

# The Brain in a Vat Argument

The Brain in a Vat thought-experiment is most commonly used to illustrate global or Cartesian skepticism. You are told to imagine the possibility that at this very moment you are actually a brain hooked up to a sophisticated computer program that can perfectly simulate experiences of the outside world. Here is the skeptical argument. If you cannot now be sure that you are not a brain in a vat, then you cannot rule out the possibility that all of your beliefs about the external world are false. Or, to put it in terms of knowledge claims, we can construct the following skeptical argument. Let “P” stand for any belief or claim about the external world, say, that snow is white.

1. If I know that P, then I know that I am not a brain in a vat
2. I do not know that I am not a brain in a vat
3. Thus, I do not know that P.

The Brain in a Vat Argument is usually taken to be a modern version of René Descartes' argument (in the *Meditations on First Philosophy*) that centers on the possibility of an evil demon who systematically deceives us. The hypothesis has been the premise behind the movie *The Matrix*, in which the entire human race has been placed into giant vats and fed a virtual reality at the hands of malignant artificial intelligence (our own creations, of course).

One of the ways some modern philosophers have tried to refute global skepticism is by showing that the Brain in a Vat scenario is not possible. In his *Reason, Truth and History* (1981), Hilary Putnam first presented the argument that we cannot be brains in a vat, which has since given rise to a large discussion with repercussions for the realism debate and for central theses in the philosophy of language and mind. As we shall see, however, it remains far from clear how exactly Putnam's argument should be taken and what it actually proves.

## Table of Contents

1. Skepticism and Realism
2. Putnam's argument
3. Reconstructions of the Argument
4. Brains in a Vat and Self-Knowledge
5. Significance of the Argument
6. References and Further Reading

# 1. Skepticism and Realism

Putnam's argument is designed to attack the possibility of global skepticism that is implied by metaphysical realism. Putnam defines metaphysical realism as the view which holds that "...the world consists of some fixed totality of mind-independent objects. There is exactly one true and complete description of 'the way the world is.' Truth involves some sort of correspondence relation between words or thought-signs and sets of things." (1981, 49). This construal brings out the idea that for metaphysical realists, truth is not reducible to epistemic notions but concerns the nature of a mind-independent reality. This characterization finds an accurate target in those scientific materialists who believe in a "ready-made" world of scientific kinds independent of human classification and conceptualization. There are, however, many self-professed metaphysical realists who are not happy with Putnam's definition; it saddles the realist with the classical difficulty of matching words to objects and of providing for a correspondence relation between sentences and mind-independent "facts." The metaphysical realist is forced to construe her thesis ontologically, as an adherence to some fixed furniture of *objects* in the world, which ignores the possibility that ontological commitment may be specified not as a commitment to a set of entities but rather to the truth of a class of sentences or even of whole theories of the world.

One proposal is to construe metaphysical realism as the position that there are no *a priori* epistemically derived constraints on reality (Gaifman, 1993). By stating the thesis negatively, the realist sidesteps the thorny problems concerning correspondence or a "ready made" world, and shifts the burden of proof on the challenger to refute the thesis. One virtue of this construal is that it defines metaphysical realism at a sufficient level of generality to apply to all philosophers who currently espouse metaphysical realism. For Putnam's metaphysical realist will also agree that truth and reality cannot be subject to "epistemically derived constraints." This general characterization of metaphysical realism is enough to provide a target for the Brains in a Vat argument. For there is a good argument to the effect that if metaphysical realism is true, then global skepticism is also true, that is, it is possible that all of our referential beliefs about the world are false. As Thomas Nagel puts it, "realism makes skepticism intelligible," (1986, 73) because once we open the gap between truth and epistemology, we must countenance the possibility that all of our beliefs, no matter how well justified, nevertheless fail to accurately depict the world as it really is. [See Fallibilism.] Donald Davidson also emphasizes this aspect of metaphysical realism: "metaphysical realism is skepticism in one of its traditional garbs. It asks: why couldn't all my beliefs hang together and yet be comprehensively false about the actual world?" (1986, 309)

The Brain in a Vat scenario is just an illustration of this kind of global skepticism: it depicts a situation where all our beliefs about the world would presumably be false, even though they are well justified. Thus if one can prove that we cannot be brains in a vat, by *modus tollens* one can prove that metaphysical realism is false. Or, to put it in more schematic form:

1. If metaphysical realism is true, then global skepticism is possible
2. If global skepticism is possible, then we can be brains in a vat
3. But we cannot be brains in a vat
4. Thus, metaphysical realism is false (1,2,3)

This article focuses mostly on claim (3), although some philosophers question (2), believing there may be ways of presenting the skeptical thesis even while granting Putnam's argument.

## 2. Putnam's argument

The major premise that underwrites Putnam's argument is what he calls a "causal constraint" on reference:

(CC) A term refers to an object only if there is an appropriate causal connection between that term and the object

To understand this criterion we need to unravel what is meant by "appropriate causal connection." If an ant were to accidentally draw a picture of Winston Churchill in the sand, few would claim that the ant represented or referred to Churchill. Similarly, if I accidentally sneeze "Genghis Khan," just because I verbalize the words does not mean that I refer to the infamous Mongolian conqueror, for I may have never heard of him before. Reference cannot simply be an *accident*: or, as Putnam puts it, words do not refer to objects "magically" or intrinsically. Now establishing just what *would* count as necessary and sufficient conditions for a term to refer to an object turns out to be tricky business, and there have been many "causal theories" of reference supplied to do just that. Many have taken the virtue of Putnam's constraint (CC) to be its generality: it merely states a necessary condition for reference and need not entail anything more controversial. Sometimes it is claimed that endorsing (CC) commits you to semantic externalism but the issues are more complex, since many internalists (for example, John Searle) appear to agree with (CC). The relation between externalism and Putnam's argument will be considered in more detail later (in the section "Brains in a Vat and Self-Knowledge").

With the causal constraint established, Putnam goes on to describe the Brain in a Vat scenario. It is important to note exactly what the thought-experiment is, for failure to appreciate the ways in which Putnam has changed the standard skeptical nightmare has led to many mistaken "refutations" of the argument. The standard picture has a mad-scientist (or race of aliens, or AI programs...) envatting brains in a laboratory then inducing a virtual reality through a sophisticated computer program. On this picture, there is an important difference between viewing the brains from a first or third person viewpoint. There is the point of view of the brains in a vat (henceforth BIVs), and the point of view of someone outside the vat. Clearly when the mad-scientist says "that is a brain in a vat" of a BIV, he would be saying something *true*, no matter the question of what the BIV means when *it* says it is a brain in a vat. Furthermore, presumably a BIV could pick up referential terms by borrowing them from the mad-scientist. Thus when a BIV says "there is a tree" referring to a simulation of a tree, it would be saying something *false*, since its term "tree," picked up from the mad-scientist to refer to an actual tree, in fact refers to something else, like his sense-impressions of the tree. Putnam thus stipulates that *all* sentient beings are brains in a vat, hooked up to one another through a powerful computer that has no programmer: "that's just how the universe is." We are then asked, given at least the *physical* possibility of this scenario, whether we could say or think it. Putnam answers that we could not: the assertion "we are brains in a vat"

would be sense self-refuting in the same way that the general statement “all general statements are false” is.

The thought-experiment stipulates that brains in a vat would have qualitatively identical thoughts to those unenvatted; or at least they have the same “notional world.” The difference is that in the vat-world, there are no external objects. When a BIV says “There is a tree in front of me,” there is in fact no tree in front of him, only a simulated tree produced by the computer’s program. However, if there are no trees, there could be no causal connection between a BIV’s tokens of trees and actual trees. By (CC), “tree” does not refer to tree. This leads to some interesting consequences.

A standard reading of a BIV’s utterance of “There is a tree” would have the statement come out false, since there are no trees for the BIV to refer to. But that would be only assuming that “tree” refers to tree in the BIV’s language. If “tree” does not refer to tree, then the semantic evaluation of the sentence becomes unclear. Sometimes Putnam suggests that a BIV’s tokens refer to images or sense-impressions. At other times he agrees with Davidson who claims that the truth-conditions would be facts about the electronic impulses of the computer that are causally responsible for producing the sense-impressions. Davidson has a good reason to choose these truth-conditions: through the principle of charity he would want to interpret the BIV’s sentences to come out true, but he would not want the truth-conditions to be phenomenalistic. Thus it turns out that when a BIV says “There is a tree in front of me,” he is saying something *true*—if in fact the computer is sending the right impulses to him.

Another suggestion is that the truth-conditions of the BIV’s utterances would be empty: the BIV asserts nothing at all. This seems to be rather strong, however: surely the BIV would mean *something* when it utters “There is a tree in front of me,” even if its statement gets evaluated differently because of the radical difference of its environment. One thing is clear, however; a BIV’s tokening of “tree” or any other such referential term would have a *different* reference assignment from that of a non-envatted person’s tokenings. According to (CC), my tokening of “tree” refers to trees because there *is* an appropriate causal link between it and actual trees (assuming of course I am not a BIV). A brain in a vat however would not be able to refer to trees since there are no trees (and even if there were trees there would not be the appropriate causal relation between its tokenings of “tree” and real trees, unless we bring back the standard fantasy and assume it picked up the terms from the mad scientist). Now one might be inclined to think that because there are at least brains and vats in the universe, a BIV would be able to refer to brains and vats. But the tokening of “brain” is never actually caused by a brain except only in the very indirect sense that its brain causes *all* of its tokenings. The minimal constraint (CC) then will ensure that “brain” and “vat” in the BIV language does not refer to brain and vat.

We are now in a position to give Putnam’s argument. It has the form of a conditional proof :

1. Assume we are brains in a vat
2. If we are brains in a vat, then “brain” does not refer to brain, and “vat” does not refer to vat (via CC)

3. If “brain in a vat” does not refer to brains in a vat, then “we are brains in a vat” is false
4. Thus, if we are brains in a vat, then the sentence “We are brains in a vat” is false (1,2,3)

Putnam adds that “we are brains in a vat” is *necessarily* false, since whenever we assume it is true we can deduce its contradictory. The argument is valid and its soundness seems to depend on the truth of (3), assuming (CC) is true. One immediate problem is determining the truth-conditions for “we are brains in a vat” on the assumption we are brains in a vat, speaking a variation of English (call it Vatese). From (CC) we know that “brains in a vat” does not refer to brain in a vat. But it doesn’t follow from this alone that “we are brains in a vat” is false. Compare:

- (A) “Grass is green” is true iff grass is green
- (B) “Grass is green” is true iff one has sense-impressions of grass being green
- (C) “Grass is green” is true iff one is in electronic state Q

On the assumption that we are brains in a vat, (CC) would appear to rule out (A): “grass” does not refer to grass since there is no appropriate causal connection between “grass” and actual grass. Thus the truth-conditions for the statement “grass is green” would be nonstandard. If we take them to be those captured in (B), then “Grass is green” as spoken by a brain in a vat would be *true*. Consequently the truth-conditions for “we are brains in a vat” would be captured by (D):

- (D) “We are brains in a vat” is true iff we have sense impressions of being brains in a vat

On this construal of the truth-conditions, “We are brains in a vat” as uttered by a BIV would presumably be *false*, since a brain in a vat would *not* have sense-impressions of being a brain in a vat: recall a BIV’s notional world would be equivalent to the unenvatted, and he would appear to himself to be a normally embodied person with a real body etc. However, if we follow Davidson and adopt the truth-conditions of (C), we would have the following:

- (E) “We are brains in a vat” is true if and only if we are in electronic state Q

Now it is no longer clear that “We are brains in a vat” is false: for if the brain is in the appropriate electronic state, the truth-conditions could well be fulfilled. There are other reconstructions of the argument that do not depend on specifying the truth-conditions of a BIV’s utterances. What is important is the idea that the truth-conditions would be non-standard, as in:

- (F) “We are brains in a vat” is true if and only if we are BIVs\*

Now since being a BIV\* (whatever that is) is not the same as being a BIV, we can construct the following conditional proof argument:

1. Assume we are BIVs
2. If we are BIVs, “we are brains in a vat” is true if and only we are BIVs\*
3. If we are BIVs, we are not BIVs\*
4. If we are BIVs, then “we are BIVs” is false (2,3)

## 5. If we are BIVs, then we are not BIVs (4)

Notice that the argument leaves the antecedent of the conditional open, what Wright calls an “open subjunctive.” We do not want the premises of the argument to be counterfactual, following the train of thought “If we *were* brains in a vat, the causal constraint would entail that my words ‘brain in a vat’ would come to denote something different, BIV\*.” For then we would be assuming that we are not brains in a vat, when that is what the argument is supposed to prove.

Nevertheless, there are still problems with the appeal to disquotation to get us from (4) to (5). Even if, by virtue of the causal constraint, the sentence “We are BIVs” is false, an intuitive objection runs that this change of language should not entail falsity of the proposition *that we are brains in a vat*. As we shall see, many recent reconstructions of Putnam’s argument are sensitive to this point and try to account for it in various ways. In the following section, I shall focus on two of the more popular reconstructions of the argument put forward by Brueckner (1986) and Wright (1994).

### 3. Reconstructions of the Argument

Brueckner (1986) argues that even if we grant the reasoning of the above argument up to (4), the most the argument proves is that if we are brains in a vat, then the sentence “We are brains in a vat” (as uttered by a BIV) is false, and that if we are *not* brains in a vat, then “We are brains in a vat” is false (now expressing a *different* false proposition). If correct then the argument would prove that whether or not we are brains in a vat, “we are brains in a vat” expresses *some* false proposition. Assuming the truth-conditions of a BIV would be those captured in (D) we could then devise the following constructive dilemma type argument:

1. Either I am a BIV or I am not a BIV
2. If I am a BIV, then “I am a BIV” is true iff I have sense impressions of being a BIV
3. If I am a BIV, then I do not have sense-impressions of being a BIV
4. If I am a BIV, then “I am a BIV” is false (2,3)
5. If I am not a BIV, then “I am a BIV” is true iff I am a BIV
6. If I am not a BIV, then “I am a BIV” is false (5)
7. “I am a BIV” is false (1, 4, 6)

If “I am a BIV” expresses the proposition that I am a brain in a vat, and we know from the argument that “I am a BIV” is false, then it follows that I know I am not a brain in a vat, thus refuting premise (2) of the skeptical argument. However, can I know that “I am a brain in a vat” expresses the proposition that I am a brain in a vat? If I am a brain in a vat, then “I am a brain in a vat” would, via the causal constraint on reference, express some different proposition (say, *that I am a brain in a vat in the image*). So even if “I am a BIV” is false whether or not I am a BIV, I might not be in the position to identify *which* false proposition I am expressing, in which case I cannot claim to know that my sentence “I am not a brain in a vat” expresses the true proposition *that I am not a brain in a vat*.

Some philosophers have gone even further, claiming that if the argument ends here, it actually can be used to strengthen skepticism. The metaphysical realist can claim that there are truths not expressible in any language: perhaps the proposition *that we are brains in a vat* is true, even if no one can meaningfully utter it. As Nagel puts it:

If I accept the argument, I must conclude that a brain in a vat can't think truly that it is a brain in a vat, even though others can think this about it. What follows? Only that I cannot express my skepticism by saying "Perhaps I am a brain in a vat." Instead I must say "Perhaps I can't even think the truth about what I am, because I lack the necessary concepts and my circumstances make it impossible for me to acquire them!" If this doesn't qualify as skepticism, I don't know what does. (Nagel, 1986)

Putnam makes it clear that he is not merely talking about semantics: he wants to provide a metaphysical argument that we cannot be brains in a vat, not just a semantic one that we cannot assert we are. If he is just proving something about meaning, it is open for the skeptic to say that the bonds between language and reality can diverge radically, perhaps in ways we can never discern.

There is yet another worry with the argument, centering once again on the appropriate characterization of the truth-conditions in (2). If one claimed in response to the above objection that in fact I do know that "I am a brain in a vat" expresses the proposition that I am a brain in a vat (whether or not I am a brain in a vat), one may have in mind some general disquotation principle:

(DQ): "Grass is green" is true iff grass is green

If it is an *a priori* truth that any meaningful sentence in my language homophonically disquotes, then we can *a priori* know that the following is also true:

(F): "I am a brain in a vat" is true iff I am a brain in a vat

Here is the obvious problem: if we are not to beg the question, we have to be open to the possibility that we are brains in a vat, speaking Vatese. Then we would get:

(G): If I am a BIV, then "I am a BIV" is true iff I am a brain in a vat.

However, (G) gives us truth-conditions that differ from premise (2) of Brueckner's argument:

(2) If I am a BIV, then "I am a BIV" is true iff I have sense-impressions of being a BIV

If we assume (CC), then (G) and (2) are inconsistent, since the term "BIV" would refer to distinct entities. No contradiction ensues if we assume we are speaking in English: for then (G) would presumably be false (appealing to CC). But the problem is that we cannot beg the question by assuming we are speaking in English: if we assume that, then we know in advance of any argument

that we are not speaking in Vatese and hence that we are not brains in a vat. But if we do not know which language we are speaking in, then we cannot properly assert (2).

One response to this is to formulate two different arguments, one whose meta-language is in English, the other whose meta-language is in Vatese, and show that distinct arguments can be run to prove that “I am a BIV” is false. Even if successful, however, these arguments run into the objection canvassed before: if I do not know which language I am speaking in, even if I know “I am a brain in a vat” is false, I do not know which false proposition I am expressing and hence cannot infer that I know that I am not a brain in a vat.

Similar worries plague Crispin’s Wright’s popular formulation of the argument (1994):

1. My language disquotes
2. In BIVese, “brains in a vat” does not refer to brains in a vat
3. In my language, “brains in a vat” is a meaningful expression
4. In my language, “brains in a vat” refers to brains in a vat
5. My language is not BIVese (2,4)
6. If I am a BIV, then my language is BIVese
7. I am not a BIV

There are several virtues to this reconstruction: first of all, it gets us to the desired conclusion without specifying what the truth-conditions of a BIV’s utterances would be. They could be sense-impressions, facts about electronic impulses, or the BIV’s sentences may not refer at all. All that is needed for the argument is that there is a *difference* between the truth-conditions for a BIV’s sentences and those of my own language. The other virtue of the argument is that it clearly brings out the appeal to the disquotation principle that was implicit in the previous arguments. If indeed (DQ) is an *a priori* truth, as many philosophers maintain, and if we accept (CC) as a condition of reference, the argument appears to be sound. So have we proven that we are not brains in a vat?

Not so fast. The previous objection can be restated: if I do not yet know whether or not I am a brain in a vat before the argument is completed, I do not know which language I am speaking (English or Vatese). If I am speaking Vatese, then so long as it is a meaningful language, I can appeal to disquotation to establish that “brains in a vat” does refer to brains in a vat. But this contradicts premise (2). The problem seems to be that (DQ) is being used too liberally. Clearly we do not want to say that every meaningful term disquotes in the strong sense required for reference. If so, we could take it to be an *a priori* truth that “Santa Claus” refers to Santa Claus. But “Santa Claus” does not refer to Santa Claus, since there is no Santa Claus. We could introduce a new term “pseudo-reference” and hold that “Santa Claus” *pseudo-refers* to Santa Claus, and then attach further conditions on reference in order to establish what it would take for the term to truly refer. One proposal (Weiss, 2000) is the following principle:

W: If “x” pseudo-refers to x in L, and x exists, then “x” refers to x in L



Thus, given the disquotation principle we know that in my language “Santa Claus” pseudo-refers to Santa Claus. Supposing to the joyful adulation of millions that Santa Claus is discovered to actually exist, then given (W) “Santa Claus” refers to Santa Claus. Now this also seems too simplistic: as Putnam pointed out, in order for a term to refer to an object we must establish more than the mere existence of the object. There has to be the *appropriate* causal relation between the word and object, or we are back to claiming that in accidentally sneezing “Genghis Khan” I am referring to Genghis Khan. But whether we accept (W) or attach stronger conditions to reference, it is clear that any such move would make Wright’s formulation invalid. For then we would have:

1. My language disquotes
2. In BIVese, “brains in a vat” does not refer to brains in a vat (CC)
3. In my language “brain in a vat” is a meaningful expression
4. In my language, “brain in a vat” pseudo-refers to brains in a vat (DQ)
5. My language is not BIVese (2,4)
6. If I am a BIV, then my language is BIVese
7. I am not a BIV

(5) no longer follows from (2) and (4) given the ambiguity of “refers” in (2) and (4). If on the other hand we insist on a univocal sense of reference, then either (2) will contradict the (DQ) principle, or we are not entitled to appeal to (1), insofar as it would beg the question that we are speaking English, a language for which the (DQ) principle applies.

## 4. Brains in a Vat and Self-Knowledge

Ted Warfield (1995) has sought to provide an argument that we are not brains in a vat based on considerations of self-knowledge. He defends two premises that seem reasonably true, and then he argues for the desired metaphysical conclusion:

1. I think that water is wet
2. No brain in a vat can think that water is wet
3. Thus, I am not a brain in a vat (2.3)

Premise (1) is said to follow from the thesis of privileged access, which holds that we can at least know the contents of our own occurring thoughts without empirical investigation of our environment or behavior. Warfield’s strategy is to present each premise as non-question begging against the global skeptic, in which case at no point can we appeal to the external environment as justification. Since the thesis of privileged access is said to be known *a priori* whether we are brains in a vat or not, premise (1) can be known non-empirically.

Premise (2) is a little trickier to establish non-empirically. The main argument for it is by analogy with other arguments in the literature that have been used to establish content externalism. The main strategy is derived from Putnam’s Twin Earth argument (1975): imagine a world that is indistinguishable from Earth except for one detail: the odorless, drinkable liquid that flows in the

rivers and oceans is composed of the chemicals XYZ and not H<sub>2</sub>O. If we take Oscar on Earth and his twin on Twin-earth, Putnam argues that they would refer to two different substances and hence mean two different things: when Oscar says “pass me some water” he refers to H<sub>2</sub>O and means *water*, but when Twin-Oscar says “pass me some water” he refers to XYZ and thus means *twin-water*. As Burge and others have pointed out, if the meaning of their words are different, then the concepts that compose their beliefs should differ as well, in which case Oscar would believe *that water is wet* whereas Twin-Oscar would believe *that twin-water is wet*. While Putnam’s original slogan was “meanings just ain’t in the head,” the argument can be extended to beliefs as well: “beliefs just ain’t in the head,” but depend crucially on the layout of one’s environment.

If we accept content externalism, then the motivation for (2) is as follows. In order for someone’s belief to be about water, there must be water in that person’s environment: externalism rejects the Cartesian idea that one can simply read off one’s belief internally (if so then we would have to say that Oscar and his twin have the same beliefs since they are internally the same). So it doesn’t seem possible that a BIV could ever come to hold a belief about water (unless of course he picked up the term from the mad-scientist or someone outside the vat, but here we must assume again Putnam’s scenario that there is no mad-scientist or anyone else he could have borrowed the term from). As Warfield puts it, premise (2) is a conceptual truth, established on the basis of Twin-earth style arguments, a matter of “armchair” *a priori* reflection and thus able to be established non-empirically.

The problem with establishing (2) non-empirically though is that the externalist arguments succeed only on the assumption that our own use of “water” refers to a substantial kind, and this seems to be a matter of empirical investigation. Imagine a world where “water” does not refer to any liquid substance but is rather a complex hallucination that never gets discovered. On this “Dry Earth,” “water” would not refer to a substantial kind but rather a superficial kind. The analogy to the BIV case is clear: since it is not an *a priori* truth that “water” refers to a substantial kind in the BIV’s language, it cannot be known non-empirically that “water” is substantial or superficial; if it is a superficial kind, then a BIV could very well think that water is wet so long as it has the relevant sense-impressions.

## 5. Significance of the Argument

Some philosophers have claimed that even if Putnam’s argument is sound, it doesn’t do much to dislodge Cartesian or global skepticism. Crispin Wright (1994) argues that the argument does not affect certain versions of the Cartesian nightmare, such as my brain being taken out of my skull last night and hooked up to a computer. Someone of a Positivist bent might argue that if there is no empirical evidence to appeal to in order to establish whether we are brains in a vat or not, then the hypothesis is meaningless, in which case we do not need an argument to refute it. While few philosophers today would hold onto such a strong verifiability theory of meaning, many would maintain that such metaphysical possibilities do not amount to real cases of doubt and thus can be summarily dismissed. Still others see the possibility of being a brain in a vat an important challenge for cognitive science and the attempt to create a computer model of the world that can

simulate human cognition. Dennett (1991) for example has argued that it is physically impossible for a brain in a vat to replicate the qualitative phenomenology of a non-envatted human being. Nevertheless, one should hesitate before making possibility claims when it comes to future technology. And as films like the *Matrix*, *Existenz*, and even the *Truman Show* indicate, the idea of living in a simulated world indistinguishable from the real one is likely to continue to fascinate the human mind for many years to come—whether or not it is a brain in a vat.

## 6. References and Further Reading

- Boghossian, Paul. 1999. What the Externalist can Know A Priori. *Philosophical Issues* 9
- Brueckner, Anthony. 1986. Brains in a Vat. *Journal of Philosophy* 83: 148-67
- Brueckner, Anthony. 1992. If I am a Brain in a vat, then I am not a Brain in a Vat. *Mind* 101:123-128
- Burge, T. 1982. Other Bodies. In A. Woodfield. Ed., *Thought and Object: Essays on Intentionality*. Oxford: Oxford University Press, 91-120.
- Casati, R. and Dokic J. 1991. *Brains in a Vat, Language and Metalanguage*. *Analysis* 51: 91-93.
- Collier, J. 1990. Could I Conceive Being a Brain in a Vat? *Australasian Journal of Philosophy* 68: 413-419.
- Davidson, Donald. 1986. "A Coherence Theory of Truth and Knowledge," in *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Oxford: Blackwell.
- Davies, D. 1995. Putnam's Thought-Teaser. *Canadian Journal of Philosophy* 25(2):203-227.
- Ebbs, G. (1992), "Skepticism, Objectivity and Brains in Vats", *Pacific Philosophical Quarterly* 73
- Forbes, G. 1995. Realism and Skepticism: Brains in a Vat Revisited. *Journal of Philosophy* 92(4): 205-222
- Gaifman, Haim. 1994. Metaphysical Realism and Vats in a Brain. (unpublished ms)
- Nagel, Thomas. 1986. *The View from Nowhere*. Cambridge: Cambridge University Press.
- Noonan, Harold. 1998. Reflections on Putnam, Wright and brains in a vat. *Analysis* 58:59-62
- Putnam, Hilary 1975. The Meaning of "Meaning." *Mind, Language and Reality: Philosophical Papers*, Vol 1. Cambridge: Cambridge University Press
- Putnam, Hilary. 1982. *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Putnam, Hilary. 1994. Reply to Wright. In P. Clark and B. Hale, eds. *Reading Putnam*. Oxford, Blackwell.
- Sawyer, Sarah. 1999. My Language Disquotes. *Analysis*, vol. 59:3: 206-211
- Smith, P. (1984), Could We Be Brains in a Vat?, *Canadian Journal of Philosophy* 14
- Steinitz, Y. Brains in a vat? Different Perspectives. *Philosophical Quarterly* 44 (175): 213-222
- Tymoczko, T. 1989. In Defense of Putnam's Brains. *Philosophical Studies* 57(3) 281-297
- Warfield, Ted. 1995. Knowing the World and Knowing our Minds. *Philosophy and Phenomenological Research* 55 (3): 525-545.
- Weiss, B. 2000. Generalizing Brains in a Vat. *Analysis* 60: 112-123
- Wright, Crispin. 1994. On Putnam's Proof that we cannot be brains in a vat. In P. Clark and B. Hale. Eds, *Reading Putnam*. Oxford: Blackwell.

## Author Information

Lance P. Hickey

Email: lance1001@optonline.net

Southern Connecticut State University

U. S. A.

