



**WILEY-  
BLACKWELL**



---

Brains in a Vat: Different Perspectives

Author(s): Yuval Steinitz

Reviewed work(s):

Source: *The Philosophical Quarterly*, Vol. 44, No. 175 (Apr., 1994), pp. 213-222

Published by: [Wiley-Blackwell](#) for [The Philosophical Quarterly](#)

Stable URL: <http://www.jstor.org/stable/2219742>

Accessed: 19/08/2012 03:29

---

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Wiley-Blackwell and *The Philosophical Quarterly* are collaborating with JSTOR to digitize, preserve and extend access to *The Philosophical Quarterly*.

<http://www.jstor.org>

in giving an exposition of one of the horns of the dilemma. One horn of the dilemma is that if we take consciousness to be inseparable from the norm-governed notion of self-consciousness and from reflective and self-critical attitudes towards oneself, then it does seem hard to accept that Dennett's framework, in which conscious contents are a selection of the contents that go to and fro between subsystems, will capture all of the phenomena; but by the same token a frog's visual experience will not count as contentful and conscious, as McDowell claims it does. But, on the other horn, if we shrink the phenomenon of consciousness to exclude self-consciousness and thereby include the experience of unself-conscious creatures such as frogs as conscious, then it seems to me the dual strategy I adopted – of (1) generalizing Dennett's idea of access relations, and (2) defusing McDowell's epistemological alarms which motivate his criticism of Dennett – provides a quite adequate defence of Dennett against McDowell.

*Columbia University*

## BRAINS IN A VAT: DIFFERENT PERSPECTIVES

BY YUVAL STEINITZ

In his 'Brains in a Vat'<sup>1</sup> Putnam argues with great *panache* against ontological, external-world scepticism. Ultimately, however, his arguments are unconvincing. In fact, they seem disingenuous in places.

This paper has two parts. In the first I argue that even if Putnam's linguistic arguments are enough to convince a modern linguistic sceptic (like Putnam himself), they are certainly insufficient to win over an old-school 'ontological sceptic' – that is, a sceptic who sees ontology as a field of philosophy more fundamental than the theory of reference (like Descartes at the outset of his *Meditations*). As I shall explain, Putnam's reliance on causal semantics to achieve his anti-sceptical result would beg the central question against certain kinds of sceptics. In the second part I argue that Putnam's arguments are not sufficient to win over his own camp of linguistic philosophers, not even when his causal semantics is taken for granted. His arguments, if valid, might undercut the possibility of what I shall call a comprehensive sceptical scenario, but they have little to say about the possibility of what I shall call partial sceptical scenarios.

Before presenting my criticism, let me briefly summarize Putnam's 'linguistic' argument against external-world scepticism. Putnam argues that sceptical propositions like 'I am a brain in a vat' (henceforth BIV), that is, a brain subjected to a sophisticated computer which performs comprehensive simulation of reality, are self-refuting.

<sup>1</sup> H. Putnam, *Reason, Truth and History* (Cambridge UP, 1981), pp. 1–21; see also the related discussion pp. 49–74. All page references are to this book.

His argument is as follows: in order for this proposition to acquire its presumed realistic meaning, the terms 'brain' and 'vat' must *refer* to *real* brains and vats. However, according to causal semantics, a word can refer to an object only if that object is somehow causally connected with the utterance or thought of the word. Thus in order for the word 'vat' to refer to an actual vat, and not to refer to a 'vat in the image' (p. 15), or be a mere cluster of random, meaningless markings, there must be some causal connection between the perception of some actual vat and the fact that the word 'vat' was written, said or conceived.

Returning now to the BIV proposition: if the speaker is a genuine BIV, then it is impossible that he should ever have observed the vat he is in, or any vat, brain or computer whatsoever. Thus the word 'vat' in his usage does not refer to real vats, but only to 'vats in the image', and the entire proposition fails to refer. If, on the other hand, the speaker is not a BIV, then his proposition has reference, but is obviously false. Hence the proposition 'I am a BIV' either fails to refer or is false.

But if it fails to refer to the vat containing the brain, i.e., if it refers to mental images rather than to real objects, then it is also false (p. 15). For what it would (actually) tell the assumed BIV is that he is a BIV in the image, while what he would (actually) be is a BIV in reality. So it is necessarily false.

Taken together, the two parts of this paper place Putnam in a dilemma. On the one hand, as the sceptical scenario (e.g., the BIV story) is rendered more *comprehensive*, it becomes increasingly necessary, from the sceptic's perspective, to discard the realistic, causal theory of reference upon which Putnam's anti-sceptical argument is based. On the other hand, as the sceptical scenario is rendered more *partial*, indirect causal connections with external reality, which are sufficient for realistic reference, become increasingly plausible.

## I. THE ONTOLOGIST'S PERSPECTIVE

The BIV hypothesis is an updated version of the sceptical arguments presented in Descartes' first *Meditation*, that the entire world might be only a self-made figment of his imagination, or rather a fiction generated by an evil genius. The first line of defence for the sceptic should be to stress the disjunctive nature of this claim:

I cannot know whether (a) I am a bodily person who perceives independent reality, *or* (b) everything is only a dream in my mind.

Adapting this formulation to the BIV hypothesis, the sceptical disjunction can be formulated thus:

I cannot know whether (a) I am a bodily person who perceives independent reality, *or* (b) I am a BIV.

An unprejudiced consideration, the sceptic will argue, will find the two disjuncts to be equally plausible. There can be no reason for preferring one over the other, since any possible experience would be compatible with both.

Does this emphasis on the disjunctive nature of scepticism pose a new challenge from Putnam's perspective? At first glance it seems not. Putnam could elegantly

tackle the above disjunction much as he responded to the BIV proposition alone: he might argue that, for a BIV, the disjunction 'I am a bodily person or a BIV' fails to acquire its assumed reference and meaning – because, for a BIV, the terms 'body', 'brain', 'computer' and 'vat' fail to refer. It follows that if the sceptical disjunction is meaningful in the presupposed common sense, i.e., if it acquires its assumed reference to a real body, brain, computer and vat, rather than referring to purely mental entities, then disjunct (a) must be true.

But here, I contend, we can see how from an old-style ontologist's perspective, e.g., from the perspective of a Cartesian for whom ontology is the most fundamental philosophical field, Putnam's argument indeed begs the question. Being a philosopher of language at heart, Putnam decides the question of the theory of reference before deciding major ontological issues, such as the existence of an external, perceivable world. As Casati and Dokic<sup>2</sup> briefly remarked, Putnam presupposes a 'meta-language' which is *inappropriate* for a real BIV, and only then addresses the sceptical disjunction. It turns out that his apparent refutation of the BIV propositions rests upon the prior assumption that the speaker is, after all, a bodily person who has a direct perceptual contact with reality, or that he at least tacitly believes himself to be such a person, or, alternatively, that he at least tacitly believes most members of his society to be such persons.<sup>3</sup> For it is only by relying on such *realistic* presuppositions (or on some *internal realistic* presuppositions, as argued by Dell' Utri<sup>4</sup>), that the assumed causal theory of reference becomes acceptable. And such realistic premises do indeed beg the question.

Of course Putnam might respond that, from his linguistic perspective, this 'begging of the question' is no liability at all, but rather a sign of the validity of his argument. For, as Geach remarks,<sup>5</sup> it is the mark of a valid deductive argument that its conclusion is implied by its premises. But if Putnam's argument is intended to convince a sceptical ontologist, he should recognize that the above 'question-begging' will, for such a sceptic, entirely undercut the force of his arguments. For the ontologist maintains, as I shall soon explain in greater detail, that choosing a theory of reference is conditioned by one's chosen ontology, and not *vice versa*.

Consider, for example, the above sceptical disjunction. The Cartesian approach sees this doubt about the nature of man, which may be compressed into the well-known conundrum 'What am I?', as a key issue of philosophy. I shall now argue that this fundamental scepticism, which Putnam claims to have rendered impossible,

<sup>2</sup> R. Casati and J. Dokic, 'Brains in a Vat, Language and Meta-Language', *Analysis* 51 (1991), pp. 91–3.

<sup>3</sup> On the social background of Putnam's linguistic argument see Massimo Dell' Utri, 'Choosing Conceptions of Realism: the Case of the Brains in a Vat', *Mind* 99 (1990), pp. 79–90 (cf. p. 83).

<sup>4</sup> Dell' Utri p. 90. His reconstruction suggests that Putnam's argument against scepticism presupposes internal realism. I have opted, however, to agree with Brueckner, 'Brains in a Vat', *Journal of Philosophy* 83 (1986), pp. 148–67, note 2, that this argument should be sharply distinguished from Putnam's internal realism as expressed in the rest of the book (all the more so since Brueckner's interpretation was confirmed by Putnam himself). I was thus surprised to discover in *Realism with a Human Face* (Harvard UP, 1992), ch. 7, that Putnam reiterates his BIV argument as an outcome of internal realism. I contend, however, that my argument against Putnam works against both interpretations.

<sup>5</sup> P. Geach, *Reason and Argument* (Oxford UP, 1976), p. 18.

assumes that any decisions concerning the right theory of reference are absolutely dependent upon decisions about central ontological issues. Therefore, from the ontologist's point of view, the ontological issues involved have *clear precedence*.

Putnam might still resist by claiming that since ontological sceptics *state* their views and argue for them, the question of whether they assert anything – whether their words have real meaning and reference – must in some sense be prior to whether what they say is true.

This, I believe, is partially true about meaning, but it must be wrong about reference. As Quine emphasized, a reasonable theory of reference must allow us to say what there is not. And since there can be no reference to what there is not, anyone who denies the existence of the external world (or who denies the existence of ghosts, unicorns, etc.) must not be disturbed by the obvious fact that his words 'external world', 'ghosts' and 'unicorns' do not refer. This seems in line with Putnam's own argument, for though he accuses the BIV-sceptic of being incapable of referring to the real vat which contains him *according to his story*, he does not accuse him of being incapable of referring to anything whose existence he denies, e.g., his hands, his friends, etc. (and we shall see in the imagined dialogue below where it would lead him if he did).

Coming back to the ontologist's point of view, I think the best way to illustrate the significance the above difference of approach might acquire is by considering Berkeley's empirical idealism. Obviously, the theory of reference *derived* from an ontological stance such as Berkeley's is incompatible with Putnam's realistic causal theory of reference.

Indeed, both Berkeley and Putnam were looking for a way to escape the scepticism Descartes unfolds in the first chapters of the *Meditations*. But Berkeley preaches a different solution. He maintains that the best method of overcoming scepticism is to neutralize it by giving up realism and then re-adjusting the theory of reference accordingly.<sup>6</sup> He would thus claim that words such as 'brain', 'vat', etc., should not be conceived as denoting *objective objects* which are independent of any perceiver, nor as denoting *subjective objects*, i.e., objects which owe their existence, *viz.*, their internal reality, to the linguistic practice of any particular society, but rather as referring to the common perceptions of God and man, perceptions projected by God onto His creatures' minds.

If Berkeley's rebuttal of scepticism is internally coherent, then Putnam's argument, claiming any surrender of realism to be self-refuting, is not responsive to it. Putnam, quite surprisingly, does not endeavour to demonstrate that idealistic or solipsistic approaches cannot be accompanied by compatible theories of reference which make their external-world scepticism self-coherent.<sup>7</sup>

These considerations suggest that there is another and indeed a major difficulty with Putnam's (and Dell' Utri's) responses to traditional scepticism: *the BIV story does*

<sup>6</sup> *Principles of Human Knowledge*, secs 87–90.

<sup>7</sup> A transcendental argument for that purpose is set forth by Kant, who asserts that our concept of self-identity acquires its meaning only against the assumed background of external reality: *Critique of Pure Reason* B 276. This view actually implies that the *Cogito* is already sufficient to confute external-world scepticism.

*not constitute an adequate version of old, external-world scepticism.* Putnam claims that a real BIV is incapable of expressing scepticism because his words fail to refer. Or, as suggested by Dell' Utri's reconstruction, an actual BIV would not have the (internal) realistic language required for expressing his fears and doubts. But what force does this claim have against, say, an idealistic solipsist, who claims the entire world to exist in his mind, without assuming a single external entity? Such a solipsist can think of himself, i.e., his soul, as the only thing that exists, without contradicting himself by appealing to realistic language.

Thus Putnam's argument against external-world scepticism (as well as Dell' Utri's version of it) is a bit tricky: the representative sceptic is held to doubt everything outside his mind, except his brain, vat and computer, which do lie outside his mind. He is then found incapable of forming direct or indirect perceptual connections with those isolated islands of external reality left behind. And this renders his scepticism self-refuting.

This shows that Putnam's refutation of external-world scepticism is heavily dependent upon, and hence limited to, a specially tooled sceptical hypothesis, in which there is an *external deceiving agent*. Where this deceiving agent is further assumed to be of a materialistic nature, e.g., the vat-computer system, Putnam can argue that it has to be *perceptually* connected to the sceptic, in order to be referred to. Now, in order to see more clearly why Putnam's argument is entirely ineffective against other, more conclusive expressions of external-world or material scepticism, consider the results were he to try to apply it to Berkeley's idealism:

*Berkeley:* All there is are souls and ideas.

*Putnam:* This is self-refuting. Indeed, any external-world scepticism is self-refuting, as I have demonstrated in the case of BIV, for if it is true, then the sceptic's words and ideas fail to refer to those material, external objects which constitute his deluding environment. So if we are brains in a vat, we cannot say or think that we are.

*Berkeley:* But my story is different! In my story there are no material, external objects, and hence no causal connection with such objects is necessary for their description.

[Putnam would do well to end the discussion at this juncture. For were he to press his point, he would be forced to argue that the existence of an external, objective referent is a necessary condition not just for reference but for meaningfulness too. He would have to do this in order to extend his linguistic argument to Berkeley's idealism, that is, in order to doubt the meaningfulness of Berkeley's entirely negative existential propositions about an objective, external reality. But such a view has unwelcome ontological consequences.]

*Putnam:* This is wrong too. True, your story does not contain any external, material objects. But have you not just spoken about such objects, while denying their very existence? Thus, according to my causal theory of meaning, there must be *some* causal, perceptual connection between your words or ideas of material objects and at least one exemplar of such an object, in order for

your words to acquire their prevalent meanings, though you believe them to have no reference at all. So either there are material objects or there are not. If there are, your external-world scepticism is false. If there are not, then you could not say and mean that there are not.

*Berkeley*: I now see the true nature of your argument. Your demand for a causal connection with the object we speak about, or at least for a causal connection with other objects of the same kind, is none other than a disguised version of the ancient Parmenidean doctrine that 'It cannot be said that anything is not'.<sup>8</sup> Indeed, you ostensibly alter it by claiming that 'It cannot be said that anything is causally unconnected'. But since there can be no causal connection with what there is not, your demand boils down to the same ontological consequences, that is, the inflated ontology of objects (or at least kinds of objects) that issues from the old riddle of non-being.<sup>9</sup>

In the last few pages we have seen why Putnam's anti-sceptical argument is ineffective against widely comprehensive sceptical scenarios which permit no external reality (solipsism), or at least no material reality (idealism). Let us summarize the discussion so far by stressing two major points:

1. Putnam argues that sceptical assertions such as 'Reality is an all-encompassing dream', or 'I am a BIV', are self-refuting, and triumphantly claims that he has revealed a substantial inconsistency in ontological scepticism. But I have suggested that Putnam's satisfaction would be warranted only if his realistic causal theory of reference were true *come what may*. I have also demonstrated that Putnam might find it quite difficult to force a stubborn ontological sceptic to accept his causal theory of reference in the first place. For the ontological doubts raised by such a sceptic are not restricted to ontology. Rather, they spread to the prevalent causal semantics as well.

Note, then, that the contradiction that Putnam claimed to have revealed, between the causal theory of reference on the one hand, and some major sceptical arguments on the other, would not immediately win over a true ontological sceptic. In fact, this contradiction might only serve to emphasize how extensive the ramifications of traditional ontological scepticism really are.

2. Instead of being true *come what may*, and instead of being an elementary hypothesis or an axiom about the theory of reference, Putnam's causal theory of reference is clearly *dependent* upon an implicit ontological stance: either upon external realism, as suggested by Brueckner, or upon internal realism, as suggested by Dell' Utri.

Well, why should not Putnam adopt one of these 'conceptions of realism' (as Dell' Utri puts it) as his premise? Take, for instance, internal realism: is it not possible, at least in principle, that this theory is indeed *necessary* for thinking and referring, as Putnam and Dell' Utri suggest?<sup>10</sup>

<sup>8</sup> See M. Furth, 'Elements of Eleatic Ontology', *Journal of the History of Philosophy* 6 (1968), p. 116.

<sup>9</sup> These unwelcome consequences of Putnam's argument were first pointed out in *Philosophia* 16 (1986) by P. Tichy, 'Putnam on Brains in a Vat', pp. 144-5, and by F.B. Farrell, 'Putnam and the Vat-People', p. 154.

<sup>10</sup> Putnam ch. 3; Dell' Utri sec. V.

This is not to the point here. Assuming that external realism and internal realism are false, then Putnam's causal semantics goes too. In that case, Putnam's anti-sceptical argument, and Dell' Utri's reconstruction of it, go as well. Assuming, on the other hand, that one of these ontological stances happens to be true, then their arguments against scepticism are valid but *useless*. For if external or internal realism is true, then all forms of external-world scepticism, e.g., the BIV story, Solipsism, Idealism, the Evil Genius story, etc., *are instantly falsified by this very assumption*. The question has been explicitly begged at the outset. There is no need for Putnam's BIV reference story.

## II. A BIV PERSPECTIVE

In the first part of this paper I showed that Putnam's argument is ineffective against old-style ontological scepticism. This is because scepticism of this sort might become so comprehensive that it must influence the theory of reference preferred by the sceptic.

In the following pages I shall evaluate the force of Putnam's arguments against a more moderate sceptic, who considers a single BIV placed in a real, rich material surrounding, and who happens to take for granted Putnam's causal theory of reference. Must he accept, under those assumptions, Putnam's solution of BIV-scepticism?

Putnam presents the BIV hypothesis as a story about a mad scientist who performs the experiment. The following story provides a reasonable sequel, in the sense that it does not contradict any major physical principles or the laws of logic. Its purpose is to show that the hermetic separation between a BIV, on the one hand, and objective reality, on the other, might not be as air-tight as Putnam maintains. It would follow that even if a sceptical BIV were to accept Putnam's realistic causal theory of reference (as no Cartesian BIV would), still there might be circumstances in which his assertion 'I am a BIV' would acquire its assumed realistic reference, and hence be meaningful and succeed in referring in the ordinary sense.

Imagine a brain in a vat named Alex. Apart from the unfortunate fact that his nerve-endings are connected to a computer, Alex is a normal person in all other respects. Like most people, Alex does not belong to the category of those who waste their time on philosophical meditations. Doubts such as 'Am I a BIV?' have never entered his mind.

One day it occurs to the mad scientist to document his sensational experiment. He buys a video camera, films the BIV system and the computer controlling it, and that very evening screens the film for his friends. Suddenly he realizes that his circle of friends includes, in some weird sense, Alex as well. It would be unfair, he thinks, to prevent Alex from seeing the film: 'After all, haven't I hurt him enough by putting him into the vat?'

But how could Alex watch the film? And how could he be prevented from realizing that he himself is the film's hero, which might jeopardize some psychological aspects of the experiment?

In order to accomplish this, the mad scientist decides that it would be best if the computer made Alex meet him somewhere, after which he would invite him to watch the film at his home. Since the film shows nothing about the way the brain was



removed from its body and put into the vat, there is no reason for Alex to realize that he is looking at himself.

And so it was. The computer performed a perfect simulation of a party at which the mad scientist invited Alex to his home. There he showed him the film about the BIV. (Obviously our amazed hero watched the film through the electrodes permanently attached to his visual nerves, just as he saw all the other things that constituted his world.)

Nevertheless, a few days later it occurred to Alex for the first time in his life to raise a sceptical question: 'Might I be a BIV myself?'

After turning this doubt over in his mind day and night, he finally went to the library (obviously a simulated library) to read an essay entitled 'Brains in a Vat'. On page 50 he came upon the following conclusion: 'So, if we are Brains in a Vat, we cannot think that we are'. 'That poor brain I saw in the film', thought Alex. 'Even if it had suspected its true predicament, that very thought on its part would be self-refuting.' And so he sighed in relief.

The moral: Alex considered the possibility of being a BIV, and conceived the concepts 'brain' and 'vat' in his mind, as a result of some causal, perceptual connection with the real vat containing him. Hence his words and doubts seem to acquire their assumed reference and meaning.

Does the fact that his observation took place by means of a film projected from a computer memory, and transmitted to his brain via electrodes, prevent the fulfilment of the required causal connection?

Putnam, quite surprisingly, claims that it does. Here he sharply differs from Kripke, another proponent of the causal theory of reference, who argues that even the most indirect and heavily mediated causal connection is sufficient for the establishment of full referential relations between the concepts or names and the objects, through what he calls 'a causal (historical) chain'.<sup>11</sup> According to Kripke, a blind man can refer to light, even though he has never seen it, or anything similar to it. Thus, if Alex opted for the Kripkean interpretation of the causal theory of reference, he would be 'allowed' to refer to vats and computers, even though he grasped their concepts through discourse with a computer, on the condition, of course, that those concepts were programmed into the computer by someone whose concepts of *vat* and *computer* do refer.

Putnam (pp. 11–12), in contrast, dismisses the referential value of such very weak causal connections:<sup>12</sup>

Although the machine does not *perceive* apples, fields, or steeples, its creator-designers did. There is *some* causal connection between the machine and the real-world apples, etc., via the perceptual experience and knowledge of the creator-designers. But such a weak connection can hardly suffice for reference.

<sup>11</sup> S. Kripke, *Naming and Necessity* (Oxford: Basil Blackwell, 1980), p. 139.

<sup>12</sup> Putnam's position here concerning the insufficiency of such weak connections contrasts not only with Kripke's, but also with his own earlier writing on this topic: cf. 'Explanation and Reference', reprinted in his *Mind, Language and Reality* (Cambridge UP, 1975), p. 202.

And he justifies this with the following argument:

The machine is utterly insensitive to the *continued* existence of apples, fields, steeples, etc. Even if all these things *ceased* to exist, the machine would still discourse just as happily in the same way. That is why the machine cannot be regarded as referring at all.

But this view is indeed very odd. It implies that when we read in a book about penguins in Antarctica, we should claim that the word *penguin* cannot refer to real penguins, as the written word and the book are no less insensitive than the computer to the continued existence of penguins. Is the assumption that this book might continue to be read 'just as happily in the same way', even if penguins 'ceased to exist', enough to disqualify the book's referential value? And if so, then what would be the reference of the entry 'penguin' in the insensitive 1930 edition of the *Encyclopaedia Britannica*?

Putnam might argue that the book could in principle be re-edited and updated, and that the 1930 edition could be replaced by the 2030 edition. But then, in exactly the same way, the mad scientist conducting the BIV experiment could update the computer memory concerning the extinctions, should they unfortunately occur, of penguins, apples, or steeples. He could also program the computer to update itself periodically, perhaps even using the *Britannica* year-books. Incidentally, these volumes are printed nowadays from information that is processed and stored in a computer memory, and transmitted to the printer via electronic mediations.

Thus we must opt for Kripke's version, and stipulate that the mere fact that sensations are processed and transmitted by a computer, or by any artificial or non-artificial means whatsoever, cannot eliminate the required causal connection between a word and its supposed *referendum*, on the condition that such a connection does exist in the first place. In 'a causal (historical) chain' the first link is what matters: through it the whole chain should be perceptually linked to reality. All other links need not be perceptual, not even the last link which conveys the term *light* to a blind man, or the term *vat* to a BIV. Therefore, if Alex suspects, as a result of his computer-mediated watching of a film (or even as a result of a computer-mediated reading of a description of vats), that he might be a BIV, his doubts are ultimately meaningful and capable of referring.

What then should our sceptic believe, if he suspects his true predicament, while assuming the existence of a rich independent world of objects outside of his own world? Surely, being sceptical, he cannot reject the possibility (a possibility acknowledged by Putnam as well) that the computer's memories and programs are at least partly a function of external reality. Therefore it would be only reasonable for him to hypothesize (even though he can never verify it) that the words 'brain' and 'vat' in his language *may* acquire a real reference after all.

I think it has become evident by now that the only way to preserve Putnam's argument against BIV scepticism is to adopt a more rigid version of the BIV story, one that eliminates any causal linkage with external reality whatever. Not surprisingly, indeed, Putnam hints at this second line of defence when he moves to consider a strange possible world in which 'every sentient being is and always was a brain in

a vat' (p. 14). Such a comprehensive sceptical scenario, he seems to believe, must be either self-refuting or false.

But this is cold comfort. First, because it returns us to the first horn of the dilemma, i.e., to our previous conclusion that such a *comprehensive* ontological scepticism necessitates changing the theory of reference accordingly. Second, because even if Alex refused to accept this, and steadfastly held to Putnam's realistic causal theory of reference, he still could not shake off scepticism. Sceptics are also concerned about *partial* sceptical scenarios. So the actual disjunction Alex is facing is not a mere double but a *triple* disjunction:

- (a) I am a person with body, perceiving objective reality;
- (b) I am a BIV controlled by a computer, but part of the 'computer's simulation of reality originates in the observations of its programmers (or, alternatively, in the observations of attached robotic sensory mechanisms), who do perceive objective reality;
- (c) I am a BIV, and there is nothing else outside my system. (Or, alternatively, all scientists and programmers are themselves BIVs.)

Of course, Putnam may happily elaborate upon the possible realization of the last comprehensive scenario (c), which permits him (on the condition of neglecting our previous arguments of section I) to rely upon his hypothesis that scepticism, in such a case, is self-refuting...

Yet this should not comfort him, nor should it comfort any BIV-sceptic. For as long as disjunct (b) cannot be excluded, scepticism seems to remain as plausible as ever.<sup>13</sup>

*Tel-Aviv University and University of Haifa*

<sup>13</sup> I would like to thank Naomi Goldblum, Noah J. Efron and Zev Bechler for helpful comments.