something is a BIV, then it cannot think that a palm tree is near. As we saw earlier, no BIV thinks such a thought when it thinks a thought using the sentence 'A palm tree is near'. I know this general proposition about BIVs on the basis of knowing the truth of content externalism, and I know the truth of that doctrine via philosophical reflection. This puts me in a position to know that, in particular, if *I* am a BIV, then I am not thinking that a palm tree is near.

Given my easy, direct, unproblematic knowledge of my own mind and its current mental states, I am now in a position to answer the following Self-Knowledge Quiz:

Q: What are you thinking ?
A: (a) That quantum entanglement is amusing.
   (b) That cows are purple.
   (c) That North America is south of South America.
   (d) That a palm tree is near.

The answer is (d), and I know that the answer is (d). This is a quiz that is very hard for me to fail.

Given these reflections, I can now construct the desired transcendental argument:

   (A) If I am a BIV, then I am not thinking that a palm tree is near. (by Content Externalism)
   (B) I am thinking that a palm tree is near. (by Self – Knowledge Quiz)
So (C) I am not a BIV.

In symbols, (C) is ~*SK*. Given my knowledge of the premises (A) and (B), and given my knowledge of the validity of the argument, I know the conclusion. So we have K(AB,~SK). So the skeptical argument is stymied, since its premise (3) – ~K(AB,~SK) – is false.

58

One way for the skeptic to respond to the argument is to reject content externalism. I will not pursue this line of response, for two reasons. First, content externalism is extremely plausible and is widely accepted: if a BIV has any thoughts at all, then they do not have contents that concern absent trees and beaches. Second, I want to see whether content externalism can indeed ground a successful transcendental argument for ~SK. So I would like to envision a situation in which the Cartesian skeptic simply grants the very plausible content externalist doctrine for the sake of argument and then see whether the transcendental argument is indeed successful.

4 Question-Begging and Epistemic Circularity

Given this dialectical situation, it is tempting to regard the transcendental argument as somehow question-begging. Charges of question-begging in philosophy are often very hard to adjudicate. The present case, as we will see, is a case in point. The argument is obviously not formally circular, in that the conclusion does not appear as a premise. Nevertheless, here is one way that the skeptic might try to substantiate the charge of question-begging. The skeptic: 'I am granting the content externalist premise (A) for the sake of argument. I want to see just how far you, the anti-skeptic, can go when fueled by the doctrine of content externalism. However, I am now suspicious about premise (B). We are obviously in dispute about the conclusion (C): you say you know it, I say you don't – -you say that you know that you are not a BIV, I say that you don't know this. But given that I have granted you (A), it turns out that whether or not you are now thinking that a palm tree is near depends upon whether or not you are a BIV. But whether or not you are a BIV is in dispute. So I will not accept your premise (B), according to which you are now thinking that a palm tree is near.'

This is a bad objection. The same "question-begging" objection can be made to *any* modus tollens argument where two disputants disagree about the conclusion while accepting the conditional premise. The disputant who wants to reject the second premise must consider what sort of case can be be made for the premise while prescinding from the fact that given the conditional premise, the second premise will yield the disputed conclusion.

At this point it might be useful to discuss the notion of *epistemic circularity*. Let us say that an argument is epistemically circular just in case having justification for accepting one or more premises depends upon having justification for accepting the conclusion. Consider this *track record argument*:

Perceptual belief B1 was true.
Perceptual belief B2 was true.
  •
  •
  •
Perceptual belief Bn was true.
So, Perception is a reliable belief-forming process (in virtue of its track record).

In order to have justification for accepting any given premise of this argument, one needs to have justification for accepting the conclusion about the reliability of perception. Otherwise, how could one have justification for believing that, say, perceptual belief B1 was *true*? Such an epistemically circular argument suffers from the following deficiency: if I start the day lacking justification for believing the track record argument's conclusion, then the argument lacks the power to engender such justification by consideration of its premises.

It might be thought that our transcendental argument is epistemically circular. We are supposing that the skeptic grants for the sake of argument that I know premise (A): so I know that if I *am* a BIV, then I am *not* thinking that a palm tree is near. Am I justified in accepting premise (B): am I justified in believing that I am thinking that a palm tree is near? It might seem that in order to have the justification in question, I must have justification for believing that I am *not* a BIV. For if I *am* a BIV, then I am *not* thinking the thought in question, and premise (B) is false.

However, this is simply a variation on the theme of the foregoing *bad* question-begging charge. The structure of the current objection is as follows. 'Consider the argument: If P, then Q; ~Q; so ~P. Suppose that I know the first premise: if P, then Q. Am I justified in accepting the second premise, ~Q? In order to have justification for believing ~Q, I must have justification for believing the conclusion, ~P. For if I had justification for *P* instead (in addition to having justification for the conditional, as we have assumed), then I'd have justification for believing that Q, rather than ~Q, is also the case. So the argument is epistemically circular, just like the track record argument'. But, again, this objection can be made against *any* Modus Tollens argument whose conditional premise is not in dispute. Surely not all such arguments are epistemically circular.[5]

It could be that the following reflections are driving the current charge of epistemic circularity. Suppose that we accept a closure principle for justification that is modeled on the closure principle for knowledge. Let 'J(S,P)' stand for 'S has justification for believing P':

For all S,P,Q: $\{J(S,P) \ \& \ J(S,[P \rightarrow Q])\} \rightarrow J(S,Q)$

[5] Claudio de Almeida suggested that, unless one is ready to argue directly against the argument in question, one may have no good reason to deny that all Modus Tollens arguments are indeed epistemically circular (at least when put forward in the above dialectical context). He bases his suggestion on the claim that the inference rule Simplification is epistemically circular as well, since it would seem that I must, e.g., have antecedent justification for the conclusion P in order to have justification for the premise P&Q. But even if all Modus Tollens arguments (when put forward in the above dialectical context) are epistemically circular, the transcendental argument under discussion in the text is no worse off than countless philosophical arguments and accordingly should not be shunned.

If I have justification for the conditional premise in the foregoing Modus Tollens argument (which I do, since, we assume, I know the premise), and I have justification for believing the second premise, ~Q, then by Closure it *follows* that I have justification for believing the conclusion, ~P. But given Closure, this result holds for *all* modus tollens arguments. Again, surely not all such arguments are epistemically circular.

Let us consider another way of trying to make a circularity, or question-begging, charge stick. Let us consider an adapted version of some reasoning given by Duncan Pritchard and Jesper Kallestrup in a recent paper.[6] As we have noted earlier, our criterion of adequacy requires that the premises of our transcendental argument must be knowable without reliance upon empirical knowledge about the external world. Otherwise, the premises will not be thin enough to be granted by the Cartesian skeptic. One way of putting this criterion of adequacy (using the terminology of Pritchard and Kallestrup) is this: the truth of premises of the transcendental argument must be *reflectively determinable*, in the sense that they must be knowable without reliance upon empirical knowledge. Is the truth of premise (B) reflectively determinable? That is, can I reflectively determine that I am thinking that a palm tree is near? Pritchard and Kallestrup argue in the negative as follows. If I am thinking that a palm tree is near, then I am not a BIV. So if I can reflectively determine that I am thinking that a palm tree is near, then I must be in a position to reflectively determine that I am not a BIV. But since my experiences do not discriminate between being a non-BIV and being a BIV, I cannot reflectively determine that I am not a BIV. Hence, I cannot reflectively determine that I am thinking that a palm tree is near.

This reasoning is flawed in several ways. It seems to presuppose that reflectively determining the truth of our argument's conclusion is somehow a *precondition* for reflectively determining the truth of premise (B) (since given (A), the conclusion follows). But the situation is otherwise. I will come to know the conclusion–-that I am not a BIV–-as a result of reflectively determining the truth of the premises and then deriving the conclusion. Can I reflectively determine the truth of premise (B)? It seems that I can. Otherwise the Self-Knowledge Quiz would be a lot harder than it is.

Let us take up another question raised by the Pritchard-Kallestrup reasoning. Can I reflectively determine that I am not a BIV? As they point out, my experiences do not discriminate normal worlds from vat worlds. So I cannot reflectively determine that I am not a BIV on the basis of my experiences. This is the idea behind premise (3) of the Cartesian skeptical argument: ~K(AB,~SK). So my experiences do not enable me to reflectively determine that I am not a BIV. But this is all right if, as I have claimed, the truth of our transcendental argument's premise can be reflectively determined. Once that is achieved, I can reflectively determine that I am not a BIV by deducing this conclusion from the reflectively determined premises.

---

[6] See "An Argument for the Inconsistency of Content Externalism and Epistemic Internalism", *Philosophia* 31 (2004), 345-54.

5 Indiscriminability

Suppose the skeptic gives up on the evanescent charge of question-begging and instead focuses more intensely upon the indiscriminability of vat worlds and normal worlds. My experiences do not enable me to discriminate between the external reality of a vat world and the external reality of a normal world. Similarly, the skeptic may suggest, nothing reflectively available to me enables me to discriminate between the divergent mental aspects of vat worlds and normal worlds. That is, by content externalism, my thought contents are different in vat worlds and normal worlds: in vat worlds, I *cannot* think that a palm tree is near, while in normal worlds, I can think that thought. Nothing reflectively available to me enables me to discriminate between vat-content-worlds and normal-content-worlds. So I am not justified in believing that I am in a normal-content-world where I think that a palm tree is near. So I do not know that I am thinking that a palm tree is near.

The question of whether content externalism is compatible with easy, direct, unproblematic knowledge of one's own mental states is notoriously difficult.[7] Here I can only explore the tip of the iceberg. I will focus upon the notion of indiscriminability. It is natural to analyze indiscriminability as follows: situations X and Y are indiscriminable to me just in case it is possible for me to be in X and yet mistakenly think that I am in Y. Is it true that vat-content-worlds and normal-content-worlds are indiscriminable to me? Imagine that I am in a vat world and think a thought using the sentence

> (T) I am thinking that a palm tree is near.

In such a world, am I mistakenly thinking that my thought-content is other than it in fact is, thereby failing to discriminate my vat-content-world from a normal-content-world? No: in such a world, my thought-content is *that PT features are running,* and in thinking a thought using (T), I *correctly* self-attribute that thought-content. So the skeptic's claim of indiscriminability is unfounded, and my claim to know premise (B) of our transcendental argument stands.

6 Conclusion

If content externalism and self-knowledge are compatible, then our transcendental argument gives us a plausible response to the BIV-based Cartesian skeptical argument. That is a big 'if', but effectively resolving the question of compatibility is too difficult a task for the present occasion. We will close on a pessimistic note. It has often been noted that there are restricted brain in a vat scenarios that are immune to our transcendental argument. Consider Sam, a

---

[7] See *Externalism and Self-Knowledge*, ed. by Peter Ludlow and Norah Martin (Standford: CSLI, 1998), *New Essays on Semantic Externalism and Self-Knowledge*, ed. by Susana Nuccetelli (Cambridge: MIT Press, 2003, and Jessica Brown, *Anti-Individualism and Knowledge* (Cambridge: MIT Press, 2004).

normal human, who is captured by the evil neuroscientists discussed in passing at beginning of this paper. Sam's brain is envatted, and his new stream of experiences fits seamlessly with his previous normal life. The semantics of Sam's language remains the same after the envatment. His words 'palm tree' continue to refer to palm trees, and he retains the capacity to think that a palm tree is near. Maybe *I* am just like Sam, freshly envatted! If so, then our transcendental argument is of no use to me.[8] That argument can, surprisingly, deliver the result that I am not a Putnamian BIV who has always been in a vat in a treeless and beachless vat world. But the argument cannot deliver the result that I am not a brand new brain in a vat in an otherwise normal world.

References and Further Reading

ALSTON, W. (1989), "Epistemic Circularity", in his *Epistemic justification: Essays in the theory of knowledge* (Ithaca: Cornell University Press).

BRUECKNER, A. (1986), "Brains in a Vat", *Journal of Philosophy* 83.

———— (1990), "Skepticism about Knowledge of Content", *Mind* 99.

———— (1991), "If I Am a Brain in a Vat, then I Am Not a Brain in a Vat", *Mind* 101.

———— (1992), "Semantic Answers to Skepticism", in K. DeRose and T. A. WARFIELD (eds.), *Skepticism: A Contemporary Reader* (Oxford: Oxford University Press).

———— (1994), "Ebbs on Skepticism, Objectivity and Brains in Vats", *Pacific Philosophical Quarterly* 75.

———— (1995a), "Trying to Get Outside Your Own Skin", *Philosophical Topics* 23.

———— (1995b), "Scepticism and the Causal Theory of Reference", *Philosophical Quarterly* 45.

———— (1997), "Is Scepticism about Self-Knowledge Incoherent?", *Analysis* (1997).

———— (1999), "Transcendental Arguments from Content Externalism", in R. Stern (ed.), *Transcendental arguments: Problems and Prospects* (Oxford: Clarendon Press).

———— (2001), "*A Priori* Knowledge of the World Not Easily Available", *Philosophical Studies* 104.

———— (2003a), "The Coherence of Scepticism about Self-Knowledge", *Analysis* 63.

———— (2003b), "Trees, Computer Program Features, and Skeptical Hypotheses", in S. Luper (ed.). *The Skeptics: Contemporary Essays* (Burlington: Ashgate).

———— (2004), "Johnsen on Brains in Vats", *Philosophical Studies*.

CHRISTENSEN, D. (1993), "Skeptical Problems, Semantical Solution", *Philosophy and Phenomenological Research* 53.

DELL'UTRI, M. (1990), "Choosing Conceptions of Realism: The Case of the Brains in a Vat", *Mind* 99.

EBBS, G. (1992), "Skepticism, Objectivity and Brains in Vats", *Pacific Philosophical Quarterly* 73.

———— (1996), "Can We Take Our Own Words at Face Value?", *Philosophy and Phenomenological Research* 56.

———— (2001), "Is Skepticisism about Self-Knowledge Coherent?", *Philosophical Studies 45.*

FORBES, G. (1992) "Realism and Skepticism: Brains in a Vat Revisited", in K. DeRose and T. A. Warfield (eds.), *Skepticism: A Contemporary Reader* (Oxford: Oxford University Press).

GALLOIS, A. (1992), "Putnam, Brains in Vats, and Arguments for Scepticism", *Mind* 101.

JOHNSEN, B. (2003), "Of Brains in Vats, Whatever Brains in Vats Might Be", *Philosophical Studies* 112.

LUDLOW, P. and Martin, N. (eds.) (1998), *Externalism and Self-Knowledge* (Stanford: CSLI).

MCINTYRE, J. (1984), "Putnam's Brains", *Analysis* 44.

NOONAN, H. (1998), "Reflections on Putnam, Wright and Brains in Vats", *Analysis* 58.

NUCCETELLI, S. (ed.) (2003), *New Essays on Semantic Externalism and Self-Knowledge* (Cambridge: MIT Press).

---

[8] For early versions of this objection, see Peter Smith's "Could We Be Brains in Vats?", *Canadian Journal of Philosophy* 14 (1984); see also my "Brains in a Vat".

Putnam, H. (1992), "Brains in a Vat", in K. DeRose and T. A. Warfield (eds.), *Skepticism: A Contemporary Reader* (Oxford: Oxford University Press).

Smith, P. (1984), "Could We Be Brains in a Vat?", *Canadian Journal of Philosophy* 14.

Steinitz, Y. (1994), "Brains in a Vat: Different Perspectives", *Philosophical Quarterly* 44.

Tymoczko, T. "Brains Don't Lie: They Don't Even Make Many Mistakes", in M. Roth and G. Ross (eds.), *Doubting* (Dordrecht: Kluwer).

Warfield, T. A. *"A priori* knowledge of the world: knowing the world by knowing our minds", in K. DeRose and T. A. Warfield (eds.), *Skepticism: A contemporary Reader* (Oxford: Oxford University Press).

Wright, C. (1992), "On Putnam's Proof That We Are Not Brains-in-a-Vat", *Proceedings of the Aristotelian Society* 92.