

Counterpossibilities and Vague Worries

By

Nathaniel Baird

A thesis submitted in partial fulfillment of the requirements  
for graduation with Honors in Philosophy.

Whitman College  
2015

*Certificate of Approval*

This is to certify that the accompanying thesis by Nathaniel Baird has been accepted in partial fulfillment of the requirements for graduation with Honors in Philosophy.

---

Mitch Clearfield

Whitman College  
May 12, 2015

## Table of Contents

Introduction.....	1
1. The Radical Counterpossibility.....	4
1.1.....	5
1.2.....	10
1.3.....	17
1.4.....	24
2. Conceptual Deficiency.....	30
2.1.....	30
2.2.....	33
2.3.....	38
Conclusion .....	45
Bibliography .....	46

## Introduction

Skeptical arguments are philosophical puzzles which are embarrassingly difficult to solve. They embarrass us because though they put forth the most outrageous conclusions (it is outrageous to suppose that we know nothing at all), discovering where these arguments go wrong is a great challenge. They incorporate premises which are not obviously wrong and often are difficult to reject. The skeptical argument, then, is a tool for bringing to light the contradictions between one's intuitions. That is, the consequence of many common presuppositions is that knowledge is impossible, a consequence which contradicts the fact that we are knowers.<sup>1</sup> Thus the analysis of skepticism is of great value even to those who refuse to consider accepting it as true.

In this paper I will investigate a style of skeptical argument used to threaten knowledge of the external world. The argument makes use of a counterpossibility. A counterpossibility is a possible state of affairs which is not known to be false. It could be false, but it also could be true. The skeptic argues that there is no way to know whether or not the counterpossibility is actually the case. If the counterpossibility *is* actually the case, the argument runs, then everything that is incompatible with the counterpossibility is not the case. For example: If I am dreaming then I am not standing up (I am in fact lying on my bed). The skeptic argues that since I do not know that the counterpossibility is false I do not know that any state of affairs incompatible with the counterpossibility is true. If I do not know that I am not dreaming then I do not know that I am standing up. The form of the skeptical argument is as follows:

### The Skeptical Argument

---

<sup>1</sup> Or perhaps that we are knowers is an *intuition* that we cannot reject.

1. I do not know that (A) is not the case
2. If (A) is the case, then (B) is not the case
3. I do not know that (B) is the case

(A) is a counterpossibility. (B) is a state of affairs incompatible with (A). The argument attacks the knowledge that (B). Just what knowledge is under attack depends on the details of (A).<sup>2</sup> It is in this sense that the force of the skeptical argument corresponds to the version of the counterpossibility which it utilizes. A weak skeptical argument utilizes a counterpossibility incompatible with relatively few states of affairs. It threatens few of our beliefs. A devastating skeptical argument is one which assaults a great deal of our knowledge. It is in the skeptic's interests, then, to utilize a counterpossibility which runs contrary to a great number of our beliefs about the external world.

This strategy can be viewed as a skeptical worry with a reason behind it. It is, in this sense, a reasonable worry. I do not mean to suggest that it is practical to take outlandish skeptical possibilities seriously, but rather, that there is a difference between a skeptical worry which is justifiable and one which is not. A reasonable skeptic, I submit, is one who backs up her worry with the sort of argument I have outlined above. In contrast, an unreasonable skeptic is one who dogmatically asserts a skeptical stance without giving any clear reason for doing so. The sort of reasonable skepticism I am

---

<sup>2</sup> The success of the skeptical strategy relies on the move from (1) to (3) via (2). This is frequently referred to as the epistemic closure principle. Some philosophers, such as Dretske (1970) and Nozick (1981), have provided compelling reasons to reject the principle. Since my anti-skeptical argument does not depend on such a maneuver I will grant the skeptic epistemic closure. One may also argue that the skeptic assumes that certainty is necessary for knowledge, and that one need not be certain that (A) in order to know that (A). Or perhaps one will suggest that there are different senses of 'to know' and that some depend on certainty while others do not. In this paper I will also grant the skeptic's assumption that knowledge demands certainty.

concerned with is more than a personal outlook or mood; the reasonable skeptic provides an argument. If the argument is good I must endorse its conclusion.

I believe that this skeptical strategy suffers from a serious limitation. My argument will initially focus on what I consider to be a ‘radical counterpossibility’, that is, a state of affairs which appears incompatible with every belief I have about the external world. I will first consider an argument against skepticism made by Hilary Putnam in the first chapter of *Reason, Truth and History*. Then I will refute a tempting counterargument made by Anthony Brueckner. My refutation will expand off of Putnam’s argument to show why a skeptical argument utilizing a radical counterpossibility must be either incoherent or a vague. Neither of these are qualities of a reasonable skeptical argument. From there I will consider other options available to the skeptic. I will show that any reasonable skeptical argument must either utilize a weaker sort of counterpossibility incompatible with considerably fewer of our beliefs about the external world or focus entirely on the possibility of conceptual deficiency rather than falsehood. I will show that there are limitations both on the degree to which our beliefs about the external world are false and the degree to which we are conceptually deficient.

## 1. The Radical Counterpossibility

Putnam provides an example of a radical counterpossibility:

“We could imagine that all human beings (perhaps all sentient beings) are brains in a vat (or nervous systems in a vat in case some beings with just a minimal nervous system already count as ‘sentient’). Of course, the evil scientist would have to be outside – or would he? Perhaps there is no evil scientist, perhaps (though this is absurd) the universe just happens to consist of automatic machinery tending a vat full of brains and nervous systems... The machinery is programmed to give us all a *collective* hallucination, rather than a number of separate unrelated hallucinations. Thus, when I seem to myself to be talking to you, you seem to yourself to be hearing my words. Of course, it is not the case that my words actually reach your ears – for you don’t have (real) ears, nor do I have a real mouth and tongue. Rather, when I produce my words, what happens is that the efferent impulses travel from my brain to the computer, which both causes me to ‘hear’ my own voice uttering those words and ‘feel’ my tongue moving, etc., and causes you to ‘hear’ my words, ‘see’ me speaking, etc. In this case, we are, in a sense, actually in communication. I am not mistaken about your real existence (only about the existence of your body and the ‘external world’, apart from brains)” (Putnam, 6).

The counterpossibility is radical in the following way: it seems to be incompatible with all knowledge of the external world.<sup>3</sup> If the counterpossibility is actually the case then seemingly every claim I have ever made or belief I have ever had about the external world has been false. That is because what I have taken to be the real world all this time has actually been a hallucination. My understanding of the world does not match the way the world really is. I think the world contains trees and dogs etc., but actually it contains brains, a vat, and automatic machinery.

---

<sup>3</sup> Perhaps there is knowledge of the external world immune to the counterpossibility due to its generality. An example of such general knowledge might be the idea that objects take up space. But even if the most general knowledge is safe from the skeptical attack such a fact hardly consoles us. We know a great deal. Unless one wishes to accept the skeptical conclusion a response to the skeptic must involve defending knowledge, not some miniscule subset of it.

## 1.1

Putnam argues that the radical counterpossibility is self-refuting (7). His argument involves an analysis of the necessary and sufficient conditions for reference. An object does not refer in virtue of its intrinsic properties (e.g. how it looks, sounds etc.).<sup>4</sup> In order for an object to refer a causal connection must hold between object and referent. This conclusion reveals what *appear* to be certain acts of reference (ones without causal connection between referent and representation<sup>5</sup>) to not involve reference at all. The radical counterpossibility, it turns out, requires such acts of reference in order to be the case.<sup>6</sup> The radical counterpossibility, then, is not the case.

Putnam asks us to imagine an ant crawling across a stretch of sand (1). As it crawls it leaves a trail in the sand which would look to us like an illustration of Winston Churchill. Intuitively we would not say that the image represents Winston Churchill. We would not say this because the ant does not know anything about Winston Churchill and did not intend to represent him in any way. It does not even know that it is leaving behind a path. If it were a very intelligent and knowledgeable ant then we would be more willing to see its path as a depiction of Churchill rather than a path which happened to seem that way to us. The consequence of this is that physical objects do not intrinsically refer (2). The qualities of a physical thing do not make it refer. Reference must depend on other factors. For, if I were to draw such an image in the sand it *would* represent Winston

---

<sup>4</sup> An 'Object' here can be anything physical (such as a drawing), mental (such as a thought) or linguistic (such as a name).

<sup>5</sup> As I use the term, an object is a 'representation' if it refers to something.

<sup>6</sup> It is true that there are meaningful linguistic expressions which do not refer, and so it is not the case that reference is necessary for meaning. The present context, however, permits me to operate with a close linkage between the two. The skeptic has in mind a counterpossibility which describes a possible state of affairs. The task of providing such a counterpossibility demands one's language be capable of stating precisely its details. It may very well be the case that the skeptic's utterances are meaningful despite not picking out a possible state of affairs. Perhaps the skeptic's utterances play a role in a language game, for example. But for the skeptic this is not good enough.

Churchill, even if it were exactly like the image which the ant was responsible for. We can, in fact, imagine a person drawing a representation of Winston Churchill which looks less to us like him than the ant's path does, and yet we would still say that the former refers to Winston Churchill while the latter does not. A physical representation need not look a certain way to refer (the words 'Winston Churchill' refer to a person, and yet they certainly look nothing like him).

What fact about my illustration of Winston Churchill that would make it refer to him? It is tempting to assume, because the ant's lack of intelligence appears to relate to the fact that it cannot refer, that what it takes for something to refer is some feature of the mind. Perhaps when I draw a picture of Winston Churchill it refers to him because in my mind there is a mental picture of Winston Churchill which accompanies the act of drawing, and perhaps mental pictures intrinsically refer to things. Putnam argues that this is not the case (3). Mental representations are subject to the same constraints as physical representations. No feature of the mind is solely responsible for getting something to refer. Putnam asks us to imagine a planet that contains absolutely no trees at all. Suppose a spacecraft is flying over this planet. In the spacecraft is a human who is familiar with trees. He has in his hand a picture that he took on Earth of a tree. Imagine that he drops the photograph down onto the planet, and suppose that one of its human-like inhabitants picks up the photograph and looks at it. We can imagine that the inhabitant wonders about what sort of object is depicted in the photograph. Suppose that the inhabitant has a mental picture of the object from the photograph, and suppose that it is exactly the same sort of mental picture that I would have if I were to look at the same photograph. Putnam argues that my mental picture would represent a tree, while the inhabitant's does not. Its

mental picture cannot be a representation of a tree because the inhabitant has never had anything to do with trees. It has never seen or heard of a tree. Its mental picture is at best a representation of the strange object found in the photograph. To see that this is so, imagine that instead of a photograph it is actually an accidental smearing of various spilled paints (4). The inhabitant would have a mental image identical to the one it had in the previous case, however the source of its mental image does not in fact depict a tree at all! The smeared paints no more refer to trees than the ant's path refers to Winston Churchill. Mental pictures, then, do not represent intrinsically either.

We can combine these examples to imagine a person who paints a picture of a what would appear *to us* to be a tree, has the same sort of mental picture we may have when we think about trees, and utters the words 'I am painting a picture of a tree' in the process (accompanied by a 'feeling of understanding'), and yet not once refers to trees. The fact that this person would seem to us at first glance to be referring to trees does not make it so. Whether or not one refers to something (or refers to anything at all) does not depend on what sorts of noises they make, what sorts of mental pictures they have, or what sorts of physical things they produce. No physical, mental or linguistic representation refers because of the intrinsic features it has. That is because for anything which is a representation something could look (or sound) identical to it and not refer. Reference depends on factors other than those intrinsic to the representation.

But what other factor could do the trick? We saw earlier that whether or not an object refers depends on who is responsible for it. The ant, we thought, cannot represent Winston Churchill since it does not know who Winston Churchill was. Yet I could draw an image of Winston Churchill (one which *seemed* to resemble him less than the ant's

accidental markings) and it would refer to him. Similarly the human who is familiar with trees can have a mental picture that refers to them. What, then, do I have that the ant does not? What is it about the inhabitant that, despite having the same mental picture as the human, prevents it from referring to trees?

The inhabitant has no significant causal connection to trees. The ant, similarly, does not have the sorts of capacities required to be causally connected to Winston Churchill in a way which allows it to refer to him. The importance of a causal connection can be brought out through considering the difference between referring and merely *seeming* to refer. The ant's picture *seemed* to represent Churchill (just as the inhabitant's mental picture seemed to refer to trees). The difference between referring to Winston Churchill and seeming to do so is in how both acts relate to the man. It is probable that the ant would have made the same path had Winston Churchill never existed. The inhabitant, in the same way, could have a mental picture of the image caused by the spilled paints even if trees had never existed at all. There is no relationship at all in these cases between what is seemingly a representation and what is seemingly being represented.<sup>7</sup> If the inhabitant's mental picture depended on the existence of trees we would be more willing to say that it is capable of referring to them (11).

If it were the case that for some reason the ant would not have made such a path had Winston Churchill never existed, it would *not* follow that it was capable of reference. It is not just any sort of causal relationship that is important here, then. An example of a relevant causal connection is one in which verbal utterances (though they need not be verbal utterances) are provided in response to some feature of the environment (perhaps

---

<sup>7</sup> Though, obviously, there is a relationship between the fact that the ant's path seems *to us* to be a representation and Winston Churchill.

they are uttered when the speaker sees an object), or when after an utterance is made the speaker then interacts with their environment (perhaps by eating a certain object).<sup>8</sup> In such a circumstance it is plausible to suppose that there is a representational relationship between the utterance and whatever feature of the environment accompanies it before or after. It is not enough that somebody is in a world with trees in order for them to refer to trees (the ant is in the same world as Winston Churchill). In order for reference to take place an individual must be causally connected in a certain way to the thing which they are referring to. The existence of trees must have some influence on one's utterances in order for those utterances to refer to trees.

What is most important for my purposes is that a certain sort of causal connection (whatever exactly it turns out to be<sup>9</sup>) is necessary for reference. Earlier I said that a certain untenable sort of reference is required for the radical counterpossibility to be the case. The radical counterpossibility presupposes a capacity to refer which is indefensible given Putnam's conclusions. It demands of a BIV that it refer to objects that it cannot possibly refer to. "The question we are interested in is this: do the BIV's utterances containing, say, the word 'tree' actually refer to trees? More generally: can it refer to external objects at all" (12)? Why would we deny that a BIV can refer to trees? The reason is simple. Though the BIV has the same internal mental life as a non-BIV, as we saw earlier, the qualities of a representation do not cause it to refer. A BIV's mental images and utterances do not refer to trees even though they are indistinguishable from the mental images and utterances of a non-BIV. In fact, there are not even any trees in the

---

<sup>8</sup> This is what Putnam has in mind when he mentions language entry and exit rules (11).

<sup>9</sup> My intentions here do not demand that I sketch out a detailed causal theory. What is most important is that, regardless of the specific relationship necessary, reference *without* any causal relationship is impossible.

world of the BIV (there are only brains, vats, and machinery). What follows is that, regardless of what a BIV refers to, it cannot refer to trees.<sup>10</sup> It also cannot refer to brains or vats. A BIV cannot not say or think that it is a BIV (7). To suggest otherwise is to believe that a BIV can refer to brains and vats despite not being in the right kind of causal contact with brains or vats.

If a BIV is incapable of saying or thinking that it is a BIV then it is incapable of worrying that it is a BIV. The skeptical worry is that I am a BIV. But that is not a worry that a BIV can have. A BIV can worry that it is whatever it means by 'BIV,' but it is not *that*. A BIV is not a BIV-in-the-image, a certain electronic impulse or the presence of a certain kind of program. Similarly, a non-BIV that worries that he is a BIV is not a BIV. 'I am a BIV,' when uttered by either the BIV or non-BIV, is false. Putnam says:

"It follows that if...we are really the brains in a vat, then what we now mean by 'we are brains in a vat' is that we are brains in a vat in the image or something of that kind (if we mean anything at all). But part of the hypothesis that we are brains in a vat is that we aren't brains in a vat in the image (i.e. what we are 'hallucinating' isn't that we are brains in a vat). So, if we are brains in a vat, then the sentence 'We are brains in a vat' says something false (if it says anything). In short, if we are brains in a vat, then 'We are brains in a vat' is false. So it is (necessarily) false" (15).

## 1.2

There is a temptation to think that Putnam's argument accomplishes less than it claims to. There is a gut feeling that Putnam is merely toying around with words. At this

---

<sup>10</sup> Putnam says: "On some theories that we shall discuss it might refer to trees in the image, or to the electronic impulses that cause tree experiences, or to the features of the program that are responsible for those electronic impulses. These theories are not ruled out by what was just said, for there is a close causal connection between the use of the word 'tree' in vat-English and the presence of trees in the image, the presence of electronic impulses of a certain kind, and the presence of certain features in the machine's program. On these theories the brain is *right*, not *wrong*, in thinking 'There is a tree in front of me'" (14).

point one might imagine a BIV reciting Putnam's argument. What then can the argument accomplish? Can I not worry that I am a BIV thinking the false thought 'I am a BIV'?

Anthony Brueckner (1986) presents an argument defending this intuition. He goes further to argue that Putnam's argument begs the question, that it presupposes that I am not a BIV before it supposedly goes on to prove such a conclusion. I will present his reasoning and then show that it misses what I take to be the central point of Putnam's argument. In order to make precise the foundations of his reasoning Brueckner provides a more intricate reconstruction of the anti-skeptical argument made by Putnam.

Brueckner's reconstruction appears as follows:

Argument E (Brueckner, 154)

1. Either I am a BIV (speaking vat-English) or I am not a BIV (speaking English)
2. If I am a BIV (speaking vat-English), then my utterances of 'I am a BIV' are true iff I have sense impressions as of being a BIV
3. If I am a BIV (speaking vat-English), then I do not have sense impressions of being a BIV
4. If I am a BIV (speaking vat-English), then my utterances of 'I am a BIV' are false
5. If I am a non-BIV (speaking English), then my utterances of 'I am a BIV' are true iff I am a BIV
6. If I am a non-BIV (speaking English), then my utterances of 'I am a BIV' are false
7. My utterances of 'I am a BIV' are false

The argument aims to show that when the non-BIV utters ‘I am a BIV’ the utterance is false, and when the BIV utters ‘I am a BIV,’ that utterance is also false. And so any utterance of ‘I am a BIV’ is false. This is because the BIV and non-BIV each mean something different by the same utterance. What the words in the utterance refer to depend on the environment which the utterer is in. The non-BIV’s utterances containing ‘vat’ refer to vats. The BIV’s utterances containing ‘vat’ refer to sense impressions of vats.<sup>11</sup> This is why the BIV and non-BIV are said to speak different languages (vat-English and English, respectively).

How is it that Putnam’s argument supposedly begs the question? Brueckner argues it does so in virtue of the implications embedded in the language which it is provided in. To see that there is something odd regarding the language of E, consider that premise (2) appears to be incorrect if I am a BIV (156). Consider the following:

(T) My utterances of ‘I am a BIV’ are true iff I am a BIV

If (T) is uttered by a non-BIV it is true. If (T) is uttered by a BIV it is true.

Brueckner says that “(T) is true so long as the metalanguage used in stating (T) is the same as (or contains) the language of the mentioned sentence” (156). The point is that (T) is true so long as the mentioned sentence (‘I am a BIV’) and the rest of (T) are both in the language of the utterer. If that is the case, (T) is true when uttered in English and when uttered in vat-English.

---

<sup>11</sup> Which is the first of Putnam’s several suggestions (mentioned in footnote 10) for what the BIV’s utterances refer to. If you are inclined to prefer another candidate (perhaps the BIV’s utterances containing ‘vat’ refer to the presence of electronic impulses of some kind) you are welcome to substitute it. So long as the BIV’s utterances of ‘I am a BIV’ remain false such preferences are compatible with the argument.

Why is this seemingly a problem for (2)? (2) is false if the language the premise is stated in is vat-English. It is only correct if the language used to state the premise is English. It is capturing the truth conditions of a statement in vat-English, *in English* (156). However, the implication of the antecedent of (2) is that (2) be in vat-English. This is because if I am a BIV I speak vat-English. It appears that if the antecedent of (2) is true I speak vat-English. But, if read in vat-English, the consequent of (2) does not accurately capture the truth conditions of my utterances of 'I am a BIV.' If a BIV were to utter the consequent of (2):

My utterances of 'I am a BIV' are true iff I have sense impressions as of being a BIV

it would be false, unlike (T). The consequent of (2) is supposedly true in a world in which I am a BIV, and yet it appears that when uttered in vat-English it is false. (T), rather than the consequent of (2), captures the truth conditions of 'I am a BIV' assuming I am a BIV. Brueckner introduces a rival premise to (2). Since (T) would be true if uttered by a BIV, we have:

8. If I am a BIV, then my utterances of 'I am a BIV' are true iff I am a BIV

Premises (2) and (8) are clearly incompatible. If (3) is correct and if I am a BIV then (2) says that my utterances of 'I am a BIV' are false, while (8) says that they are true. So which premise accurately captures the truth conditions of my utterances if I am a BIV? For (8) to be true the language specified in its antecedent must be the language which its

consequent is in. Vat-English is the language specified in the antecedent of (8). (8) is only true if we read its consequent as being in vat-English. “But this fact about (T) establishes (8) only if, in evaluating the truth value of (8) at a world in which its antecedent is true, we must interpret its consequent as being in vat-English and evaluate its truth value accordingly” (158). Brueckner argues that to read (8) such that the language specified in its antecedent is the language the consequent is stated in is a “nonstandard” (157) reading. To show this he considers:

9. If I am speaking a language in which ‘tail’ refers to legs, then horses have four tails

According to Brueckner, the reading of (9) that makes it true is just as nonstandard. The nonstandard reading takes the language specified in the antecedent (the language in which ‘tail’ refers to legs) as that which the consequent is stated in. Indeed, (9) is true if we read it in this way. What is wrong with the nonstandard reading of (8) and (9)? The nonstandard reading demands that two languages are involved in (8) and (9). To see this, consider that (9) cannot be read in one language throughout. The language used in *stating* the antecedent is not the same as the language *specified* in the antecedent. The language specified is a language in which ‘tail’ refers to legs. However the language used in *stating* the antecedent cannot be the same language as the specified language since the *act* of specifying the language occurs in a language in which ‘tail’ refers to tails. Imagine a speaker of the specified language. The speaker’s utterance of ‘horses have four tails’ is true. However the speaker would not utter ‘I speak a language in which ‘tail’

refers to legs'.<sup>12</sup> Such an utterance would be false if uttered in the specified language (unless he was explaining his language to a speaker of ordinary English *in English*). To see that it would be false, consider what 'leg' means in the specified language. Likewise, a vat-English speaker's utterance of the consequent of (8) is true. But while vat-English is the language which the antecedent of (8) *specifies*, it is not the language which *states* the antecedent. The language used in stating the antecedent of (8) is English. That is because English (not vat-English) can refer to BIVs.

Yet (8) cannot be half in English and half in vat-English. One cannot speak both English and vat-English (though presumably one could be bilingual in English and the language specified in (9)). One is, after all, either a BIV or not. One either speaks vat-English or English (but not both!). With this in mind (8) must be false, since if we take the metalanguage used in stating (8) to be the same throughout it provides the incorrect truth conditions of the utterance in question. That is because, on the standard reading, it is *not* the case that if I am a BIV my utterances of 'I am a BIV' are true if I am a BIV (just as, on the standard reading of (9), horses have only one tail regardless of the language I might be speaking).

Since I cannot be bilingual in both English and vat-English Brueckner suggests (2) be understood as an English description of the truth conditions of an utterance in vat-English. But according to Brueckner I cannot assume that the argument is in English. That begs the question. For if I know that the argument is in English then I know that I am not a BIV, since if I were a BIV I would be speaking vat-English. The argument is

---

<sup>12</sup> I assume that in the language specified in the antecedent of (9), 'leg' either means nothing at all or it means what *we* mean by 'tail.' If the latter case is right then still the only way for a speaker's utterance of 'I speak a language in which 'tail' refers to legs' to be true is for it to *not* be spoken in that very language.

supposed to prove that I am not a BIV, not conceal that fact within its premises. I must be neutral with regard to whichever language the argument is in.

Brueckner recognizes that this, however, does not threaten the soundness of the argument. Though I cannot know whether or not the argument is in English or vat-English, I can know that whichever language it turns out to be in it is a sound argument. Regardless of whichever language I am speaking, 'I am a BIV' is false (159).<sup>13</sup> Whoever is reading E can conclude that their utterances of 'I am a BIV' are false. I know that, without considering which language I am speaking, my utterances of 'I am a BIV' are false. Brueckner returns to (T). If my utterances of 'I am a BIV' were true I would be a BIV. Since my utterances are false I am not a BIV. I know that I am not a BIV (164).

However, Brueckner does not believe that this accomplishes much. Since I do not know which argument it is that I am expressing when I utter E, I do not know what the conclusion of E means. I merely know that a sentence, 'I am a BIV,' is false (165). He says:

“Even though I know that my sentence 'I am a BIV' expresses a false proposition, I do not know whether or not it expresses the skeptic's counterpossibility proposition. If I do not know whether or not the sentence expresses that problematic proposition, then our anti-skeptical argument has not enabled me to conclude that I know that I am not a BIV” (166).

If this is right, it is difficult to see what the value of the anti-skeptical argument is. I do not know whether or not the argument counters the skeptical counterpossibility. It *could* counter that assertion, but I don't know that it does. For all I know, my utterances of 'I am not a BIV' are in vat-English. If that is the case then I am expressing something

---

<sup>13</sup> To confirm this Brueckner translates E into E', a parallel argument which expresses in English what a vat-English speaker would express were they to read E (160). I do not wish to question whether or not this move is available to Brueckner here as my retort does not depend on its being available.

which does not counter the skeptic's assertion at all. Since a BIV and non-BIV mean different false propositions through their utterances of 'I am a BIV,' all I can know for sure is that the proposition I express by 'I am a BIV' is false. *Which* proposition that is I cannot know.

### 1.3

The sort of argument that Brueckner defends requires a reading of Putnam which fails to take its central lesson seriously. That lesson is that BIVs are not capable of referring to 'brains' or 'vats'. Brueckner grants this conclusion. He does so when he acknowledges that the BIV and the non-BIV each mean something different by 'I am a BIV.' "If I were a BIV, the proposition that my sentence 'I am a BIV' would express would be false, and, if I were not a BIV, the *different* proposition that my sentence 'I am a BIV' would express in that case would be false as well" (Brueckner, 153). This follows from Putnam's argument. However to suggest that both of these propositions are candidates for the meaning of *my* utterances of 'I am a BIV' is to make a critical mistake. It makes no sense to provide alternatives for what our words refer to. One cannot do so in words, and one cannot think such thoughts. Brueckner is wrong in thinking that the "central point" (152) of Putnam's argument is that a BIV and non-BIV express different false propositions by the same utterance. What is most important for Putnam's argument, rather, is that a BIV's language, vat-English, is not capable of referring to BIVs.

Returning to Brueckner's discussion of premises (2) and (8) reveals the fundamental difficulty in thinking that *I* could be a BIV. (2) and (8) are two candidates for the way in which I can conceive of the skeptical counterpossibility as actually being the case. As Brueckner suggested, for (8) to be true it must be read in two languages. The

antecedent, which specifies that I speak vat-English, must be stated in English. However, as I discussed earlier, I cannot possibly speak English if I am a BIV. I cannot specify vat-English as the language that I speak if *that very language* lacks the capacity to make such a specification. We cannot see the radical counterpossibility as true in same way that we see (8) as being true. The radical counterpossibility, the worry that I could have always been a BIV, specifies a state of affairs in which *that very worry* is not specifiable. I cannot simultaneously worry that I am a BIV *and* be a BIV, just as I cannot specify that I speak vat-English *and* actually speak it. (2), then, seems to be the best candidate for understanding the truth of the counterpossibility.

What I want to argue, however, is that (2) is merely a counterfactual consideration. It is merely what I *would* have meant by an utterance if I *had been* a BIV. Its antecedent certainly appears to be concerned with what is actually the case. That is because it is pitched in present tense. ‘I am...’ appears to relate to my current actual situation, not to the past or to my situation at some other possible world. But really it does not. The ‘I am...’ is used in the same way as ‘If I were...’ or ‘If I had been...’ It is only concerned with those other worlds in which I am a BIV. In being concerned with those other worlds, it is not concerned with this world. At best, then, the skeptical counterpossibility is merely a statement about what could have been the case. If that is all it is, however, it does not succeed in threatening my knowledge of the external world. The fact that I am a BIV at some other world does not threaten my knowledge regarding this world.

The trouble with the skeptic's counterpossibility is that its intelligibility unravels its actually being the case.<sup>14</sup> Being able to refer to the skeptical counterpossibility demands that it in fact be false. The skeptic, while presenting such a counterpossibility to me, demonstrates that I am not a BIV. The story is only intelligible to a non-BIV, in the same way that trees are only intelligible to those who have a relevant sort of causal connection to them. To suppose that a BIV can worry that it is a BIV or know what it means to be a BIV is equivalent to believing that the inhabitant that lacks a causal connection to trees can still somehow form a worry about trees. Yes, my being a BIV is true in certain worlds. 'I am a non-BIV' is a contingent truth. There is nothing illogical about such a possibility. It breaks no natural laws. I can coherently entertain what I *would* mean by various utterances if I *were* a BIV. It is not unreasonable to suppose that: If I *were* a BIV, my utterances of 'I am a BIV' would be true iff I had sense impressions of being a BIV. Musing about what could be the case at other worlds, however, is not the focus of the skeptical strategy. The skeptical worry is not that I could have been a BIV, or that I am at other worlds. It is that *I could be* a BIV. Any statement of the form: 'If *I am* a BIV...' is bound to be confused so long as it is taken as a statement about *this* world. Yet that is exactly what the skeptic believes herself to be doing. Her worry is not some sort of counterfactual one, but one about *this* situation. Such a worry is incoherent.

I cannot take seriously the notion that, of the language that I am speaking: Either my utterances of 'I am a BIV' mean 'I am a BIV' or 'I have sense impressions of being a BIV.' Only a confused non-BIV can make this sort of statement. I can talk about what a BIV would mean by the utterances that I make, but I cannot coherently suggest that *my*

---

<sup>14</sup> I can present an intelligible possibility so long as I can refer to those elements which it describes. 'The book is on the table' is intelligible to me because I can refer to books and tables. The possibility is intelligible to you so long as you occupy the needed causal relationships as well.

utterances mean *that*. Here it is tempting to exclaim: “Yes, but a BIV could say exactly *that!*” Such is not the case! A BIV could utter those same words but would mean something entirely different. The lesson to take away from Putnam’s argument is not: my meanings (whichever ones they are) of ‘I am a BIV’ must be false. It is, rather, in order for the skeptic’s counterpossibility to be an intelligible one it must not actually be the case.

My focus here concerns reasonable skepticism. The reasonable skeptic presents an argument which supports a skeptical worry. It provides a reason to endorse the skeptical worry. The skeptical worry involving the radical counterpossibility is unreasonable. After considering Putnam’s views on reference it is clear that its being the case is incompatible with its being intelligible. The radical counterpossibility is intelligible (and the skeptic who aims to convince us her worry is reasonable better hope that it is intelligible!) and so is not the case. What, then, is the skeptic worried about? No sense can be made of the worry; there is no reason to take the worry seriously.

Putnam describes the radical counterpossibility as being self-refuting. “If we can consider whether it is true or false, then it is not true (I shall show). Hence it is not true” (8). This follows from the BIV’s inability to refer to BIVs. The following is my extraction of Putnam’s main point from his discussion:

- A. Either I am a BIV or I am not a BIV
- B. A BIV cannot refer to BIVs
- C. I can refer to BIVs
- D. I am not a BIV

The argument is simple. What is important is that when the BIV thinks ‘I am a BIV,’ whatever that thought is about (it can be about nothing, if you like) it is not about BIVs. It does not matter at all what the BIV refers to, so long as it cannot refer to ‘brains’ or ‘vats.’ If a BIV cannot refer to BIVs then a BIV cannot entertain the skeptical counterpossibility. A BIV cannot worry that it is a BIV. A BIV, then, cannot refer to what I can refer to. Since I am capable of referring to the skeptical counterpossibility I am not a BIV.

Brueckner grants the truth of (A) and (B). (C) is the most vulnerable premise in the reconstruction. Brueckner attacks (C) when he introduces his form of semantic skepticism. He argues that although I know that ‘I am not a BIV’ is false, I do not know what it means. I do not know what I refer to by ‘BIV.’ For all I know, according to this skeptical approach, I cannot refer to BIVs.

Is the worry reasonable? I do not believe that it is. Imagine the semantic skeptic saying: “You don’t know whether or not, by your utterance ‘dog,’ you are referring to dogs or sense impressions of dogs.” How can I be worried that I mean ‘sense impressions of dogs’ when I utter ‘dog?’ The worry consists in my use of ‘dog’ not actually meaning dog. But how can I capture that very worry in that same language which the worry is concerned with? Recall the trouble we encountered with respect to premise (8) which attempted to state the truth conditions of ‘I am a BIV’ in two languages. The person who uses the word ‘dog’ to refer something other than dogs cannot refer to dogs at all! So how can I worry that I’m not referring to dogs? The worry *specifies* a language (one which does not refer to BIVs) while *stating* that which the specified language cannot (BIVs).

The argument suffers from the same difficulty that the initial skeptical counterpossibility encountered. If I can express the semantic worry, the worry that I cannot refer to BIVs, then the worry has been defused. Does the worry take the form: Either my utterances of 'X' are referring to Y's or referring to Z's? If so, the capacity to have such a worry demands that both Y's and Z's can be referred to. The worry that I cannot refer to BIVs demands that I can, in fact, refer to BIVs! It is hard to see how the skeptic can attack (C). It is outrageous for me to think that I cannot refer to BIVs. Such a thought refers to what it claims I cannot refer to. It is also outrageous for me to think that I could be referring to either BIVs or sense impressions of BIVs by my utterances of 'I am a BIV.' For if I am able to consider such possibilities then I can in fact refer to *both* BIVs and sense impressions of BIVs. Nor can the skeptic suggest merely that I may refer to sense impressions of BIVs (as opposed to suggesting that I might not refer to BIVs), for such a suggestion, as shown to be the case with regard to (2), must take place in English. 'Sense impressions of BIVs' is a translation of what a BIV means by 'BIV' *in English*. If a BIV refers to sense impressions of BIVs (perhaps by uttering 'BIV') then it does not do so by uttering 'sense impressions of BIVs'. That is how a non-BIV refers to those sense impressions. How can the skeptic attack (3) without the attack contradicting herself? The worry appears to be incoherent.

Perhaps rather than providing an argument against (3) the skeptic instead might ask me to support the premise. "How do you know that (C) is true?" the skeptic might ask. A careful skeptic, at this point, must not provide candidates for what I could mean by 'BIV.' As noted above, such suggestions have an incoherency built into them. Instead the skeptic must operate solely from the confines of a vague worry. It is vague in that there is

no target of worry. Usually when one worries, she worries that something might be the case, either now or in the future. But here that is not so. Rather the worry is concerned with some potential yet unspecifiable state of affairs. It is a state of affairs in which I don't refer to BIVs.

In order to see just how peculiar this challenge is one must avoid seeing it as following from a sort of skeptical regress argument. For example, one could challenge premise (B) by asking: "How do you know?" I would then return to Putnam's analysis on what a BIV is incapable of referring to and why. I can imagine the follow up question: "And how do you know that is right?" At this point it would be foolish for me to try to explain how it is that I know an argument to be a good one, since that too would surely face the same sort of skeptical retort. But this is not what is occurring with regard to the challenge to (C) mentioned above. Here, instead, I am stunned by the absurd suggestion that I might not refer to BIVs, or that I may not refer to anything at all. It is tempting to exclaim: "BIV! See? I just did it!" What else is there to say? Bredo Johnsen (2003) captures the dilemma clearly: "Attempting to treat it as an open question whether 'our language' is meaningful is fraught with mind-numbing difficulties – In what language, or manner, should we imagine ourselves trying to address it, or to determine whether it can be addressed, or ...?" (Bredo, 228).

The skeptic cannot be suggesting that I might, instead of referring to the counterpossibility, *not* refer to the counterpossibility. The worry cannot be framed as an inability to distinguish between various possibilities (me referring to BIVs or me not referring to BIVs). What then is the worry exactly? Is it possible to capture exactly what the worry is getting at? I do not think so. There is not a single situation I could be in that

would contradict (C) and which I cannot rule out in the ways noted above. The support for (C) then comes from the fact that no other possibility can be coherently entertained. The self-undermining nature of these semantic skeptical retorts, then, is not avoided by the skeptic when she presses them in her language vaguely. Why should I have any reason to think that I cannot refer to BIVs? A vague worry is a worry without a reason. These are either reasonable worries or they are not, and the less detailed these possibilities are the less they say anything at all.<sup>15</sup>

The skeptical argument utilizing the radical counterpossibility is ultimately self-undermining. If the counterpossibility is intelligible then it reveals itself to be false. Additionally, one cannot reasonably worry that they do not refer to the skeptical counterpossibility. Such a worry involves reference to that which is framed by the worry as something one cannot refer to. It is, then, a worry that a BIV cannot have. These worries are incoherent. Furthermore, the skeptic cannot hope to avoid this problem by presenting a vague worry. A vague worry is an unjustified skeptical itch. It is nothing more than a background vibe which motivates the skeptic toward states of uncertainty.

## 1.4

Is there a counterpossibility which the skeptic can utilize to deflect the charge that her argument is incoherent or vague? I will now sketch out an answer to this question.

The change that must be made to the skeptic's counterpossibility involves its being seen as a reasonable worry. The counterpossibility must be one which is potentially true of *this*

---

<sup>15</sup> It is important to keep in mind that my goal here is to respond to the skeptical argument. For that reason it is a mistake to suggest that I beg the question against the skeptic through presupposing that my language refers (though I do wonder what would involve *not* presupposing such a fact). A meaningful language (specifically a language capable of reference) is necessary for the skeptic to formulate her counterpossibility in the first place. So a skeptic cannot challenge me by referring to possible state of affairs in which reference is impossible! Bernhard Weiss (2000) makes this same observation (p. 115).

world. Its actually being the case must not threaten its intelligibility. To stick to the theme of envatted brains: the counterpossibility must be intelligible to a BIV. A BIV must be able to worry that it is a BIV. This means that the details of the counterpossibility must allow for the BIV to refer to brains and vats.

The counterpossibility offered by the skeptic must be less radical in order to account for this change. The most obvious way to go about this is to introduce an envattment procedure. No longer has the BIV in the skeptical story been permanently envatted. Let  $t_1$  be prior to envattment. At this time, Sam, I will call the individual, has not yet been envatted. At  $t_1$  Sam has the sort of knowledge of the external world that we typically assume we have. We can suppose that at  $t_1$  Sam has mostly true beliefs about the world. Let  $t_2$  be directly after the envattment procedure. The important feature of this weakened counterpossibility is that at  $t_2$  Sam can refer to BIVs despite the fact that he himself is a BIV. If it is tempting to think that reference would shift quickly (to the meanings found in the radical counterpossibility) then we can imagine that  $t_2$  is the instant after envattment. At  $t_2$ , then, when Sam utters 'I am a BIV' the utterance is *true*.

This is all assuming that the act of envattment does not bring with it an instantaneous change in reference. I am currently assuming that the meanings of a BIV's utterances do not instantly shift the moment envattment has occurred. I think there is a good reason to assume that they would not. That is because, really, the act of envattment has a great deal in common with transportation. It is rather like being transported to a place (against your knowing) which seems to be the previous location but really is not. In such a circumstance it is difficult to see why one would be immediately incapable of referring back to the features of the previous location. Why would *the act* of being moved

change the meanings of one's words? Additionally I can talk about past occurrences which no longer hold in the present, and I can talk about objects that used to exist but no longer do (and in that sense have been moved away from me). And so it seems to me that a BIV (of the sort described by my weakened counterpossibility) could refer to BIVs despite not being in a situation in which it can point one out. Here it is worth considering what Putnam says about the sort of relationship one needs to have with objects in order to refer to them:

“Our talk of apples and fields is intimately connected with our nonverbal transactions with apples and fields. There are ‘language entry rules’ which take us from experiences of apples to such utterances as ‘I see an apple’, and ‘language exit rules’ which take us from decisions expressed in linguistic form (‘I am going to buy some apples’) to actions other than speaking. Lacking either language entry rules or language exit rules, there is no reason to regard the conversation... as more than syntactic play. Syntactic play that resembles intelligent discourse, to be sure; but only as (and no more than) the ant’s curve resembles a biting caricature” (11).

As a statement meant to provide a reason to reject the possibility of reference without causal connection the above quotation is surely right. However it is obvious that I can refer to apples without first seeing one or proceeding to seek one out. If it is true that my talk of apples *always* shares an intimate connection to my interactions with apples it is not at all obvious what sort of interactions those are. Am I interacting with apples when I am reminded of one after seeing a red book? Since one does not, before *every* act of reference, need to see or hear about a brain or vat in order to go on to refer to brains and vats, it is hard to see how the fact that an envatted individual will never again interact with brains or vats influences his *immediate* capacity to refer to them. For, imagine that after the instant that is  $t_2$  an un-envattment procedure were to occur. Surely we (along with Sam) would say that, at  $t_2$ , Sam’s thought ‘This is a tree’ was false. Then, at  $t_2$ ,

Sam's thought 'I am a BIV' is true. It is just the same way when we consider our own dreams. While unknowingly dreaming my utterances of 'I see a clown wielding a chainsaw' are false. My utterances would not be true due to 'clown' referring to dream-clowns and 'chainsaw' referring to dream-chainsaws.<sup>16</sup> What this indicates is that at  $t_2$ , which is at least a moment in time, Sam can refer to brains and vats. It is worth considering that if I am wrong, if with the act of envattment goes the capacity to refer to brains and vats, one arrives at a more dramatic anti-skeptical conclusion than the one I am currently providing. Such a fact would render all varieties of BIV counterpossibility unreasonable.

This weakened counterpossibility's utility in a skeptical argument, however, is not as straightforward as it might first appear. The skeptic cannot abandon the radical counterpossibility in favor of a weakened counterpossibility without a significant tradeoff. For example: Sam's belief at  $t_2$  that trees have branches is true. Furthermore, when Sam expresses at  $t_2$  the beliefs he had at  $t_1$  he does so correctly. Sam utterances of 'At  $t_1$  I believed that trees have branches' at  $t_2$  are true.<sup>17</sup> Just as Sam's utterance 'I am a BIV' at  $t_2$  is true, so is his utterance 'trees have branches' at  $t_2$  also true. Recall that, with regard to the radical counterpossibility, it first appeared that these sorts of statements would

---

<sup>16</sup> Such utterances would be true for that reason, however, if the dream lasted for several years (or however long it takes for reference to shift). In such a situation my utterances of 'clown' and 'chainsaw' would be intimately connected with dream-clowns and dream-chainsaws.

<sup>17</sup> This is because the language Sam uses at  $t_2$  to express his beliefs at  $t_1$  is the same as his language at  $t_1$ . Were these languages not the same Sam would express his past beliefs incorrectly. Crispin Wright (1992) provides a weakened counterpossibility in which such is the case. Let  $t_3$  occur long enough after the envattment procedure that Sam's meanings have shifted such that his utterances involving 'tree' refer to, say, sense impressions of trees. Sam's utterances of 'At  $t_1$  I believed that trees have branches' at  $t_3$  are false, since at  $t_1$  Sam's belief was about trees, not sense impressions of trees. The trouble with Wright's weakened counterpossibility is that a skeptical argument which utilizes it is vulnerable to the same charge I made against the skeptical argument which utilizes a radical counterpossibility. At  $t_3$  Sam cannot refer to brains or vats. So at  $t_3$  Sam cannot entertain Wright's weakened counterpossibility. I am not, then, a BIV of the sort Sam is at  $t_3$ .

come out false. Here we see that when the skeptic constructs a scenario in which a Sam *can* refer to BIVs he also can refer to anything else in the real world which he previously referred to. This means that, if the envattment procedure does not co-occur with a radical change to the world, the true claims Sam made at  $t_1$  about the world remain true if made at  $t_2$ . It appears that the weakened counterpossibility has built in a lot less skeptical power than the skeptic may have initially hoped.

There is, though, good reason to suppose that some of Sam's beliefs at  $t_2$  are threatened by his situation. If Sam remains capable of referring to the objects he referred to at  $t_1$  it is not unreasonable to think that when Sam encounters what he *supposes* to be a tree at  $t_2$  and utters 'This is a tree' his utterance is false.<sup>18</sup> It is true that over time reference to vat-objects would become fixed due to Sam's "nonverbal transactions" with them. However, assuming that this process takes time, the skeptic can always assert that, for all I know, I could be a BIV of the sort Sam is at  $t_2$ . If that is the case, then my belief that I am typing on a computer is likely false. But, again, it is only false in virtue of my capacity to refer to actual computers. All of my knowledge regarding what computers are and how they work remains intact.

What this shows is that when the skeptic abandons the radical counterpossibility in favor of a weakened counterpossibility her skeptical argument is considerably less potent. Sam can, seemingly, have true beliefs at  $t_2$  *even about the present* so long as they are rooted to his meanings at  $t_1$ . Sam's envattment does not instantaneously prevent him from referring to the external world. At  $t_2$  Sam can refer to brains, vats (and so the

---

<sup>18</sup> Perhaps whether or not his beliefs at  $t_2$  are true or false depends a great deal on why they are being entertained in the first place. For example, we may consider Sam's utterances at  $t_2$  of 'There is a tree in my yard' to be true, if there is a tree in Sam's yard. And yet if by 'my yard' he is speaking of *that* place in the vat-World, immediately in front of him, then *what is there* is not a tree at all. His utterances, in that case it seems, would be false.

skeptical argument which utilizes it is neither vague nor incoherent) trees and computers. Sam's utterances at  $t_2$  of 'That is a tree' and 'That is a computer' are false. It appears, then, that the sort of knowledge threatened by the counterpossibility is that which concerns immediate experiences of the world. At  $t_2$  Sam is unable to correctly identify, for example, what he is looking at. Knowledge not concerned with immediate experience, on the other hand, does not appear threatened by the counterpossibility. The counterpossibility does not threaten knowledge about what computers are nor does it threaten knowledge about where Seattle is located. If Sam had such knowledge at  $t_1$  he has it at  $t_2$ .

Recall just how threatening the radical counterpossibility appeared. It seemed to threaten not only knowledge concerning experiences of the external world, but general knowledge about what the external world is as well. Seattle does not exist if the radical counterpossibility is actually the case. All that exists are brains, a vat and automatic machinery. So although the skeptic's use of the weakened counterpossibility escapes my previous arguments she does not get it for free. In utilizing the weakened counterpossibility her argument is less threatening; if I am a BIV of the sort Sam is at  $t_2$  then I know a great deal about the external world (assuming that I did at  $t_1$ ).

## 2. Conceptual Deficiency

I hope I have adequately captured the difficulties which arise for the skeptic when she suggests that the radical counterpossibility could actually be the case. In order for the skeptic to avoid these difficulties the counterpossibility provided must accommodate Putnam's views on reference. It must be a story expressible by its own characters. If its characters are BIVs, and if it is suggested that I could be such a character, then it must be expressible by BIVs. I have provided an example of a counterpossibility which meets this demand. Swapping counterpossibilities, I have argued, is not free. When the skeptic abandons the radical counterpossibility in favor of a weaker counterpossibility the range of her argument narrows. That is because the details of the counterpossibility determine just what knowledge is under attack. The failure of the radical counterpossibility, then, reveals something interesting about the necessary conditions for falsehood. In order for a belief about the world to be false it must *be about* that which it gets wrong. But for a belief to be about something its believer must have acquired that *concept*.

### 2.1

A concept, according to Putnam, cannot be a word or a mental image. Remember that Putnam's thought experiments illustrate the difference between referring and merely *seeming* to refer. The inhabitant's mental picture does not refer to trees even if it is exactly like mine when I think about trees. Since I can imagine my internal mental life as being identical to the inhabitant's concepts are not things we perceive through introspection (Putnam, 17). Having the concept necessary to refer to trees involves a causal relationship with trees. To attribute a concept to a person is to attribute to them the *ability* "to use sentences in situationally appropriate ways" (19). It is to say that one is

*sensitive* to those features of the world which his meanings are about (11). Reference demands more than making certain sorts of noises with one's mouth. One must make certain noises in response to his environment, and one must interact with his environment after making certain noises. This is what it means to have language entry and exit rules. We can imagine a machine designed to fool others into thinking that it refers to features of the world (9). It could churn out those same English sentences we produce when we talk about the world. But it would only be fooling us if it lacked sense organs. Regardless of the noises it produces it cannot refer to apples if it could not notice one placed before it (10). If it would continue to make those same noises when all the apples disappeared, or if there had never been any apples to begin with, then it is not *sensitive* in any way to the existence of what it seems to be talking about (11). Noises which *do* refer to apples precede or proceed after non-verbal interactions with apples. Those who are sensitive to apples in such a way can be said to have the concept required to refer to them.

One cannot be wrong about some feature of the world if they entirely lack the concept required to refer to it. Whether one has such an ability depends on how successful they are both at using sentences at the appropriate times and acting appropriately after using sentences. Lacking language entry and exit rules *entirely* is the difference between some degree of sensitivity to a feature of the world and none at all. If I occasionally mistake a pear for an apple I have the concepts (relating to apples and pears) which I need to make that mistake. Such cases would show that I am *slightly* insensitive to apples and pears to the extent that error arises. But if one is insensitive enough error altogether no longer applies. It is this point which is illustrated by the difference between the radical and weaker counterpossibility. Pack too much of what

appears to be skeptical power into the argument and suddenly the entire project fails. If referring to features of the world demands an ability so does getting the world wrong. Permanent envatment, which initially appeared to deliver permanent undetectable error, severs the possibility of error altogether.<sup>19</sup> A skeptic cannot threaten *all* of our knowledge of the external world. It cannot be that all of our beliefs about the external world are false because falsehood demands some degree of sensitivity and radical falsehood is a complete lack of sensitivity.

As is far more obvious in Putnam's stating of his argument than my own presentation of it, this is a conclusion which equally applies to the BIV.<sup>20</sup> Whatever the BIV means by 'tree', "the brain is *right*, not *wrong* in thinking 'There is a tree in front of me'" (14). The BIV lacks the concepts required to get such things wrong. Thus according to Putnam's argument it might be tempting to think that we were wrong to worry about the skeptical argument on two counts. Firstly it was wrong to think that the radical counterpossibility could actually be the case. The suggestion that I am a BIV is unreasonable. The second temptation is to think that it was wrong to *worry* that it might be the case (even if we were to grant that such a suggestion is reasonable). Not only do we lack the sort of massive error the skeptic envisioned, but equally so does the BIV since it lacks the required concepts. If skepticism is troubling because the prospect of massive delusion is a horror, being a BIV no longer seems so horrifying. It is this second temptation which is flawed. The BIV's epistemic state is undesirable despite the fact that it is not massively in error about the world.

---

<sup>19</sup> Or severs the possibility of *skeptical* error. A permanently envatted BIV can go on to be wrong about those matters which it is sensitive to (whether they are objects of the image or computer programs or whatever).

<sup>20</sup> Here I restrict my usage to only include those permanently envatted brains.

My goal here is to articulate precisely why the BIV suffers from a sort of epistemic misfortune and how that reveals another way for the skeptic to respond to the argument made in the first chapter of this paper. It is the skeptical response illustrated by Graeme Forbes (1995) when he says: “a classical skeptic, however, can still get mileage out of the thought that we are perhaps relevantly like brains in vats... While our inability to rule out the hypothesis that we are relevantly like them does not threaten to show that we have no knowledge, it does threaten to show that we have no assurance that knowledge is worth having: it grants us the capacity to acquire it but destroys its value” (Forbes, 222).

The skeptic, I will argue, can provide the possibility of conceptual deficiency, rather than radical delusion. I will first show that the response is compatible with the skeptic’s aims. Skepticism is threatening because it argues that we lack knowledge. To put it another way, we value knowing and skepticism threatens to show that we do not know. But what does it mean to say that we value knowledge? In what follows I will present two types of worry which external world skepticism supports. These worries are responses to the value we place on knowledge of the external world.

## 2.2

Since knowledge requires true belief, what does it mean to say that we value believing the truth? One possible answer is:

(BC) Believing  $p$  is valuable if and only if  $p$  is true

The biconditional is meant to capture our valuing believing the truth. Though it is a useful starting point, I will argue that (BC) does not adequately capture what we mean when we say that we value truth. Both halves of the biconditional demand refinements in

order to successfully account for our valuing truth. The two halves of the biconditional are:

(C1) If believing  $p$  is valuable, then  $p$  is true

(C2) If  $p$  is true, then believing  $p$  is valuable

It is worth mentioning here that the value I am concerned with is epistemic value. Jason Baehr (2012) says that “for something to be epistemically valuable is for it to be valuable from the standpoint of a good intellectual life, or from the standpoint of what one might think of as intellectual or cognitive flourishing” (Baehr, p. 4). The idea is that in addition to having desires or values relating to other matters (such as the desire to be healthy) there are desires and values relating specifically to the formation of beliefs. These domains of value can cross paths. Surely there are times when it is good to believe something which is false. If I am a soldier fighting in an unwinnable battle I might be better off with unrealistic optimism than realistic dread. Additionally, if believing some truth will lead to my assassination then it appears that believing it is bad. These possibilities are not counterexamples to (BC), but rather, are clear cases in which our non-epistemic interests outweigh our epistemic interests. The fact that the value I place on being alive outweighs the value I place on knowing the truth in certain circumstances does not mean that I don’t value knowing the truth at all (Kvanvig 2008, 201).

Furthermore a desire to believe the truth is not the same as a desire for what I believe to be true.<sup>21</sup> (BC) says that if  $p$  is true then it is good to believe it (from an epistemic point of view) and if  $p$  is not true it is not good to believe it. (BC) does not suggest that if I believe a loved one to have cancer then I want it to be the truth. Rather it suggests a desire that one’s beliefs match the way the world really is.

---

<sup>21</sup> See Piller, 2009, p. 195, Sosa, 2003, p. 50, and Whiting, 2013, p. 224.

(C1) says that we don't value believing falsehoods.<sup>22</sup> (C2) says that we value believing truths. Upon reflection it is obvious why two conditionals are needed. Being epistemically well off demands both believing truths and avoiding false beliefs. If one accepted (C1) but not (C2) then a person who believed nothing at all would be as epistemically well off as somebody who believed both no falsehoods *and* a great deal of truths. But that's obviously not right. Accepting only (C2), on the other hand, betrays our intuition against epistemic carelessness. Compare the careful believer to the careless believer. We think that it is better to believe a number of truths and no falsehoods than to believe that same number of truths *and* a great deal of falsehoods. We already see, however, that (C1) fails to account for this. That is because, according to (C1), believing falsehoods is merely not-valuable. We think that believing falsehoods is *bad* rather than merely *not-good*. (C1) accounts for our thinking that the careless believer isn't better off than the careful believer, but it fails to account for our thinking that the careless believer is *worse* off than the careful believer. Or compare a person who believes neither *p* nor *not-p* to a person who believes both *p* and *not-p*. Imagine *p* is true. We might be inclined to think that the latter individual is worse off epistemically than the former (Whiting 2013, 238). But even if we think that believing a truth is more valuable than believing a contradiction is disvaluable (which is hard to imagine), it would obviously be better to believe solely *p* than both *p* and *not-p*. (C1) *merely* claims that believing both *p* and *not-p* is not better than believing just *p*. In actuality it is *worse* than believing just *p*.

(C1) should be replaced with (C1\*):

(C1\*) If *p* is false, believing *p* is disvaluable

---

<sup>22</sup> The contrapositive of (C1) is actually:  
 (CP-C1) If *p* is not true, then believing *p* is not valuable  
 I am assuming presently that if *p* is not true then *p* is false.

(C1\*) accounts for why we find, at first glance, the radical counterpossibility to be so horrifying. If it is bad to believe falsehoods then it is supremely bad to be massively deluded about the nature of the world. This skeptical worry which I have dealt with thus far is a skeptical response to (C1\*). If (C1\*) was sufficient for explaining what we value epistemically (which I have suggested above, is not the case) then it would appear that Putnam's considerations show us that that the BIV is not in an epistemically poor state. That is because, if we translate (C1\*) into the goal of avoiding falsehoods, there is no reason to think that the BIV must be less capable of satisfying the goal than the non-BIV is. But despite this it is a bad thing to be a BIV. To see this one must consider (C2).

(C2) says that if something is true we value its being believed. According to (C2) the more true beliefs the better. There is something clearly wrong about this. We do not want to say that an effective way to improve one's epistemic state is to search high and low for any truth at all. To borrow an example from Ernest Sosa: "Here is a trove of facts, of the form *grain x is so many millimeters in direction D from grain y*, than which few can be of less interest" (Sosa 2000, p. 49). It cannot be merely a matter of quantity. We don't think that the best epistemic course of action is to collect as many true beliefs about grains of sand as we can. It seems as though the epistemic value of believing a truth is not independent of what the truth is about.<sup>23</sup>

To extend Sosa's example, imagine two men who are specifically concerned with acquiring truths regarding the external world (suppose prior to their respective

---

<sup>23</sup> Kvanvig (2003, 2008) would dispute this conclusion. He believes that ultimately all truths are of equal epistemic value and that those which are useful to us are prioritized over those which are not. Here I am not concerned with *just how* epistemically valuable certain truths are. I am not interested in whether or not truths about grains of sand have little epistemic value or lack epistemic value entirely. Rather, I am concerned with whether or not our interest in acquiring truths is just an interest in acquiring as many truths as possible. I will argue that this is not the case, that mere quantity of truth is not good enough. It is not obvious that Kvanvig would disagree with this.

investigations they know little about the world). The two men have very different strategies. The first man's strategy is to organize an adventure across the globe. He plans on visiting a variety of different locations and interviewing a large number of experts in various scientific fields. These include astronomers, biologists, physicists and social scientists. The second man's strategy is to study grains of sand in his son's sandbox out in the yard. Without needing to travel anywhere he can devote as much time as possible toward acquiring truths. While the first man is waiting for his airplane to land in Egypt (suppose he chooses not to carefully study the flight safety catalogue) the second man is diligently counting and measuring grains of sand and acquiring yet more truths. The second man will acquire more truths than the first. Despite knowing fewer truths about the external world it is apparent that the first man, more so than the second, gets at what is meant by our valuing knowledge of the external world. This shows that when it comes to valuing knowledge of the external world we are not interested in raw *quantity* of truth but in knowledge of what sorts of truths there are. Though the second man will have a small understanding of, say, the cosmos, if he has performed his investigation properly he will *know* that. Despite having an incomplete picture of the world, it is an incompleteness that *he has a sense of*. The man who acquires a great number of truths regarding his sandbox does not have any grasp of the limits of his understanding (assuming he is ignorant about what else there is to know about the external world).

How should (C2) be refined considering that our desire to acquire truth is not a matter of quantity? Recall that ultimately I am concerned with establishing those skeptical worries which respond to our epistemic goals. The above story illustrates that we want a sense of what sorts of truths there are. Regardless of how many truths about

the external world one has acquired it is desirable to know their place relative to what there is to know. If the man who investigates the sandbox thinks that he is getting at the external world in a particularly valuable way he is wrong. The skeptical worry pertaining to our desire to acquire valuable truths is the worry that we are cut off from the valuable truths which there are. It is the worry that my knowledge of the external world resembles that of the man who acquires large quantity of truths regarding his sandbox.

### 2.3

I said that skepticism is threatening because we value believing the truth. If we did not care about the truth we would not care about the skeptic's argument. We would be happy to concede the skeptical premise that we lack knowledge. I am not happy to concede that. That is because I do care whether or not I know the truth. Specifically, I value my knowledge regarding the external world. Not only do I wish to avoid falsehoods, but additionally I would like to occupy a location from which I am not clueless with respect to the sorts of truths there are. This can be put in terms of Putnam's views. Firstly, I want to be as sensitive as possible to those features of my environment which I am involved with in causal relationships. The better my ability to, say, point to an apple when I see one, the less often I mistake pears for apples. Secondly, I want to be the sort of individual and located in such a place such that I am capable of acquiring a diversity of concepts. Our epistemic preferences, then, concern:

1. The degree to which we are sensitive to objects
2. The range of our ability to form concepts

It is now clear why the BIV is an unfortunate sort of creature even if it completely lacks false beliefs. Though it is adequately sensitive to its environment (it makes certain

utterances when certain computer programs are activated or whatever) and thus has an ability to mostly 'get it right', it can't get *most* of it right. The BIV occupies a conceptually deficient location in the world. It lacks the concepts required to put its body of knowledge into perspective. Lacking falsehoods is not enough. It matters both *that* we can know and also *what* we can know. Though it might be that the BIV knows a large quantity of truths it does not have any sense of those truths which are forever out of reach due to its cognitive limitations and its environmental location.

The skeptical argument utilizing the radical counterpossibility aims to maximize the degree to which sensitivity and conceptual range are diminished. Yet radical falsehood, or radical insensitivity, is not possible. Falsehood requires concepts, and radical insensitivity suggests complete conceptual lacking. The weakened counterpossibility, which is one direction the skeptic can go in, is less threatening to both concerns. The individual that is envatted overnight is only significantly in error at  $t_2$  due of the range of concepts he had acquired by  $t_1$ . There is another way in which a skeptic might choose to balance the skeptical worry in light of our epistemic preferences. Returning to the radical counterpossibility, given the fact that radical falsehood is out of the question the skeptic might focus instead on the fact that the BIV is radically conceptually deficient. This skeptical worry acknowledges knower's ability to refer but threatens the range of the ability. It attacks *not* the fact that I know, but the extent of what I can know. It is the worry not that I am a BIV, but that I might be similarly conceptually deficient. For all I know there is an out of reach perspective which reveals me to be a severely cognitively deficient creature. Perhaps there are truths which are, in principle,

un-acquirable. Truths which, as Crispin Wright suggests, “would make a mockery of mankind and its place in nature” (Wright, 93).

It might be that a consequence of ‘semantic externalism’ is that any cognitive perspective is guaranteed some degree of conceptual incompleteness. Bernard Weiss states: “Semantic externalism sponsors the idea that the availability of a thought is (heavily) dependent on context. Thus it is likely that there is no context from which all thoughts are available. So any theory (that we or similar beings construct) will be incomplete since it will be bound by its context” (Weiss, 122). It is important to keep in mind that there are magnitudes of conceptual incompleteness. The existence of certain out of reach concepts, certain causal relationships I am unable to form, need not itself spell out doom and gloom. The ‘twin earth’ story (Putnam, 1975) provides an example of two sorts of conceptually similar sorts of creatures who, though we can imagine to be unable to refer to what each other means by ‘water’, are not to be considered conceptually deficient because of it (or at least not the sort of conceptual deficiency of a terrifying variety). Thus one might be able to make peace with the possibility of some degree of undetectable conceptual incompleteness. Still, if radical conceptual deficiency is a possibility it is too quick to celebrate the conclusion of Putnam’s anti-skeptical argument. Even a complete lack of falsehood is compatible with a serious skeptical worry.

Crispin Wright (1992) and Graeme Forbes (1995) make similar arguments. Both see the alternate skeptical strategy as responsible for the “persistent unpersuasiveness” (Forbes, 219) and “dissatisfaction many have felt” (Wright, 93) with regards to Putnam’s argument. “The trouble is that we easily slip into inept formulations of it” (Wright, 93). Wright describes the conceptually deficient individual as one who lacks the concepts

“which reflect the *real* kinds that there are” (91, italics mine), which “are on to the *right* categories...of the world” (93, italics mine). Is Wright’s usage meant to signal that the individual is unable to capture the predicament that he is in, that he lacks access to the *right* categories to do so, or is it that those categories he has access to aren’t the *right* or *real* categories that there are? Take the BIV as an example. It is correct to suggest that the BIV lacks the concepts necessary to capture its own epistemic situation. But it is wrong to suggest that the concepts which the BIV *does* have are not real, or that the BIV does not refer to real things.

In a response Putnam restates Wright’s skeptical worry: “we are so situated that we are unable to conceive of the true laws governing the most fundamental physical magnitudes” (Hale, 286). If the skeptical worry no longer concerns falsehood Putnam presumably means something like *actual* laws. His response to Wright is to demonstrate that the BIV is incapable of forming this worry. That is because the BIV lacks access to the notion of a “physical magnitude” in a way which non-BIVs do not. Just as is the case with brains and vats the BIV lacks the required causal relationships to form the concepts necessary for stating the skeptical worry. In fact, the BIV does not even have access to “causation” (287), so it cannot worry that it lacks concepts (or the appropriate causal relationships necessary to form concepts) to the degree which it does. Putnam appears to be suggesting that a transition from a specific possibility (being a BIV) to a *type* of possibility (being *like* a BIV) is either vague or incoherent. It is vague if all that the skeptic can say is that it is some sort of epistemically unfortunate state of affairs (and when pressured, cannot specify just what such a state of affairs might involve) and it is incoherent if the worry demands that it cannot be offered by the skeptic in the first place.

If the worry is that I am in a certain *type* of scenario (illustrated by the BIV's scenario) then its being expressible cannot hinge on its not being the case (otherwise it is an incoherent worry). As Putnam suggests, this limitation on the skeptic's argument does not go away when one transitions from specific scenarios to types of scenarios. But one *need not* understand the type of scenario as one whose characters are incapable of grasping it. It is in this sense, the way in which we are meant to see the use of a *type* of scenario, that Wright and Forbes offer different approaches. Wright suggests that Putnam's argument delivers a blow to any "specific description" (92) of the skeptical story, where specific descriptions are precisely those counterpossibilities "*which we can understand*" (92). Wright's concern is that Putnam's argument cannot dispose of the general type of counterpossibility, including the specific forms of which are unintelligible to us. But it is precisely the unintelligibility of the skeptical story which renders it unreasonable.

Forbes, on the other hand, sees the type of scenario as picked out by the "*contrast*" (219) between the BIV and non-BIV. It is a contrast between having access to certain sorts of truths and being precluded from acquiring them. Forbes says: "We cannot avoid raising the question for ourselves whether there may not be other instances of the same contrast differing precisely from the given one by having us occupy the cognitively deficient type of context, one whose agents are precluded in principle from knowing certain truths" (220). The specifics of the truths we would be precluded from knowing would be, of course, "beyond our grasp" (220), but the *notion* of un-acquirable truths would not be. It is the contrast illustrated by the two men, one of whom chooses to investigate the details of his sandbox. He lacks truths regarding the external world *and*,

since he is familiar with the same notions of ‘physical’ and ‘causation’ as the other man, can form a worry regarding his actual epistemic lacking.<sup>24</sup> We can adjust the thought experiment such that his investigating grains of sand is due to his cognitive and environmental situation (rather than by his foolish conception of what it means to value truth).

The BIV’s diminished conceptual scope is what provides justification for thinking that its epistemic status is undesirable. Putnam is correct to suggest that the worry that I am just as conceptually deficient as the BIV is incoherent. That is because the BIV is conceptually deficient to the degree that it cannot formulate the worry that it is conceptually deficient. The man who has lived his entire life inside of a sandbox (and is therefore conceptually deficient), however, *can* formulate the worry. Putnam claims that if a worry about diminished conceptual scope is merely a “worry about the limits of human science...the hypothesis may be true, but does not imply any radical skepticism” (Hale, 286). It is hard to see how the difference between the BIV and the man in the sandbox is the difference between a worry being skeptical or not. If skepticism threatens what we epistemically value it need not show that we lack truth entirely. The skeptic may instead aim to show that the truth we *do* have is severely limited. This can be done by solely focusing on conceptual deficiency and the possibility of being conceptually deficient need not be unintelligible in order to be skeptical. It is likely foolish to doubt that finite beings such as ourselves are conceptually deficient in some respect. As I suggested earlier, there is clearly a spectrum of conceptual deficiency, and the skeptical worry involves locating oneself as close to a complete conceptual lacking as possible

---

<sup>24</sup> Remember that even though *both* men lack truths regarding the external world it is the way in which their respective lacking is perceived of which is of importance. The man who plans out his grand adventure has some sense of his own lacking.

without rendering the worry unreasonable due to its being the case hinging on its intelligibility. We might say that the BIV is *radically* conceptually deficient to the extent that it cannot even worry that it is conceptually deficient. Less conceptually deficient is the man who studies grains of sand. He at least has the concepts required to worry that he is severely conceptually deficient. The skeptical worry that we are severely conceptually deficient is not the worry that our knowledge is incomplete (that is obvious). It is the worry that the extent to which we are conceptually deficient may never manifest itself as a noticeable gap in any of our theories. The massive degree to which the man who counts grains of sand is conceptually deficient will never influence his capacity to carry on his investigation. His conceptual deficiency will never make itself known.

## Conclusion

In order for the skeptic's argument to be a good one she must *not* provide a skeptical argument which aims to maximize both forms of epistemic worry. She must either utilize a weakened counterpossibility, which reduces both the degree of conceptual deficiency and falsehood from that found in the radical counterpossibility, or she must let go of the worry regarding falsehood altogether and focus solely on conceptual deficiency, in which case there is a limit to how conceptually deficient we can be. Although Putnam's argument does not show that there are no reasonable skeptical worries regarding knowledge of the external world it does illustrate a significant limitation to them. A skeptical argument's being reasonable demands that the truth of the counterpossibility (or possibility of conceptual deficiency) it provides not hinge on its being intelligible. If the worry is that I am a BIV, or that I am a severely conceptually deficient being, it cannot be that such an entity lacks the semantic resources to form that very worry. If such is the case the worry is unreasonable.

Currently it seems to me that skepticism, of the sort I have dealt with here, should be seen as a tool for bringing to light contradictions internal to one's background presuppositions. The skeptical argument which utilizes a radical counterpossibility did, at one point, seem to me to be a good argument. This was before I had considered the preconditions for stating such a skeptical worry. Being troubled by the radical counterpossibility, I now know, demands accepting an untenable sort of reference. Worries which maximize falsehood or conceptual deficiency fail for this very reason. Being troubled by skepticism, then, is a method of generating questions which would have otherwise been missed. It is a way to justify questioning what one would otherwise

take as given. For this reason I see skepticism as an essential tool for the philosopher (including the stubbornly anti-skeptical ones).

## **Bibliography**

Baehr, J. (2012). Credit Theories and the Value of Knowledge. *The Philosophical Quarterly*, 62, 1-22.

Brueckner, A. (1986). Brains in a Vat. *The Journal of Philosophy*, 83 (3), 148-167.

Brueckner, A. (1992). If I Am a Brain in a Vat, Then I Am Not a Brain in a Vat. *Mind*, 101 (401), 123-28.

Brueckner, A. (1992). Semantic Answers to Skepticism. *Philosophical Quarterly*, 73, 200-219.

Forbes, G. (1995). Realism and Skepticism: Brains in a Vat Revisited. *The Journal of Philosophy*, 92 (4), 205-222.

Gallois, A. (1992). Brains in Vats, and Arguments for Scepticism. *Mind*, 101 (402) 273-86.

Grimm, S. (2008). Epistemic Goals and Epistemic Values. *Philosophy and Phenomenological Research*, 77 (3), 725-744.

Grimm, S. (2012). The Value of Understanding. *Philosophy Compass*, 7 (2), 103-117.

Hale, B. & Clark, P. (1994). Comments and Replies. *Reading Putnam*. Basil Blackwell Ltd. Pp. 284-288.

Johnsen, B. (2003). Of Brains In Vats, Whatever Brains In Vats May Be.” *Philosophical Studies*, 112, 225-249.

Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.

- Kvanvig, J. (2008). Pointless Truth. *Midwest Studies in Philosophy*, 32, 199-212.
- Ludwig, K. (1992). Brains in a Vat: Subjectivity, and the Causal Theory of Reference. *Journal of Philosophical Research*, 17, 314-45.
- Piller, C. (2009). Desiring the Truth and Nothing But the Truth. *Nous*, 43 (2), 193-213.
- Putnam, H. (1975). The Meaning of 'Meaning'. *Minnesota Studies in the Philosophy of Science*, 7, 131-193.
- Putnam, H. (1981). *Reason, Truth, and History*. Cambridge University Press, 1-21.
- Putnam, H. (2006). Philosophy as the Education of Grownups. *Reading Cavell*.  
Routledge, London and New York, pp. 119-130.
- Riggs, W. (2003). Balancing Our Epistemic Goals. *Nous*, 37 (2), 342-352.
- Sosa, E. (2000). For the Love of Truth? In Linda Zagzebski (ed.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Oxford University Press, pp. 49-62.
- Sosa, E. (2003). The Place of Truth in Epistemology. In DePaul (ed.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Clarendon Press, pp. 155-179.
- Warfield, T. (1998). A Priori Knowledge of the World: Knowing the World by Knowing Our Minds. *Philosophical Studies*, 92, 127-147.
- Weiss, B. (2000). Generalizing Brains in Vats. *Analysis*, 60, 112-23.
- Whiting, D. (2012). Does Belief Aim (Only) at Truth? *Pacific Philosophical Quarterly*, 93 (2), 279-300.
- Whiting, D. (2013). The Good and the True (or the Bad and the False). *Philosophy*, 88, 219-242.

Wright, C. (1992). On Putnam's Proof That We Are Not Brains-in-a-Vat. *Proceedings of the Aristotelian Society*, New Series, 92, 67-94.