

Semantic Pragmatism and *A Priori* Knowledge
(or “Yes we could all be brains in a vat”)*

1. Introduction

Hilary Putnam has famously argued that we can know that we are not brains in a vat because the hypothesis that we are is self-refuting.¹ While Putnam’s argument has generated interest primarily as a novel response to skepticism, his original use of the brain in a vat scenario was meant to illustrate a point about the “mind/world relationship.”² In particular, he intended it to be part of an argument against the coherence of metaphysical realism, and thus to be part of a defense of his conception of truth as idealized rational acceptability. Putnam’s argument has drawn a good deal of criticism already, but it will be argued here that these criticisms fail to capture the central problem with Putnam’s argument. Putnam’s conclusions about the self refuting character of the brain in a vat hypothesis, rather than simply being a consequence of his semantic externalism, will be shown to be actually out of line with central and plausible aspects of his own account of the relationship between our minds and the world. Reflections on intentionality and semantics ultimately give us no compelling reason to suppose that the beliefs expressed by claims like “I am a brain in a vat” could not be true,³ but (*pace* Putnam) this supports neither skepticism nor metaphysical realism.

2. Putnam’s Argument

Putnam’s attempt to show that we could not be brains in a vat begins with his asking us to imagine the following scenario:

[A] human being (you can imagine this to be yourself) has been subjected to an operation by an evil scientist. The person’s brain (your brain) has been removed from the body and placed in a vat of nutrients which keeps the

* I’d like to thank Robert Brandom, Joe Camp, Jonathan Cohen, Brian Garrett, Mark McCullagh, John McDowell, Ram Neta, Deborah Smith, and audiences at the Mid-South Philosophy conference and York University for comments on earlier versions of this paper.

¹ Putnam 1981, ch. 1.

² Putnam 1981, p. 6

³ To simplify the presentation, the paper will typically focus on the types of claims that we can make, but the point should be understood as extending to the beliefs expressed by those claims. The relevant issues deal with the limitations on what we can think about as much as they do with what we can coherently talk about.

brain alive. The nerve endings have been connected to a super-scientific computer which causes the person whose brain it is to have the illusion that everything is perfectly normal. . . .⁴

While this scenario “violates no physical law,” and is “perfectly consistent with everything we have experienced,” Putnam still insists that it “cannot possibly be true, because it is, in a certain way, self-refuting.”⁵ Putnam takes the hypothesis to be self-refuting because it purports to state a possibility that (according to his understanding of semantic externalism) should be unstateable. What our words refer to is determined by what their usage is causally connected to, and a brain in a vat’s usage of “vat” would not have the sorts of causal connections to vats needed for it to designate them. As Putnam puts it, “‘Vat’ refers to vats in the image in vat-English, or something related (electronic impulses or program features), but certainly not to real vats, since the use of ‘vat’ in vat-English has no causal connection to real vats.”⁶ If we were brains in a vat, then our word “vat” wouldn’t refer to *vats*. Consequently, the mere fact that we can raise the possibility that we are brains in a vat shows that we are not. In other words, “If we can consider whether it is true or false, then it is not true . . . Hence it is not true”⁷

Putnam’s assumption that the use of “vat” in vat-English has no causal connection to real vats is, of course, essential to his argument, and this leads him to seriously modify his original scenario. Someone “subjected to an operation by an evil scientist” could have had plenty of causal contact with vats before being envatted, and even someone who was always a brain in a vat could have such causal connections indirectly (through artificial sense receptors or changes in the virtual environment based upon what takes place outside of it). It is, then, not surprising that Putnam adds a number of further embellishments to his story. In particular, one’s brain is supposed *always* to have been envatted, and the vat and automated machinery are no longer designed by an

⁴ Putnam 1981, pp. 5-6.

⁵ Putnam 1981, p. 7.

⁶ Putnam 1981, p. 14. See also, “Although the people in that possible world can think and ‘say’ any words we can think and say, they cannot (I claim) *refer* to what we can refer to. In particular, they cannot think or say that they are brains in a vat (*even by thinking ‘we are brains in a vat’*).” (Putnam 1981, p.8.)

⁷ Putnam 1981, p. 8. See also: “It follows that if . . . we are really 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 So, if we are brains in a vat, 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.” (Putnam 1981, p. 15.)

intelligent scientist, but rather are “supposed to have come into existence by some kind of cosmic coincidence” so that they “have no intelligent creators or designers.”⁸ It is only the hypothesis so embellished that is supposed to be self-refuting. Consequently, in spite of Putnam’s tendency to say things such as “I am claiming that there is an argument we can give that shows that we are not brains in a vat,”⁹ Putnam presents no such argument unless “brains in a vat” is understood as shorthand for the modified scenario.¹⁰ Nevertheless, the secondary literature has followed Putnam’s lead in using the expression “brain in a vat” to refer to the modified scenario, and this paper will, henceforth, do the same.

3. Traditional Objections and the Shared Assumption

Even those who typically see little point in worrying about whether or not we might all be brains in a vat have been surprised by Putnam’s claim to be able to *prove* (indeed, prove *a priori*) that we couldn’t be. Consequently, many have attempted to reconstruct his arguments more formally in order to identify just what assumptions and inferences are being made. There have been many such reconstructions, but Putnam’s argument can, for present purposes, be reconstructed as follows:¹¹

- (i) My language disquotes.
- (ii) In vat-English, “brain in a vat” does not refer to brains in vats.
- (iii) In my language “brains in a vat” is a meaningful expression.
- (iv) In my language, “brains in a vat” refers to brains in a vat. [From (i) and (iii).]
- (v) My language is not vat-English. [From (ii) and (iv).]
- (vi) If I am a brain in a vat, my language, if any, is vat-English. [Definition of vat-English.]
- (vii) I am not a brain in a vat. [From (v) and (vi).]

⁸ Putnam 1981, p.12.

⁹ Putnam 1981, p. 8.

¹⁰ For a discussion of this, see Brueckner 1986, p. 152.

¹¹ I am here following Wright 1994 (p. 224), since Putnam himself seems to endorse this reconstruction (Putnam 1994, p. 284). In any case, the objections considered below should be locatable in any of the many acceptable formulations of Putnam’s argument available.

While this argument seems valid, questions about the types of semantic self-knowledge compatible with the externalist semantic framework Putnam presupposes have led some to challenge our *a priori* entitlement to a number of its steps.

Possibly the earliest and most influential line of objection to Putnam's argument questions our *a priori* entitlement to (iv).¹² We do not have introspective access to those 'external' factors that, according to Putnam's externalistic semantic framework, determine what our expression "brain in a vat" refers to. Consequently, even if Putnam's argument lets one know that one's sentence "I am not a brain in a vat" must be true, it doesn't let one know that one is not a brain in a vat. One may know the sentence's truth-value, but one still lacks *a priori* access to its *content*. Disquotation alone is not enough to insure that one knows what one's sentences mean, since mastery of the disquotation schema does not require understanding all of the terms found within it. To *really* know that one was not a brain in a vat, one would have to know whether one was speaking English or vat-English, and one could only know that if one already knew whether or not one was a brain in a vat.¹³

A second, and more radical, line of attack can be directed at step (iii). Critics can argue that our lack of semantic self-knowledge extends to the point that one cannot even tell by introspection whether one's words and 'thoughts' are contentful at all. For all one knows, one's words and thoughts may have none of the causal connections to the external world needed to make them meaningful.¹⁴ How, the critic of our entitlement to (iii) might ask, does one know that one is not just 'thinking' with words that are utterly lacking in content? Internalists could at least be sure that they were thinking, but within Putnam's semantic framework, one's claim to know *a priori* that the

¹² For the best known exposition of this line, see Brueckner 1986.

¹³ "I can conclude ... that I am a normal human being rather than a BIV... only if I can assume that I mean by 'I may be a BIV' what normal human beings mean by it. But I am entitled to that assumption only if I am entitled to assume that I am a normal human being speaking English rather than a BIV speaking vat-English. This must be shown by an anti-skeptical argument, not assumed in advance." (Brueckner 1986, p. 103.)

¹⁴ The connection between externalism and this possibility is made very vivid in the discussion of the 'Swampman' in Davidson 1987. Putnam discusses a related possibility in Putnam 1981, p.17. The possibility of externalist anti-skeptical arguments backfiring in this way is discussed by Brueckner 1986, p.159, Falvey and Owens 1994, p. 126, Klein 1986, p.385, and Stich 1994 p. 356.

expressions running through one's consciousness are meaningful, and thus one's entitlement to step (iii), seems undermined.

The force of such attacks on Putnam's argument is a matter of some controversy.¹⁵ However, even if such criticisms are sound, they still allow Putnam's argument to establish a number of surprising and non-trivial conclusions. The first objection allows that any attempt to *formulate* the skeptical hypothesis will be false. The second allows that a denial of the skeptical hypothesis is presupposed by our assumption that we are thinking at all. Both objections thus leave in place the conclusion that there is something fundamentally problematic with attempts to *claim* that one might be a brain in a vat. Both concede that if one is entitled to the claim that one is thinking (and that one knows *what* one is thinking), then one is entitled to the claim that one is not a brain in a vat. Such allowances to Putnam concede too much, and the problems with Putnam's argument are more fundamental than the two previous objections suggest.

In particular, the most serious problem with Putnam's argument is with step (ii), namely the assumption that:

(ii) In vat-English, "brain in a vat" does not refer to brains in vats.

Putnam, his supporters, and his critics typically agree that one can know *a priori* that, if one were a brain in a vat, then one's word "vat" would not refer to vats. Consequently, they all assume that the claim "I am a brain in a vat" couldn't *possibly* be true. Their disagreements are over what this purported 'semantic' fact is supposed to show. Putnam and his sympathizers take it to show that we can know that we are not brains in a vat, while his critics take it to show only that we can know that a certain type of utterance, if meaningful, must be false.

It is this shared assumption, that a brain in a vat could not refer to brains in vats, and thus could not truly think "I am a brain in a vat," that should be questioned. Indeed, the problems with (ii) are considerably more serious than those with (iii) and (iv). I know of no one who seriously questions the *truth* of (iii) or (iv). All that is questioned is our *a priori* entitlement to them. I can doubt (iii) and (iv) only in some very limited 'philosophic' sense. I recognize that I may not be

¹⁵ See, for instance, Brueckner 1986, Nagel 1986, Falvey and Owens 1994, Wright 1994, Davies 1995, Forbes 1995, Ebbs 1997, Noonan 1998, Sawyer 1999, Ludlow and Martin 1998, and the numerous papers cited therein.

entitled to them while in a philosophical argument with the skeptic, but I have no *real* doubt about their truth. No one really doubts that the word “vat” is a meaningful expression, or that it refers to vats. On the other hand, many people’s naïve intuitions about (ii) seem to be that it is false. This is why the Evil Demon and brain in a vat hypotheses have seemed coherent, if implausible, to so many. Such naïve intuitions can, of course, turn out to be incorrect. Nevertheless, it will be argued here that not only do we have little reason to think that (ii) can be known *a priori*, but we also have good reason to doubt that it is true at all. The following criticisms of (ii) (unlike those of (iii) and (iv)) thus question not only the *a priori* availability of Putnam’s argument, but also its *soundness*.

The hypothesis that I am a brain in a vat *seems* like an intelligible (if unlikely) one, and the *prima facie* intelligibility of the hypothesis partially explains the intuitive discomfort that many people have with Putnam’s purported proof of its incoherence.¹⁶ Indeed, given the rather unpromising history of attempts to rule out such ‘skeptical’ scenarios on semantic grounds,¹⁷ one might think that any account of meaning that entailed that the brain in a vat hypothesis was unintelligible would, thereby, cast serious doubts upon its own acceptability.¹⁸ Questioning (iii) and (iv) does not, however, get at what is intuitively suspect about Putnam’s argument, since there is nothing unintuitive about Putnam’s assumptions that our words are meaningful and that *our* word “vat” refers to vats. Rather, what is unintuitive is the claim that a brain in a vat would be saying something false were it to say “I am a brain in a vat.” The objections that focus on (iii) and (iv) typically endorse (or at least ignore) this claim and only question what sort of knowledge can be derived from it. By contrast, an attack on (ii) gets to the heart of the matter by defending the conceivability of the brain in a vat’s ability to make a true claim about its condition.

¹⁶ After all, most people (including Putnam himself – Putnam 1981, p.7) typically *do* feel that there must be something ultimately wrong with Putnam’s argument when it is first presented to them. (Or at least that is my experience with students when they are presented with the argument, and with most of my colleagues with whom I have discussed it.) Their intuition is that there must be *something* wrong with the argument even if they cannot pin down what that something might be.

¹⁷ For a discussion of some of these, see Stroud 1984.

¹⁸ See Falvey and Owens 1994, and McGinn 1986 for claims of this sort. However, it will ultimately be argued that a commitment to ‘semantic externalism’ is perfectly compatible with the hypothesis’s intelligibility, and thus that no *reductio* of semantic externalism is in the offing.

4. Switching and De-vatting

Before evaluating the plausibility of premise (ii), consider the following two cases.

I. A speaker discovers that seven days ago, while sleeping he was transported (it doesn't matter how) to Earth from his own planet (hereafter "Earth2"). Earth2 looks just like Earth though every substance on it has a different atomic structure than does its Earth-counterpart. On looking back at what he said over the past week, he is inclined to say that assertions like "I'm in Toronto", "That's Hilary Putnam", "here is a rabbit", and "This is water" were mistaken. He considers his terms "Toronto", "Hilary Putnam", "rabbit" and "water" to not (yet) refer to the people, places, animals and substances that go by those names here on Earth. Rather, he takes them to refer to their counterparts on Earth2. Still, he thinks that, say, the things he called "phones," "cars" and "spoons" here on Earth were, in fact, phones, cars and spoons. Indeed, he thinks that he was correct to call the vats on Earth "vats," even if he could not truly apply any of the terms for what Earth2-vats are made of ("copper," "steel," "iron," etc.) to the vats on Earth.

II. A speaker discovers that seven days ago (it doesn't matter how) his sleeping brain was scooped out of the vat it had always floated in and placed in a human body. He discovers that while his new environment seems exactly like his old one, his experiences of his old environment were dependent upon the states of a computer in this new environment. Indeed, the whole set-up responsible for his previous experience seems to have come together through some sort of 'cosmic coincidence'. Looking back at what he said over the past week, he is inclined to say that assertions like "I'm in Toronto", "That's Hilary Putnam" "here is a rabbit", and "This is water" were mistaken. He takes his terms "Toronto," "Hilary Putnam", "rabbit" and "water" not to refer to the people, places, animals and substances in this environment. Rather, he takes them to refer to their counterparts in his previous computer-generated environment. Nevertheless, he still thinks that

could correctly identify the phones, cars, spoons and vats in this new environment as “phones,” “cars,” “spoons” and “vats.”¹⁹

These two cases may represent how speakers would describe themselves and their usage upon discovering that they had been recently ‘switched’ or ‘de-vatted.’ The question remains, however, whether we should endorse such descriptions.

5. Externalism and Non-Natural Kinds

Of course the intuition that someone could correctly identify vats as “vats” upon being ‘de-vatted’ is precisely what Putnam claims semantic externalism gives us compelling reasons to reject. However, the intuition can be understood as compatible with semantic externalism if we understand “vat” to pick out some type of kind that has instances in both ‘real’ and ‘virtual’ contexts. For instance, while a natural kind term like “water” may be inapplicable in contexts where the functionally/experientially equivalent substances lack water’s molecular structure, terms like “spoon” or “vat” can be applied ‘across contexts’ provided that the differently constituted ‘spoons’ or ‘vats’ play the same ‘role’ in the alternate environments.

The sampling of objects one’s usage is causally connected to constrains what one’s terms can refer to, but it does not, in itself, determine what sortals they should be interpreted as falling under.²⁰ A term in a language can denote objects that its users have not had causal contact with if it picks out a category/kind that encompasses both those unexperienced objects and whatever instances of the kind that the speakers do experience. After all, while our use of “vat” has no causal connection to the vats on Earth2, that hardly means that they cannot fall within the term’s extension. If our term “vat” were interpreted as picking out, say, some sort of functional or interactional kind, then “vat” would pick out both the vats on Earth that we have experienced and the vats on Earth2 that we have not. In much the same way, the vat-English term “vat” might pick

¹⁹ I should note that this case departs from Putnam’s example slightly since it allows that there are other conscious creatures outside of the vat. Nothing, however, should turn on this, since the creatures outside of the vat are taken to have nothing to do with the vat in which the brain sits and the virtual world generated by the computer.

²⁰ A point that should be familiar from Putnam’s own discussions of the feasibility of *purely* causal accounts of reference. (See, for instance, Putnam 1981, p.53.)

out a kind that includes both the vats-in-the-image (or ‘virtual vats’)²¹ that the speakers of vat-English have experienced, and the ‘real vats’ that they have not.²²

After all, consider the following three sets of objects: the set of all physical vats (hereafter P-Vats), the set of all virtual vats (hereafter V-Vats), and the combined set of all P- and V-Vats (hereafter C-Vats).²³ While the brain in a vat’s term “vat” couldn’t be interpreted as picking out just the set of P-Vats, it is far less clear that it must be interpreted as picking out the set of V-Vats rather than C-Vats. Since all V-Vats are C-Vats, the brain in a vat’s usage has had just as much causal contact with C-Vats as it has with V-Vats (while it has had none with P-Vats).²⁴ In light of this, we should keep in mind that Putnam’s premise (ii) is:

(ii) In vat-English, “brain in a vat” does not refer to brains in vats.

Which is incompatible with the claim that their term picks out C-Vats. It is *not* the more plausible premise

(ii)* In vat-English, “brain in a vat” does not refer *exclusively* to P-Vats.

Which is compatible with the expression picking out C-Vats, but not with its picking out just P-vats.²⁵ Unfortunately for Putnam, while (ii)* is more defensible than (ii), his argument is invalid if (ii)* is substituted for (ii).²⁶

²¹ Putnam is less than clear about what these “virtual vats” should themselves be understood to be. Vats in the image, electronic impulses, and program features have all been suggested by Putnam (1981, p.14) and others. I will try not to take a stand on this issue, and thus treat ‘virtual vats’ to pick out whatever is *causally responsible* for the brain’s ‘vat utterances’. Consequently, the ‘virtual vats should not be understood as in any way fictional in the way that we think of unicorns as fictional, since, whatever they are, they have causes and effects.

²² One might question both this use of ‘real vat’ and thus Putnam’s claim that “the use of ‘vat’ in vat-English has no causal connection to real vats,” (Putnam 1981, p. 14). Both may seem to beg the question at hand by assuming that the vats in the image could not be ‘real.’ On the other hand, one might try to preserve Putnam’s claim by arguing that ‘real’ could be used as a comparative term picking out a *type of* vat, and something that was not a ‘real’ vat could still ‘truly’ be a vat. (On the various uses of ‘real’ see Austin 1962, ch.7.)

²³ For ease of exposition, assume that we are not brains in vats and that ‘physical’ refers to *this* environment, and “virtual” is virtual *relative to* this environment.

²⁴ Of course, one might try to argue that the categories of P- and V- Vats are somehow more ‘natural’ than that of C-vats, and that the initial samples only ‘project’ to such ‘natural’ properties (see, for instance, Lewis 1983, 1984). However, such a line could hardly be appealed to by Putnam, since such an interest-independent ‘ranking’ of properties is one of the characteristics of Metaphysical Realism he is most anxious to reject.

²⁵ Note that Putnam assumes (in Putnam 1978, p. 127) that the metaphysical realist would describe the brain in the vat as referring to P-Vats rather than C-Vats by “vat”.

²⁶ The issue of how to understand premise (ii) is actually more complex than this. One might argue that all that Putnam’s argument requires is that the phrase “brains in a vat” have different extensions in English and Vat-English. (Indeed, Wright suggests something like this (Wright 1994, pp. 221-3).) Consequently, as long as the reference of “vat” in English was the set of P-Vats, then the argument would go through whether the Vat English expression referred to either C or V-Vats. However, if one takes this line (and I would argue that

In light of this, consider Putnam’s analysis of the extension of “water”:

We can understand the relation *same_L* (same liquid as) as a cross-world relation by understanding it so that a liquid in W_1 [World 1] which has the same important physical properties (in W_1) that a liquid in W_2 possesses (in W_2) bears the *same_L* to the latter liquid.... an entity x , in an arbitrary possible world, is *water* if and only if it bears the relation *same_L* (construed as a cross-world relation) to the stuff *we* call ‘water’ in the *actual* world.²⁷

While Putnam may be right to claim that the ‘*same_L*’ relation has to do with physical/micro-structural properties in the case of “water,” the same-kind relation for “vat” is not best understood this way.²⁸ Putnam claims that the ‘hidden structures’ determine the reference of natural kind terms not because only such hidden structures could serve in the same-kind relation, but rather because “normally the ‘important’ properties of a liquid or a solid, etc., are the one’s that are structurally important.”²⁹ However, while Putnam’s claim may be true for terms like “water” and “gold”, importance is, as Putnam himself stresses, “an interest relative notion,”³⁰ and for vats it is how we are able to interact with them rather than micro-structural properties that are important. Consequently, we might give the following account of the extension of the brain in a vat’s term “vat”:³¹

We can understand the relation *same_R* (same role as) as a cross-environmental relation by understanding it so that an object in E_1 [Environment 1] which has the same important interactional properties (in E_1) that an object in E_2 possesses (in E_2) bears the *same_R* to the latter object.... an entity x , in an arbitrary possible environment, is *a vat* if and only if it bears the relation *same_R* (construed as a cross-environmental relation) to the things *we actually* call ‘vats.’

How one is able to interact with an object depends upon one’s body as well as the object itself. Virtual vats have the same interactional properties as non-virtual vats because the subjects in the

establishing that the English expression picks out just P-vats is a non-trivial task) the problem reemerges in terms of the question of determining what the referent of “vat-English” is supposed to be. If “vat-English” is simply the language spoken by any brains in vats that I would encounter in *my* environment, then the argument is not an interesting one, since there was never a worry about whether I was a brain in one of the vats in *my current* environment. On the other hand, if “Vat English” is just a general term for English-like languages spoken by anything my expression “brain in vat” can truly apply to, it is less clear that it picks out a single language which can be identified with the version of English that would be spoken by any brains in vats found in *this* environment.

²⁷ Putnam 1975, p. 232

²⁸ After all, ‘vat’ would seem to mean the same thing in English and its Earth2 counterpart, even if the term was applied to an entirely different set of objects on Earth2.

²⁹ Putnam 1975a, p. 239. This emphasis on our interests separates him, to his credit, from Devitt 1980 and Lewis 1983, 1984.

³⁰ Putnam 1975a, p. 239.

³¹ Cross-environmental relations replacing cross-world ones here, since the same object could have different interactional properties in different environments within the same world.

vat-world have virtual bodies that interact with them in just the way that non-virtual bodies interact with non-virtual vats. Consequently, if the brain in a vat's term "vat" picks out this sort of cross-environmental interactional kind, then it would pick out both the 'virtual' vats in its own environment and the 'real' vats in ours, because when the subjects change environments, they change bodies as well. A brain in a vat would thus be able to truly claim "I am a brain in a vat."³²

Of course, one might suggest that a term like "vat" need not be analyzed in quite this way. In particular, one could argue that it picks out a functional kind that makes reference not only to the function played, but also to its being able to play it in a *particular* environment.³³ For instance, consider the following account of the extension of "vat":

We can understand the relation *same_F* (same function as) as an environment-specific relation by understanding it so that an object in E_1 [Environment 1] bears the *same_F* relation to an object in E_2 if it would have the same important functional properties were it in E_2 that the latter object possesses (in E_2)... an entity x , in an arbitrary possible environment, is a *vat* if and only if it bears the relation *same_F* (construed as an environment-specific relation) to the things *we actually* call 'vats.'

If "vat" picks out this sort of environmentally specific functional kind rather than a cross-environmental one, then "vat" would have completely different extensions in English and vat-English. Vats would have the function of holding, for instance, *water*, while virtual vats would lack the ability to hold any such non-virtual liquids. The functional roles played would thus be very different. However, it seems unclear why we should believe that by "vat" the brains in the vat must intend to pick out this more restrictive sort of functional kind rather than the more expansive sort of interactional kind suggested above.³⁴ Some functional kinds are 'objective' in

³² There are, of course, some 'skeptical' hypotheses that *might* still be self-refuting. For instance, the claim "I have always been a brain in a vat on the dark side of the moon" may be self-refuting. Even if a brain in the vat could talk about our moon as "*a* moon," when it uses the term "the moon" it refers to something in the image, not in our world. Much the same could be said of the hypothesis "I have always been a brain in a vat sitting in *Hilary Putnam's* basement." Such hypotheses make reference to certain *particulars* in our environment and thus require *names* for their formulation. However, the interest of *these* skeptical hypotheses is obscure to me. Furthermore, there may very well be future experience (discovering massive and constant switching, etc.) that would lead us to conclude that our 'proper names' actually were multi-realizable kind terms.

³³ Some seem to think that the term "functional kind" should only be used this way, hence my preference for "interactional" kind.

³⁴ Such an understanding of functional kinds would, after all, seem to miss out on how we understand even such basic functional kinds such as "heart." While human hearts and mouse hearts play similar roles in their respective environments, they could not, to put it mildly, play their roles adequately if their environments were switched. Such an interpretation would also require that the meaning of "phone" has changed over the last 20 years, since many phones we now use (cell phones in particular) would not be able to function in a remote or past environment where there were no satellites to support them. A less environmentally restrictive account of functional kinds, on the other hand, could easily explain why we are entitled to consider cell phones to be a

that the relevant functions can be specified without making any reference to our activities.³⁵ For other kinds, the relevant functional role makes essential reference not just to other objects in the world but to how they interact with the subjects (and possibly their social practices) as well. The resulting kinds may seem very ‘unnatural’ since, ‘objectively’, real and virtual vats (large metal containers and states of a computer) seem to have nothing in common. However, while the interactional kind might seem like a ‘funny’ disjunction of properties if the experience of the subject is left out, that makes them no less legitimate. After all, it has been argued that much the same is true of our color terms. Objectively, the things that are, say, blue (my shirt, the sea, the sky) may seem to have little in common, and it is only how they interact with us and our optical apparatus that grounds their falling under a single kind.³⁶

At this point, one should note that, even if the brain in a vat’s term “vat” could refer to *our* vats, it need not follow that *our* term “vat” must also refer to the virtual vats.³⁷ After all, our inclination to understand our term as a cross-environmental one upon discovering a ‘virtual world’ may be considerably less than the inclination of the recently de-vatted speakers to see their own terms this way.³⁸ We may ultimately decide that the terms in vat-English typically have different

type of phone. These considerations are hardly conclusive, but as will soon become clear, for the purposes of the current argument, the cross-environmental interpretation of “vat” need only be established as *possibly* correct to cause problems for Putnam’s argument. Furthermore, there is nothing about the second definition’s reference to an actual environment that makes it more in keeping with Putnam’s ‘indexical’ account of meaning than the first. Each case allows that the objects experienced in the initial environment help determine the reference, the disagreement is just over which type of sortal these initial samples should be understood in terms of.

³⁵ The restrictive nature of the alternate definition can also be seen in the fact that, if the molecular differences between Earth and Earth2 prevented Earth2 vats from holding Earth liquids, then we would have to say that, appearances to the contrary, the ‘vats’ on Earth2 weren’t really vats. This is a bullet that could be bitten, but it isn’t one we should be eager to bite.

³⁶ For a discussion of such cases, see Thompson 1995, Lakoff and Johnson 1999. Furthermore, there is now considerable evidence that classification is often not carried out in terms of categories defined in terms of shared sets of properties (Rosch 1975, Lakoff 1987). I do not have the space to pursue this point here, but if one accepts such ‘prototype driven’ accounts of concepts and categories, it seems quite likely that it would be even easier to defend the claim that a term like “vat” could be truly applied within the new environment.

³⁷ Though it is far from clear that it shouldn’t. After all, if we were to enter into the vat’s virtual world, we probably would use regular English words to describe the ‘virtual’ phones, cars and vats that we experienced. This raises the question of why we shouldn’t simply understand these words in terms of the experienced similarities that lead us to apply them cross-environmentally. (Once again, for ease of exposition, I’m assuming here that we are not brains in vats.)

³⁸ Indeed, we may be, when traveling between such ‘orders of reality’, generally more willing to ‘trade up’ than ‘trade down’.

extensions than the terms in English do, but that may only be because the terms in vat-English apply in both environments while the terms in English apply in just one.

One could say much the same for the other terms in the skeptical hypothesis such as “brain,” “in,” and “cause.” For instance, Putnam suggests that a brain in the vat would have had no experience of one thing actually causing another, and so could not mean what *we* (purportedly) do by “cause.”³⁹ This may be so, but upon being de-vatted, the former brain in a vat might come to realize that what *it* meant by “cause” was, after all, something (roughly) like lawlike correlation. While what the brain in a vat means by “cause” is not what *we* mean by the term, it could still truly claim that its experiences were being ‘caused’ by a computer, since its experiences would be correlated with the computer’s states in a lawlike fashion. Even if it didn’t have *our concepts*, it could still use *its* concepts to make claims that were true of *our world*. “Brain” is also plausibly a type of functional/interactional kind term, and brains in a vat might plausibly be able to refer to their actual brain. This is not only because brains are the source of their thinking, which is what they presume the ‘brains’ to be, but also because, just as one’s “body” could be understood in relation to its interactions as the locus of agency within an environment, one’s brain can be understood as whatever plays a certain role vis a vis that body. “In” is also easily understood in experiential terms, and would thus lend itself to a cross-environmental interpretation.

This claim that a brain in a vat could refer to brains and vats thus involves neither a retreat from the ‘externalism’ upon which Putnam bases his argument, nor the acceptance of any sort of ‘magical’ or ‘transcendental’ theory of the mind’s relation to the world.⁴⁰ ‘Interactional’ kind terms are still externalistically understood, and the kinds that they pick out are constrained by the *actual* role played by the items referred to in the primary experiential environment. What determines the term’s extension is thus not simply the speaker’s *conception* of the term. If “vats” in the virtual environment played a different role than our vats (say, they played the role of industrial colanders whose tiny holes would prevent them from effectively holding liquids), then

³⁹ Putnam 1994, p.287.

⁴⁰ For Putnam’s use of these terms for any account that would allow the brains in the vat to refer to vats with “vat”, see Putnam 1981, pp. 3, 5, 15, 16, and Putnam 1994, p.287.

someone whose experience was limited to those virtual ‘vats’ would not refer to vats with their term “vat.” Furthermore, this might hold true even if they were unaware of the tiny holes or straining function of the ‘vats’ in their own environment.⁴¹ It is the vats we *actually experience* that help determine the interactional kind picked out by “vat,” not just our *conception* of them.⁴²

6. Semantic Pragmatism and Self-Knowledge

Once again, even if the brain in the vat’s term “vat” picked out a cross-environmental interactional kind, it need not follow that *our* term “vat” does so. Nevertheless, there may be no way to tell *a priori* that it doesn’t, and if we can’t know *a priori* what kind of kind “vat” is, then we can’t know *a priori* that the claim “I am a brain in a vat” must be false. Indeed, given that some sort of cross-environmental interpretation of “vat” would seem extremely compelling if we were to suddenly find ourselves de-vatted, the claim that “vat” is not such a kind presupposes that we couldn’t experience a de-vatting.⁴³

Knowing what kind of kind a term picks out involves knowing what sorts of properties are *essential* to its application, and this is not something that can be conclusively determined *a priori*. For instance, I’m reasonably confident that “vat” picks out *some* sort of interactional kind.⁴⁴ However, I could, in principle, be mistaken about this. One could imagine it turning out that, upon more careful investigation, all of the ‘vats’ we had ever encountered were living beings that

⁴¹ If one of them were to say “we might all be brains in a vat,” another could correctly reply, “we couldn’t be, since all the nutrient fluid would flow out of the holes.”

⁴² This is why the analysis should not be viewed as ‘phenomenalistic’ (though a phenomenalistic analysis of the terms in a language might seem more plausible if one didn’t have a single experiential environment). The *cause* of the phenomena still determines the term’s referent. Some have argued that non-natural kinds should not be viewed as ‘indexical’ in this way (Schwartz 1978, 1980, Devitt & Sterelny 1987), and that the extensions of such ‘nominal kinds’ are “determined by an analytical specification of superficial features such as phenomenal properties, and/or form, function, or origin” (Schwartz 1980, p. 182). Such accounts, however, must assume that we (or at least some member of our community) can know *a priori* what the relevant functional or formal properties are, and there is no reason to think that this must be (even if it often is) the case. I won’t defend this last claim at length here, but the point is developed Kornblith 1980, and Jackman 1996.

⁴³ Which is more than just assuming that we *won’t* experience a de-vatting. The relevant dispositions are present as long as a de-vatting is physically possible.

⁴⁴ Precisely what sort of kind it may be is less clear. It may, for instance, be an *artifact kind* term, whose deliberate construction to play a certain role in our lives is essential to its being a member of the kind. In such a case, there could be no ‘naturally occurring’ vats. Treating “vat” as a term for an artifact kind rather than for the more generic sort of interactional kind suggested above would still, however, allow the term to be applied cross-environmentally.

moved around when they thought themselves unobserved, and that produced offspring that looked nothing like ‘mature’ vats. (These ‘baby vats’ were hidden away in ‘vat factories’ which were really just secret installations for the vats to grow up in.)⁴⁵ If this turned out to be the case, we might conclude that what “vat” picked out was, after all, a *natural* rather than an interactional kind.⁴⁶ The ‘baby vats’ would still be vats even if they were small and couldn’t hold liquids, and any vat-like object that *we* went on to construct would just be a mass of metal that looked like a vat.

The kinds of kind that our terms ultimately pick out will depend to a large extent on what kinds of kind ‘work’ best with our past, current and future experience. We may, for instance, have originally taken “air” to be a natural kind term of a sort that we still take “water” to be, but such an understanding of the term proved unworkable. Experience has a way of ‘boiling over’ our current understanding of our environment, and an understanding of our terms that seems adequate at a time may have to be radically changed as our experience unfolds.⁴⁷ Treating “vat” as a kind for objects that could play a given role within our *currently experienced* environment might seem plausible to many, but such an understanding might collapse quickly if we suddenly found ourselves ‘de-vatted.’ Our current experience simply may not settle just what kind of kind “vat” is.

The fact that we have only dealt with our current experiential environment can lead us to assume that its having a ‘physical’ make-up is essential to being a vat. Nevertheless, having a ‘physical’ instantiation may be no more essential to being a vat than being white is essential to being a swan. While we typically use “vat” to refer to vats in *this* experiential environment, future experience (involving either descending into virtual worlds, or emerging into a ‘realer’ world)

⁴⁵ This example is, of course, an adaptation of Putnam’s own discussion of “pencil.” (Putnam 1975a, pp.242-3.)

⁴⁶ Or, one could take the more extreme position that such a case would amount to our discovering that there were no, and never had been any, vats (for something like this view, see Katz 1975). However, Putnam himself clearly seems unsympathetic with this approach to such cases (see Putnam 1975a). More plausibly, one might think that the term should be understood as indeterminate between the natural and the interactional kind in the way that a term like “dog” might be understood as indeterminate between an ‘evolutionary’ and a ‘genetic’ understanding. The latter of these would allow a ‘synthetic dog’ which had no genealogical connection to our dogs, but an identical physical and genetic make-up, to be a dog, while the former would not (see Putnam 1994, pp. 76-7).

⁴⁷ The echo from *Pragmatism* (James 1907, p. 106) is deliberate, and it is found in Putnam’s work as well. (See Putnam 1995, p.8.)

might make our current assumptions about vats' 'physical' instantiation seem inessential to the term's meaning.⁴⁸ I know of no virtual worlds, and "All vats are physical objects" may be true. Nevertheless, it may be true partially because of the way the world turns out to be, not simply in virtue of the fact that we take certain properties to be essential to being a "vat." What kinds of kinds our terms ultimately pick out can be neither infallibly determined by introspection, nor conclusively settled by convention. Indeed, this is something that Putnam himself has stressed more than just about anyone else.⁴⁹ By requiring that we treat certain current beliefs or presuppositions as essential to a term's meaning (and thus as unrevisable in the face of future experience without changing the meanings of the terms involved), Putnam's own argument thus presupposes a type of 'semantic essentialism' at odds with his generally 'pragmatic' picture of meaning.

Putnam's argument thus requires that we have *a priori* knowledge of the kinds of kinds that our terms pick out, and there is little reason to think that we must have such knowledge. The potential failure of semantic self-knowledge involved here is more robust than that relating to criticisms of (iii) and (iv). While we typically are not mistaken as to whether we are thinking or not, we often *are* mistaken about what properties are *essential* to the application of our terms.⁵⁰ Consequently there is nothing unintuitive in suggesting that we occasionally lack this sort of self-knowledge.

7. Skepticism and Metaphysical Realism

Nevertheless, even if the brain in a vat hypothesis is a coherent one, the suggestion that we are radically mistaken about the world (in the sense of having mainly false beliefs) still seems hard to defend from within an externalist framework. The "transcendental" and "magical" conceptions of reference that Putnam criticizes would allow a brain in a vat to have a term "vat" which referred to

⁴⁸ One can see this in recent definitions of "life", where the properties essential to the kind are all of a functional/interactional sort that can be shared by various 'objects' found within the running of an appropriately programmed computer. (See, for instance, the discussion of "artificial life" in Turkle 1995.) Furthermore, one might argue that the experience of 'devoting' would lead us to reshape our conception of the 'physical' in a way that would allow 'natural kind' terms, and even the term "physical" itself, to apply cross-environmentally.

⁴⁹ See, for instance, his discussion of "cat," "energy" and "pencil" in Putnam 1962a, 1962b, and 1975.

⁵⁰ Indeed, our fallibility with respect so such questions has been evident in philosophic discourse from Socrates down to the present day.

vats but *didn't* refer to the 'vats' it experienced. By contrast, the position outlined here suggests that, while the reference of one's terms can extend *beyond* the sources of one's experience, it typically cannot be *divorced* from them. For instance, the brain in the vat's term "vat" would be instantiated in *both* the experienced and the unexperienced domain. A reinterpretation of current experience in the light of future experience typically will still leave the reanalyzed kinds applying to most of the currently experienced objects. The possible truth of the brain-in-a-vat hypothesis thus cannot be used to establish any sort of global skepticism. One can, for instance, still know that, say, one is looking at an apple, without knowing that one is not a brain in a vat because, even if one were a brain in a vat, one's claim to be looking at an 'apple' would still be true. The brain in a vat hypothesis is not a 'relevant alternative' that must be ruled out to be assured that one's claims are true.⁵¹ Consequently, the position defended here is still 'anti-skeptical' to the extent that it suggests that even if we were brains in a vat, most of our beliefs about the world we experience could still be true.

This would not, however, be enough to satisfy Putnam. The mere assurance that (even if we were brains in a vat) most of our beliefs would be true still leaves room for a considerable amount of epistemic disquiet. Brains in a vat, even if they typically have true beliefs, are fundamentally out of touch with reality's ultimate structure. In this sense, they still are 'radically mistaken' about the world. It is this worry that may ultimately be the target of Putnam's argument. After all, while Putnam admits that the possibility of our being brains in vats is typically used "to raise the classical problem of skepticism with respect to the external world in a modern way,"⁵² he claims that the possibility would be of interest only as a sort of "logical paradox" if it were not for the sharp way in which it brings out the difference between 'metaphysical' and 'internal' realism.⁵³ The brain in the vat is supposed to illustrate the metaphysical realist's worry that even our best theory could be radically out of touch with the world's fundamental structure, and Putnam's argument is meant to show how this worry characteristic of metaphysical realism is incoherent.

⁵¹ Once again, this is assuming that "brain in a vat" stands for the second scenario Putnam describes, which does not involve recent envating or 'switches' between the 'real' and 'virtual' environments.

⁵² Putnam 1981, p. 6.

⁵³ Putnam 1981, p. 49.

Many have found this attack on metaphysical realism unconvincing, and the following passage from Wright is a typical (if unusually clear) expression of the intuition that Putnam's argument shows, at best, that we may not even be able to state how bad our epistemic position is.

The difficulty is that Putnam's proof does not represent a general method for disproving *any* specific version of the relevant kind of possibility; at best, it represents a general method for disproving any specific version *which we can understand....* But the sort of dislocation whose possibility is arguably implicit in metaphysical realism does not involve that its victims can conceptualize their predicament; quite to the contrary -- their predicament consists in part precisely in the fact that they are debarred from arriving at the concepts necessary to capture the most fundamental features of their world and their place in it.... the real specter to be exorcised concerns the idea of a thought standing behind our thought that we are not brains in a vat, in just the way that our thought that they *are* mere brains in a vat would stand behind the thought -- could they indeed think anything -- of actual brains in a vat that "We are not brains in a vat." The specter is that of a thought whose truth would make a mockery of humankind and its place in nature, just as our true thought that they are merely brains in a vat makes a mockery of the "cognitive" activity of the envatted brains.⁵⁴

Putnam, in his reply to Wright, argues that this more abstract worry is, while perhaps less obviously so, as incoherent as the original assumption that we are brains in a vat.

It is, perhaps, the vagueness of terms like "fundamental categories," "real kinds," etc., that conceals from Wright the fact that he is tacitly assuming conceptual access to such general notions as "physical" and "causation." But I take it that what we mean by "fundamental categories" and "real kinds" is kinds and categories that play a fundamental role in the description of physical things and their causal relations; if not, then I will ask Wright to give me an account compatible with externalism of how a being whose position is analogous to that of a brain in a vat would refer to the property of being "fundamental."⁵⁵

Putnam's reply may indeed work as a response to Wright. Wright accepts Putnam's claim that a brain in a vat could not refer to vats, so he is not well placed to claim that it could refer to causation and fundamental categories. Nevertheless, Putnam's response to Wright ultimately runs into the same sort of trouble as his original argument. As with the case of "cause," even if the brains in the vat did not mean quite what *we* did by "fundamental" (which I doubt), it is quite plausible to think that their term could be interpreted so that it applied to the world outside the vat as well.

After all, the "fundamental categories" for the brains in the vat may be the categories by which their experiences were ultimately explained. If they were de-vatted, they would come to view certain categories relating to the computer as being 'more fundamental' than the categories they used before. Truths about the computer would, after all, explain why certain apparently 'fundamental laws' in their virtual environment were true. "Fundamental categories" apply to the

⁵⁴ Wright 1994, pp. 239-40. (See also Forbes 1995.) The same sort of intuition, though more explicitly tied to the traditional problem of skepticism, is expressed in Nagel 1986, p.73.

⁵⁵ Putnam 1994, pp.287-8.

total range of potential experience, and are not limited to the experiences available at a particular point. A brain in a vat can refer to categories in our world by its expression “most fundamental” because *future* experience could (though not necessarily ‘will’) reveal them to be so.⁵⁶ This is why the brains could describe the world outside the vat as ‘more real’ than the one they are currently experiencing.⁵⁷ Putnam wishes to understand truth and reference in terms of rational acceptability under “epistemically ideal conditions,”⁵⁸ and there is no reason why (given that a de-vatting is physically possible) potential out-of-vat experiences should be excluded from these ‘ideal conditions.’

Putnam’s intuitions about the incoherence of the brain in a vat scenario may be driven by a misplaced assimilation of it to the metaphysical realist’s worry about our never getting hold of things as they are in themselves. In formulating the metaphysical realism/internal realism issue in the way that he does, Putnam seems to be trying to use the brains in a vat to give a ‘naturalized’ version of a Kantian phenomena/noumena distinction. The vat’s virtual world would here take the place of the world of experience, and the world outside of the vat would play the role of the world as it is in itself. This understanding of the brain in a vat’s predicament is suggested by Putnam’s initial formulation of the problem in his “Realism and Reason.” In that paper, he asks how, if we were brains in a vat, would it come about “that *our* word ‘vat’ refers to *noumenal* vats and not to vats in the image?”⁵⁹ However, unlike Kant’s noumena, experience of the world outside the vat is not, *in principle*, inaccessible to the brains in the vat.⁶⁰ The world outside the

⁵⁶ Putnam seems to suggest that a future de-vatting could only be relevant by allowing the brain in the vat access to descriptions such as “the things I will refer to as ‘vats’ at such and such a future time” (Putnam 1981, p. 16). In much the same way, Putnam claims that the brain in a vat hypothesis would be a coherent one if it predicted a de-vatting some time in the future (Putnam 1981, p. 131). However, even the *potential* for de-vatting, whether it is actualized or not, is relevant to the interpretation of one’s terms. Even if the brain is not de-vatted, it is still *disposed* to respond to its de-vatting in a particular way, and these dispositions are relevant to the interpretation of its terms.

⁵⁷ The temptation to rule such statements about a world that is ‘more real’ as nonsensical presupposes that if an explication of what an expression means can’t be funded within *current* experience, then the expression is meaningless.

⁵⁸ Putnam, 1981, p. 55. Or at least the Putnam of *Reason, Truth and History* did. Putnam’s views on the topic of truth have changed since then.

⁵⁹ Putnam 1978, p. 127 (Italics, as elsewhere, are Putnam’s). Note that not only are the unexperienced vats outside the vat treated as ‘noumenal’, but it is also assumed that the hypothesis requires that vats in the image are *not* referred to by ‘vat.’

⁶⁰ Exactly how Kant’s distinction between appearances and things in themselves should be understood is a notoriously difficult topic. Nevertheless, it seems fairly certain that, whatever the proper understanding of

envatted brains is beyond the reach of their current experience, but there is nothing necessarily *inexperientiable* about it. The world outside the vat is not ‘the world as it is in itself.’ It is a world that the brains in a vat could (but unfortunately don’t) experience. If it were de-vatted (and given its new and expanded range of experience), the former brain in a vat would rightly deny that the theory of the world developed in the vat meets its ‘highest cognitive standards,’ since accounting for past, current, and future experience is manifestly such a cognitive standard. A brain in a vat might rationally inquire indefinitely without discovering its predicament, but ‘the best theory possible’ for the brain in a vat would not be one that left out the world external to the vat, since experience of that world is at least potentially available to it.⁶¹ Saying that we might be brains in vats is, after all, compatible with saying that we could, in principle, come to *recognize* that we had always been brains in vats. Even *if* the fundamental features of the world must be *experientiable*, they need not be *experienced*.⁶²

8. Conclusion

Putnam considered the brain in a vat hypothesis to be philosophically significant because he took the purported self-refuting character of this sort of skeptical worry to undermine the plausibility of metaphysical realism. However, Putnam’s argument against the possibility of our being brains in a vat relies upon treating the terms in the ‘skeptical’ hypothesis as if they picked out sortals that could not be applied cross-environmentally. There is, however, no compelling reason to think that such an assumption, even if true, could be established *a priori*. Consequently, there is no way to establish that a brain in a vat couldn’t truly think that it was brain in a vat. It may thus be, in some sense, possible that we are all brains in a vat. Nevertheless, the possibility that we are

Kant’s distinction between the noumena and phenomena is, it will be significantly different from the relation between experiences within the vat and experience of the world outside of it. For instance, the noumena are not distinct objects from the phenomena that could, in fact, eventually be objects of experience themselves.

⁶¹ Say, the ‘de-vatting’ also occurs through some ‘cosmic’ coincidence. If the first coincidence is ‘physically possible’ then it should seem as if the second would be as well. There should, then, be no problem tying such potential experiences to idealized, if not actual, inquiry.

⁶² No commitment need be taken here on here about the truth-value of this conditional’s antecedent.

brains in a vat, so understood, supports neither the skeptic's suggestion that most of our beliefs could be false, nor the metaphysical realist's worry that the best theory we could *possibly* come up with might still be radically out of touch with the world's fundamental structure. The coherence of the brain in a vat scenario simply does not have the philosophical implications that Putnam fears, and semantic externalism may thus lack many of the metaphysical and epistemic consequences that Putnam hoped for it. Since such purported metaphysical and epistemic consequences provided many with good reason to be wary of semantic externalism, such a result should ultimately make the view more plausible.

Works Cited

- Austin, J.L. 1962. *Sense and Sensibilia*, Oxford: Oxford University Press.
- Butler, K. 1997. "Externalism, Internalism, and Knowledge or Content", *Philosophy and Phenomenological Research* vol. LVII no. 4.
- Brueckner, A. 1986. "Brains in a Vat" *The Journal of Philosophy*,
- Brueckner, A. 1990. "Skepticism about Knowledge of Content," *Mind* 99: 447-51.
- Clark, P. and Hale, B. (eds.), 1994. *Reading Putnam*, Cambridge: Blackwell.
- Davidson, D. 1987. "Knowing One's Own Mind" in Ludlow and Martin 1998.
- Davies, D. 1995. "Putnam's Brain-Teaser", *Canadian Journal of Philosophy*, Vol 25, No. 2.
- Devitt, M. 1980. *Designation*, New York: Columbia University Press.
- Devitt, M. and Sterelny, K. 1987. *Language and Reality*, Oxford: Blackwell.
- Ebbs, G. 1997. *Rule Following and Realism*. Cambridge: Harvard University Press.
- Falvey and Owens, 1994. "Externalism, Self-Knowledge, and Skepticism" *The Philosophical Review*.
- Forbes, G. 1995. "Realism and Skepticism: Brains in a Vat Revisited" *The Journal of Philosophy*, vol. XCII, no.4.
- Jackman, H. 1996 *Semantic Norms and Temporal Externalism*, Ph.D. Thesis, University of Pittsburgh.
- James, W. 1907, 1975. *Pragmatism*. Cambridge: Harvard University Press.
- Katz, J. 1975. "Logic and Language: An Examination of Recent Criticisms of Intentionalism", in Gunderson, K. (ed.). *Minnesota Studies in the Philosophy of Science VII*, Minneapolis: University of Minnesota Press, 1975.
- Klein, P. 1986. "Radical Interpretation and Global Skepticism" in LePore (ed.) *Truth and Interpretation*. Oxford: Blackwell, 1986.
- Kornblith, H., 1980. "Referring to artifacts" *The Philosophical Review*, 89, 109-14.
- Lakoff, G. 1987. *Women, Fire and Dangerous Things*, Chicago: The University of Chicago Press.
- Lakoff, G and Johnson, M. 1999. *Philosophy in the Flesh*. New York: Basic Books.
- Lewis, D. 1983. "New Work for a Theory of Universals", *Australasian Journal of Philosophy*, v.61, n.4.
- Lewis, D. 1984. "Putnam's Paradox", *Australasian Journal of Philosophy*, v.62, n.3.

- Ludlow, P. 1995. "Externalism, Self-Knowledge, and the Prevalence of Slow Switching" in Ludlow and Martin 1998.
- Ludlow, P. and Martin, N. 1998, *Externalism and Self-Knowledge*, Stanford: CLSI Publications.
- McGinn, C. 1986, "Radical Interpretation and Epistemology" in LePore (ed.) *Truth and Interpretation*, Oxford: Blackwell, 1986.
- Nagel, T. 1986. *The View from Nowhere*, New York: Oxford University Press.
- Noonan, H. 1998. "Reflections on Putnam, Wright and brains in vats" *Analysis*, Vol 58, No. 1.
- Putnam, H. 1962a. "It Ain't Necessarily So", in Putnam 1975b.
- Putnam, H. 1962b. "The Analytic and the Synthetic", in Putnam 1975c.
- Putnam, H. 1975a. "The Meaning of 'Meaning'", in Putnam 1975c.
- Putnam, H. 1975b. *Mathematics, Matter, and Method*, New York: Cambridge University Press.
- Putnam, H. 1975c. *Mind Language and Reality*, New York: Cambridge University Press.
- Putnam, H. 1978. "Realism and Reason" in his *Meaning and the Moral Sciences*, Boston: Routledge & Kegan Paul, 1978.
- Putnam, H. 1981. *Reason Truth and History*, New York: Cambridge University Press.
- Putnam, H. 1994. "Comments and Replies", in Clark and Hale 1994.
- Putnam, H. 1994b. "Aristotle after Wittgenstein", in his *Words and Life*, Cambridge: Harvard University Press, 1994.
- Putnam, H. 1995. *Pragmatism, an open question*. Cambridge: Blackwell.
- Rosch, E. 1975. "Family Resemblances: studies in the internal structure of Categories", *Cognitive Psychology*, 7.
- Sawyer, Sarah, 1999. "My language disquotes", *Analysis*, v. 59. n.3.
- Schwartz, S. 1978. "Putnam on Artifacts" *The Philosophical Review*, 87, 566-74.
- Schwartz, S. 1980. "Natural kinds and nominal kinds", *Mind*, 89, 182-95.
- Stich, S. 1994. "Might Man be an Irrational Animal?" in Kornblith (ed.) *Naturalizing Epistemology*, Cambridge: MIT, 1994.
- Stroud, B. 1984. *The Significance of Philosophical Skepticism*, New York: Oxford University Press.
- Thompson, E. 1995. *Color Vision*, New York, Routledge.
- Turkle, S. 1995. *Life on the Screen*, New York: Simon & Shuster.
- Wright, C. 1994. "On Putnam's Proof that We Are Not Brains in a Vat", in Clark and Hale 1994.