

Closure Scepticism and The Vat Argument

JOSHUA ROWAN THORPE

University of Campinas

joshthop@gmail.com

If it works, I can use Putnam's vat argument to show that I have not always been a brain-in-a-vat (BIV). It is widely thought that the vat argument is of no use against closure scepticism – that is, scepticism motivated by arguments that appeal to a closure principle. This is because, even if I can use the vat argument to show that I have not always been a BIV, I cannot use it to show that I was not recently envatted, and it is thought that the claim that I am not justified in thinking that I was not recently envatted is all that the closure sceptic requires. In this paper I first argue that scenarios in which I have been recently envatted are inadequate for the sceptic's purposes, and so the standard argument that the vat argument is of no use against closure scepticism fails. I then argue that it is not possible to revise the standard argument to meet my objection. I conclude that, if it works, I can use the vat argument as a defence against closure scepticism.

If it works, I can use Putnam's (1981, pp. 1-22) vat argument to show that I have not always been a brain-in-a-vat (BIV). Irrespective of whether the argument is adequate for Putnam's purposes, it is widely thought that the vat argument is of no use against scepticism of the sort most commonly considered in contemporary epistemology, namely closure scepticism, so called because it is motivated by arguments that appeal to a closure principle.¹ The vat argument is thought to be of no use against closure scepticism because even if I can use it to show that I have not always been a BIV I cannot use it to show that I was not recently envatted, and it is thought that the claim that I am not justified in thinking that I was not recently envatted is all that the closure sceptic requires.² As a result, the vat argument is not much discussed by contemporary epistemologists.

In this paper I will first show that, surprisingly, scenarios in which I have recently been envatted are inadequate for the closure sceptic's purposes. This is because a sceptical argument that is based on such a

¹ Putnam intends the vat argument to undermine a position he calls metaphysical realism. There will be no discussion in this paper of whether Putnam succeeds in realising this aim.

² This argument can be found in Smith (1984), Glymour (1982), Farrell (1986), Tymoczko (1989), Wright (1992), Christensen (1993), Forbes (1995), and Brueckner (1999).

scenario leaves the justification of some of my empirical beliefs intact, and so it is possible for me to appeal to these beliefs to show that I have not recently been envatted. The sceptic thus faces a dilemma: either she makes use of a scenario in which I have recently been envatted, and so enables me to appeal to my empirical beliefs to show that I have not recently been envatted; or she makes use of a scenario in which I have always been a BIV, and so enables me to use the vat argument to show that I have not always been a BIV.

I will then argue that a more general version of this dilemma holds: either the sceptic makes use of what I will call a non-radical sceptical scenario, and so enables me to appeal to my empirical beliefs to show that I am not in that scenario; or she makes use of what I will call a radical sceptical scenario, in which case I can use the vat argument to show that I am not in that scenario.

I conclude that if it works, the vat argument is of use against closure scepticism. I hope that this conditional conclusion will promote discussion of whether the vat argument does work in contemporary epistemology.

1. Closure scepticism

I am now going to describe the general form of the argument for closure scepticism. My aim here is to be as generous as possible to the sceptic so as to offer the prospect of defeating her on her own terms; hence, I shall grant her much that she perhaps should not be granted.

Let us suppose that the sceptic frames her argument in terms of justification rather than knowledge.³ She proposes the following necessary condition for justified belief:

- (J) For any belief *b* that I hold, *b* is justified only if my evidence indicates to me that *b* is more probable than not-*b*.

The evidence that I may have for an empirical belief consists of sensory experiences, and of other justified beliefs.

The notion that the justification of a belief depends, at least in part, upon whether my evidence indicates to me that the belief is more likely than its negation is not universally accepted. However, thinking about justification in this way gives the sceptical argument that I am about to discuss a great deal of its initial plausibility. In fact, one

³ I think that, at least on some conceptions of knowledge, much of what I have to say in this paper would be applicable to a sceptic who framed her argument in terms of knowledge. However, I do not have the space to argue this here.

reason that is often given for preferring some other notion of justification is that doing so makes it harder to motivate scepticism.⁴ It is worth considering whether the sceptic can be defeated even whilst assuming an account of justification that seems to be particularly conducive to her argument, even if one does not subscribe to that account.

One might accept the notion that the justification of one of my beliefs depends, at least in part, upon whether my evidence indicates to me that the belief is more likely than its negation, whilst questioning whether the necessary condition that I propose above is appropriate in this context. If my aim is to be concessive to the sceptic why not work with something stronger? Suppose, for example, that we accept the following strengthened version of (J) as a necessary condition for justified belief:

- (SJ) For any belief *b* that I hold, *b* is justified only if my evidence indicates to me that *b* is *to some very high degree* more probable than not-*b*.

Won't it be easier for the sceptic to show that my beliefs do not meet (SJ) than for her to show that they do not meet (J)? After all, my beliefs might meet (J) without meeting (SJ).⁵

It is a limitation of this paper that the discussion is framed in terms of a concept of justification according to which (J) is a necessary condition, and so any progress that the paper makes against a sceptic who works with such a concept of justification will not immediately be progress against a sceptic who works with a conception of justification according to which (SJ) is a necessary condition.⁶ However, I do not think that this limitation is severe. One reason for working with (J) rather than (SJ) is that the latter invites the question: to what degree must my evidence indicate to me that *b* is more probable than not-*b* in order for *b* to be justified? It may seem that there is no non-arbitrary answer to this question, which may in turn lead one to doubt that (SJ) is a condition on justification.

However, my main reason for working with (J) rather than something like (SJ) is that I do not wish to deprive the sceptical conclusion

⁴ Sosa (1997) suggests that we should endorse a reliabilist account of justification in order to avoid scepticism. Pryor (2000) motivates a dogmatist account of justification in a similar way.

⁵ I am grateful to an anonymous Mind referee for raising this point.

⁶ In fact I think it is likely that, with some tweaks, the discussion in this paper could be framed in terms of a stronger concept of justification.

of its teeth. Suppose a sceptic showed that my beliefs do not meet (SJ). This might be a surprising and disappointing conclusion, but for all that had been said, I might still claim that, nevertheless, many of my beliefs meet (J). If this is correct, then many of my beliefs possess a positive epistemic property, namely, the property of being supported by evidence that indicates to me that they are more likely to be true than their negations. For this reason, one might think that it would still be reasonable for me to continue to believe by and large what I do. But if the sceptic shows that my beliefs lack (J) this is a disaster: for all that my evidence indicates to me about the probabilities of my beliefs being true, I might just as well believe their negations. In that case it would be completely unreasonable for me to continue to believe as I do. That is a disastrous conclusion, and the fact that it is disastrous is what provides a lot of the motivation to resist the sceptical argument. For this reason, I will assume in this paper that (J), but not (SJ), is a necessary condition for justification.⁷

I now want to make a distinction between two sorts of sceptic. One sort of sceptic denies that I am entitled to think that most of my empirical beliefs are justified until I *prove to her*, relying only on non-empirical premises, that most of my empirical beliefs are justified. It is not immediately clear why I should respond to such a sceptic. Why should I have to prove an obvious claim before I am entitled to it? The sort of sceptic that I want to consider in this paper does not make such an unreasonable demand, regarding my empirical beliefs as ‘innocent until proven guilty’. She is happy to admit that I am entitled to think that most of my empirical beliefs are justified until someone gives me a good argument to think otherwise. If I cannot say what is wrong with this argument, I lose my entitlement to think that I have a lot of justified beliefs.⁸

The sceptic’s attempt to give me a good argument that my empirical beliefs are not justified proceeds as follows. Let *SS* stand for a sceptical scenario, and let *p* be an arbitrary empirical proposition that I

⁷ Similar reasons might be given for considering scepticism about justification, as we are in this paper, rather than scepticism about knowledge. One response to the sceptical claim that I have no knowledge is to accept it, but to continue to maintain that at least many of my beliefs meet (J), so it is reasonable for me to hold them. If this retreat can be made, one might think that scepticism about knowledge is not as worrying as it may at first have seemed.

⁸ See Byrne (2004, pp. 300-3) for discussion of the distinction between these two kinds of sceptic, and for a convincing argument that we need only be concerned with the second kind.

Indeed, it is unclear what the first sort of sceptic would need a closure argument for, since she does not need a sceptical *argument* at all; she operates by demanding that I make an argument that my empirical beliefs are justified, and then attempts to pick holes in it.

justifiably believe to be logically incompatible with that scenario. We can then give the following argument schema:

- (C1) If I do not have a justified belief that *SS* is false, then I do not have a justified belief that *p*, but
- (C2) I do not have a justified belief that *SS* is false; so
- (C3) I do not have a justified belief that *p*.

Call any argument that is an instance of this schema a *closure argument*. Any sceptical conclusion resulting from a closure argument is an instance of *closure scepticism*.

In order to support (C1) the sceptic appeals to a closure principle such as the following:

Closure principle: For any subject *S*, and propositions *p*, *q*: if *S* has a justified belief that *p*, and *S* has a justified belief that *q* is logically incompatible with *p*, then *S* has a justified belief that not-*q*.

This principle may stand in need of refinement if it is to avoid counterexamples. However, we will concede to the sceptic that none of these refinements will prevent her from using the principle to support (C1).⁹

SS is, by definition, a scenario that I justifiably believe to be logically incompatible with *p*, and so (C1) is implied by the closure principle. The sceptic's appeal to the closure principle highlights the fact that the conclusion of the closure argument is *only* that the beliefs that I justifiably believe to be incompatible with *SS* are not justified. This will be important later.

In order to support (C2) the sceptic makes *SS* a scenario in which my sensory evidence would be exactly the same as it currently is.¹⁰ Thus, the sceptic claims, my sensory evidence does not indicate to me that it is more probable that *SS* than that not-*SS*, and so I am not justified in believing not-*SS* on the basis of my sensory evidence alone.

We might immediately object that it does not follow from the fact that my sensory evidence would be the same in each of two

⁹ Brueckner (2010, pp. 280-305) argues convincingly that the prospects for a plausible version of the closure principle that does not support (C1) are dim. Someone who wanted to halt the sceptic here would probably have to give up on the closure principle altogether; see for example Dretske (1992) and Nozick (1981). But most people do not find this option attractive, and if the response to closure scepticism that I discuss in this paper is correct there is no need to take it.

¹⁰ Thus, one way to resist the closure argument is to deny that my evidence could be the same if I were in a sceptical scenario. See, for example, Williamson (2000, pp. 164-208).

incompatible scenarios that my sensory evidence does not indicate to me that I am in one scenario rather than the other. Consider the following scenario:

Laptop: when I turned around to stare out of the window just now my laptop spontaneously disappeared. I turn back to look at the spot previously occupied by my laptop. However, by some trick of the late evening light reflecting off the glass table at which I am sitting, I have the sensory experiences that I would have had if my laptop were still in front of me.

By stipulation, in *Laptop* my sensory evidence would be exactly the same as it currently is (at least so long as I do not try to touch the laptop). However, it is natural to think that I have a justified belief that *Laptop* is false.

This thought is correct, and as it will play an important role later on it is worth making explicit why it is correct. Recall that my evidence for a belief can consist of my sensory evidence, and of other justified beliefs. My sensory evidence alone may not indicate to me that *Laptop* is false; however, the conjunction of my sensory evidence and evidence consisting of certain justified empirical beliefs that I have does indicate to me that *Laptop* is false. For example, I have a justified belief that it is vastly improbable that the light should reflect off the table so as to create the illusion of a laptop. This belief, together with my sensory evidence, does indicate to me that *Laptop* is false. Thus I have a justified belief that *Laptop* is false.

The sceptic is aware that it does not follow from the fact that my sensory evidence is logically compatible with a scenario that my evidence doesn't indicate to me that that scenario does not obtain. She responds by ensuring that *SS* is a scenario that is logically incompatible with any empirical belief that I might appeal to in an attempt to justify my belief that *SS* is false. The conclusion of the relevant closure argument will then imply that these empirical beliefs are not justified. Only justified beliefs can be part of my evidence; thus I cannot assume that I can appeal to these beliefs without assuming that the conclusion of the sceptic's argument is false, and thereby begging the question against the sceptic. So I cannot claim that my sensory evidence in conjunction with my empirical beliefs indicates that *SS* is false. If the sceptic is also careful to choose a sceptical scenario that I cannot rule out by appealing to my non-empirical beliefs then it seems that my belief that *SS* is false is not justified.

The following scenario seems to fulfil the above conditions on an effective sceptical scenario:

Lifelong Envatment: I always have been and always will be a BIV. A computer stimulates my brain in such a way that I have all the experiences that I would have had if I had lived the life I take myself to be living. However, nothing except my brain and its vat exists, and so almost all my empirical beliefs are false.

Lifelong Envatment is an example of a radical sceptical scenario; that is, it is a scenario that falsifies nearly all my empirical beliefs.

Substituting Lifelong Envatment into the closure argument schema we get the following argument:

- (LE₁) If I do not have a justified belief that Lifelong Envatment is false, then I do not have a justified belief that *p*, but
- (LE₂) I do not have a justified belief that Lifelong Envatment is false; so
- (LE₃) I do not have a justified belief that *p*.

Call this the Lifelong Envatment argument

When most of my empirical beliefs are substituted for *p* the closure principle implies (LE₁). I would have the same sensory experiences in Lifelong Envatment as I would have if I were living the life I take myself to be living, so my sensory evidence alone does not indicate that Lifelong Envatment is false. Nor can I appeal to any of my empirical beliefs in order to justify my belief that Lifelong Envatment is false, because the conclusion of the Lifelong Envatment argument implies that all the empirical beliefs that I could otherwise appeal to here are not justified. I cannot, for example, appeal to my belief that the technology to envat people does not exist, unless that belief is justified. This is one of the beliefs that the Lifelong Envatment argument says is not justified, and so to appeal to it and thereby assume that it is justified would be to beg the question against the sceptic. Finally, it does not seem that I can appeal to my non-empirical beliefs in order to justify my belief that Lifelong Envatment is false. So it seems that (LE₂) is true, and the sceptical conclusion follows.

There are varieties of scepticism other than closure scepticism; for example, scepticism motivated by the idea that justification is iterative,¹¹ and scepticism motivated by worries about the impossibility of justification that is not ultimately viciously circular, or that does not

¹¹ Wright (1991) discusses scepticism motivated by the idea that justification is iterative.

entail an infinite regress.¹² However, in this paper I am concerned only with the consequences of the vat argument for closure scepticism.

2. The vat argument

I am now going to sketch what I think is the best version of Putnam's vat argument. I believe that the argument works, but I will not attempt to give a full defence of it here.¹³ My purpose is to discuss the consequences for closure scepticism if the argument works. However, I will make some remarks about features of the vat argument that will be important later.

Here is the version of the vat argument that targets the Lifelong Envatment argument:

- (V1) If Lifelong Envatment is true, then I cannot entertain Lifelong Envatment, but
- (V2) I can entertain Lifelong Envatment; so
- (V3) Lifelong Envatment is false.

(To entertain here means to represent, whether in thought, linguistically, or otherwise.)

(V1) is supported by an appeal to what I will call the *causal constraint on concept possession* (or *causal constraint* for short). The causal constraint says that there are some things such that, in order for a subject to have the concept of that thing, the subject must have been in a particular kind of causal contact with it. Plausibly, the causal constraint is a consequence of any semantic externalist view that aims to endorse the judgements that we naturally make about what is being referred to in certain cases.

Putnam's 'Twin Earth' scenario is an example of one such case. Putnam imagines a planet, Twin Earth, which is identical to earth at the macro level. However, at the micro level, wherever there is H₂O on Earth, there is XYZ on Twin Earth. We are asked to consider the question, what concept does my doppelganger on Twin Earth express with the word 'water'? Most people are inclined to say that he expresses the concept of twater, which has in its extension all and only XYZ, rather than the concept of water, which has in its extension all and only H₂O. A natural explanation for why this judgement is correct is that my doppelganger has had the right sort of causal

¹² Foglin (1994) discusses scepticism motivated by both circularity and regress worries.

¹³ See Button (2013) and Wright (1992) for defences of the soundness of the vat argument.

contact with twater to acquire the concept of it, whilst he has not had the right kind of causal contact with water.¹⁴

The idea is that, without getting into the details of exactly what kind of causal contact is required, it is clear that someone in Lifelong Envatment would not have been in the sort of causal contact with brains necessary to have the concept of a brain. If Lifelong Envatment is true then I have never seen, heard, felt, smelt, or touched a brain. Of course, the causal constraint does not imply that I have to have come into direct causal contact with a thing in order to have the concept of it. Often people acquire the concept of a thing that they have never encountered by talking to someone who has the concept of that thing. However, if Lifelong Envatment is true, I cannot ever have met anyone who has the concept of a brain, because the universe contains nothing except me in my vat. For the same reasons, if Lifelong Envatment is true then I do not possess any of the concepts that would be required to refer to brains via a definite description. However, I need the concept of a brain in order to entertain Lifelong Envatment, so if Lifelong Envatment is true I cannot entertain Lifelong Envatment.¹⁵

Ordinarily (V₂) would not need any defence. However, the sceptic may object that the argument made for (V₁) suggests that I am not justified in thinking that I have the concept of a brain unless I am already justified in thinking that Lifelong Envatment is false, and so I am not justified in thinking that I can entertain Lifelong Envatment unless I am already justified in thinking that Lifelong Envatment is false. If this is true then the vat argument is epistemically circular: I must have justification for its conclusion before I can have justification for one of its premises, and so I cannot use the argument to justify my belief that its conclusion is true.¹⁶

¹⁴ This paragraph is a summary of an argument that appears in Putnam (1973, 1975). Putnam invokes this argument in support of the vat argument in his (1981, pp. 18-9).

I believe that there are good reasons to think that most (perhaps all) plausible views of semantics designed to be compatible with the judgements we make about cases like the Twin Earth case will imply (V₁), whether or not they are externalist. For example, the causal descriptivist views endorsed in Jackson (1998) and Lewis (1984) plausibly entail (V₁). If this is correct, the dialectical appeal of the vat argument is significantly strengthened. However, this is not the place to develop this thought.

¹⁵ See Button (2013, pp. 118-23) for a particularly comprehensive defence of the first premise of the vat argument.

¹⁶ Johnsen (2003) makes a similar objection.

However, if I am entertaining the Lifelong Envatment argument, then I can entertain the Lifelong Envatment scenario. Thus, it follows from the fact that I am entertaining the Lifelong Envatment argument that I have the concept of a brain.¹⁷ At this point it may be tempting to think that there is a sceptical worry about a scenario that I cannot entertain. This is an interesting possibility, but I am not going to pursue it here.¹⁸ I will assume that if there is a sceptical worry that I am in a scenario it is a worry that I can be confronted with directly in a way that involves my entertainment of that scenario.

3. Recent envatment and embodiment

Let us suppose that the vat argument works as a reply to the Lifelong Envatment argument for scepticism. What progress have we made against the sceptic? One might think: none at all, because the sceptic can simply use a different closure argument to secure her conclusion.

Consider the following sceptical scenario:

Recent Envatment: One year ago I was sedated, and my brain was removed from my body and placed in a vat. From then on, a computer stimulated my brain in such a way that I had the experiences that I would have had if I had been living the life that I took myself to be living. Just after I was envatted everything in the universe except my brain in its vat was destroyed, so many of my empirical beliefs are false.

In contrast to Lifelong Envatment, Recent Envatment is an example of a sceptical scenario that is non-radical; that is, there are large chunks of my empirical beliefs that Recent Envatment does not falsify. For example, Recent Envatment does not falsify my empirical beliefs about how things were up to one year ago.

Substituting Recent Envatment for SS we get the following closure argument:

- (RE1) If I am not justified in believing that Recent Envatment is false, then I do not have a justified belief that *p*, but
- (RE2) I do not have a justified belief that Recent Envatment is false; so
- (RE3) I do not have a justified belief that *p*.

¹⁷ Bruckner (2006) makes a similar reply to Johnsen (2003).

¹⁸ Moore (1996, 2011) and Button (2013) argue that we should reject the idea that we should worry that we are in a sceptical scenario that we cannot entertain. Wright (1992, p. 93) can be read as suggesting that we should take the idea seriously.

Call this the Recent Envatment argument

Whatever the exact nature of the causal constraint, it is not going to prevent me from having the concept of a brain if I have spent most of my life interacting with brains and other people who have interacted with brains in the way I ordinarily take myself to. Thus, I cannot reject (RE₂) by using a version of the vat argument to justify my belief that Recent Envatment is false. Smith (1984, p. 117), Glymour (1982, pp. 173-5), Farrell (1986, p. 150), Tymoczko (1989, pp. 294-5), Wright (1992, pp. 86-90), Christensen (1993, pp. 314-5), Forbes (1995, p. 207), and Brueckner (1999, p. 237) have taken this to mean that the vat argument cannot save us from closure scepticism. According to Wright (1992, p. 86), for example, so far as closure scepticism is concerned, the conclusion of the vat argument 'should merely be that some traditional sceptical arguments employ inept examples'. In a similar vein, Christensen (1993, p. 314) writes that the Recent Envatment argument seems to show that the vat argument 'amounts to little more than a philosophical curiosity, without great significance for epistemology'.¹⁹

Although most authors are content to draw this conclusion about the vat argument immediately after reflection on the Recent Envatment argument, more needs to be said in order to make it plausible that the vat argument has no consequences for closure scepticism. When the schema of the closure argument was described in §1 it was stipulated that any proposition can be substituted for *p* in the argument, *so long as I justifiably believe that proposition to be logically incompatible with SS*. This stipulation is necessary if the first premise of the closure argument, which says that if I do not have a justified belief that *SS* is false then I do not have a justified belief that *p*, is to follow from the closure principle.

Recent Envatment is a non-radical sceptical scenario: it is logically compatible with many of the empirical beliefs that I hold. These empirical beliefs cannot be substituted for *p* in the Recent Envatment argument, on pain of preventing the sceptic from supporting (RE₁) by appealing to the closure principle. Thus, whilst the conclusion of the Lifelong Envatment argument seems to imply that I have no justified empirical beliefs (or at least, hardly any) the conclusion of the Recent Envatment argument is compatible with the claim that many of my empirical beliefs are justified. For example, it is compatible with the claim that the beliefs I have about how things have been up to one year

¹⁹ It should be noted that both Wright and Christensen go on to consider the possibility that the vat argument has significance for issues other than closure scepticism.

ago are justified. Since I can appeal to the beliefs I have about how things have been in the past in order to obtain inductive justification for my more general beliefs about, for example, the laws of physics, it appears to be compatible with the claim that these more general beliefs are justified as well.

If I can use the vat argument to defend myself from a sceptic who uses the Lifelong Envatment argument to show that I do not have *any* (or hardly any) justified empirical beliefs, then the vat argument does have significant consequences for closure scepticism. This is so even if I cannot use the vat argument to prevent the sceptic from using the Recent Envatment argument to show that *many* of my empirical beliefs are not justified. The conclusion of the Recent Envatment argument that my beliefs about how things have been this year and about how things are going to be in the future are not justified is a form of scepticism, but it is not nearly as strong a form of scepticism as the conclusion that I have no (or hardly any) justified empirical beliefs at all.

However, it seems that the vat argument cannot save us from the strong form of scepticism after all. The sceptic claims to be able to mount another closure argument that, as it were, mops up the justification for my empirical beliefs that the Recent Envatment argument is unable to wash away. As usual, the sceptic starts by describing a sceptical scenario:

Embodiment: For most of my life I have been a BIV, stimulated by a computer in such a way that I have had the experiences that I would have had if I had been living the life that I took myself to be living. In fact, however, nothing but my brain in its vat existed. Luckily for me things changed dramatically this time last year. My brain was placed in a normal human body, and the world came to be by and large the way that I believe it to be.

This scenario can be substituted for SS, resulting in the following closure argument:

- (E1) If I do not have a justified belief that Embodiment is false, then I do not have a justified belief that *p*, but
- (E2) I do not have a justified belief that Embodiment is false; so
- (E3) I do not have a justified belief that *p*.

Call this the Embodiment argument

If I have spent a year interacting with brains in the way that I ordinarily take myself to, no plausible version of the causal constraint

will prevent me from having the concept of a brain. Thus, the vat argument cannot be used to block the Embodiment argument. However, Embodiment is logically incompatible with precisely those beliefs that Recent Envatment is compatible with (that is, the beliefs I have about how things were before the time of my envatment in Recent Envatment). Therefore, the conclusion of the Embodiment argument is that the beliefs that the Recent Envatment argument did not touch are, in fact, not justified. Thus, the conjunction of the conclusions of the Recent Envatment argument and the Embodiment argument is equivalent to the conclusion of the Lifelong Envatment Argument.²⁰

It seems that, notwithstanding the vat argument, the sceptic can get the strong sceptical conclusion of the original Lifelong Envatment argument by employing both the Recent Envatment argument and the Embodiment argument. If this is true, then the widely held view that the vat argument cannot be used as a defence against closure scepticism is correct. At best, the vat argument only forces the sceptic to be careful about the scenarios she makes use of, and forces her to employ two closure arguments rather than one.

4. The inadequacy of Recent Envatment and Embodiment

I am now going to show that the sceptic's response to the vat argument, as described in the last section, does not work. The Recent Envatment argument and the Embodiment argument are vulnerable to an objection to which the Lifelong Envatment argument is not vulnerable.

The second premise of the Lifelong Envatment argument, (LE2), says that I do not have a justified belief that Lifelong Envatment is false. In §1 we saw that the Lifelong Envatment argument threatens all the empirical beliefs that I could otherwise have appealed to in order to justify my belief that Lifelong Envatment is false. The idea that I cannot appeal to these beliefs in order to justify my belief that Lifelong Envatment is false on pain of begging the question against the sceptic was a crucial part of the sceptic's defence of the claim that I do not have a justified belief that Lifelong Envatment is false.

Now consider the second premise of the Recent Envatment argument, (RE2), which says that I do not have a justified belief that

²⁰ Brueckner (2010, pp. 174-7) suggests a similar strategy for mopping up the justification left over by the Recent Envatment argument.

Recent Envatment is false. Unlike the conclusion of the Lifelong Envatment argument, the conclusion of the Recent Envatment argument is compatible with the claim that many of the empirical beliefs I have about how things were up until one year ago are justified. Thus, it would not be question begging for me to appeal to any of these beliefs in an attempt to justify my belief that Recent Envatment is false.

In fact, I *can* justify my belief that Recent Envatment is false by appealing to some of these beliefs, despite the fact that Recent Envatment is compatible with my sensory evidence. For example, I believe that one year ago no one was even close to having the technology required to envat someone, and I believe that no one had any sort of motivation to envat me. I can appeal to these beliefs, and to similar beliefs that I have, in order to justify my belief that Recent Envatment is false. If I do so, the second premise of the Recent Envatment argument, which says that I do not have a justified belief that Recent Envatment is false, is no longer true. The basic idea here is that my belief that Recent Envatment is false and my belief that Laptop is false are justified for the same reason. In both cases my sensory evidence alone does not indicate to me that the scenario in question is false, but my sensory evidence in conjunction with other justified beliefs that I have does.

The sceptic might try to prevent me from appealing to these beliefs by altering Recent Envatment so that my envatment is not quite so recent. This will result in one of two things. Either the sceptic will not move the time of my envatment back far enough, in which case I will still be able to appeal to some of my empirical beliefs in order to justify my belief that I have not been envatted, or the sceptic will move the time of my envatment back far enough that there is no possibility that I can justify my belief that I am not envatted by appealing to my empirical beliefs. However, by moving the time of my envatment back this far, the sceptic will deprive me of the kind of causal contact with brains required to have the concept of a brain. If this is the case, I will be able to use the vat argument to justify my belief that I am not in the altered version of Recent Envatment.

But couldn't the sceptic specify that my envatment took place at a time when I have formed no beliefs that I can appeal to in order to justify my belief that I have not been envatted, and yet I have acquired the concept of a brain? No. It is a biographical fact about me that by the time that I had acquired the concept of a brain (not to mention the other concepts required in order to entertain any of the sceptic's envatment scenarios) I had already been around for some time. I now

have many beliefs about how things were before I acquired the concept of a brain that I can appeal to in order to justify my beliefs that I was not subsequently envatted. For example, I believe that there was no indication that the technology to envat people existed at that time. I expect that what I am saying here will apply to the reader as well.

Let us turn our attention to the second premise of the Embodiment Argument, which says that I do not have a justified belief that Embodiment is false. Unlike the Lifelong Envatment argument, the Embodiment argument is compatible with the claim that many of the empirical beliefs I have about how things have been this year, and about how things are now, are justified. Thus it would not be question begging for me to appeal to any of these beliefs in an attempt to justify my belief that Embodiment is false.

Again, I can justify my belief that Embodiment is false by appealing to these beliefs, despite the fact that Embodiment is compatible with my sensory evidence. I believe that nothing this year indicates that anyone has ever had the technology to envat someone at birth and then de-envat them later in life without their noticing; there is no old envatting machinery lying around, for example. I can appeal to these beliefs in order to justify my belief that Embodiment is false. If I do so, the second premise of the Embodiment argument will be false.

The sceptic may re-describe Embodiment so that the time of my embodiment is moved forward in the hope that this will prevent me from appealing to my empirical beliefs in order to justify my belief that Embodiment is false. Again, one of two things may happen: either the time of my embodiment will not have been moved forward far enough, and I will still be able to appeal to my empirical beliefs in order to justify my belief that Embodiment is false, or the time of my embodiment will have been moved forward enough to prevent me from appealing to my empirical beliefs, but I will not have been in contact with brains for long enough to acquire the concept of a brain. As a result, I will be able to use the vat argument to justify my belief that I am not in the re-described version of Embodiment.

One natural response that the sceptic might make at this point is that I am right that neither the Recent Envatment argument nor the Embodiment argument is sound on its own, but that both arguments are sound when they are run at the same time. There are two ways in which these arguments might be run 'at the same time'. However, neither way will serve to save the arguments.

First, suppose that the sceptic runs the Recent Envatment argument and I object by saying that I can appeal to some of my beliefs about

how things were before last year in order to justify my belief that Recent Envatment is false. The sceptic replies that she has another argument to show that the beliefs to which I appeal are not justified either; namely, the Embodiment argument. Because they are not justified, I cannot appeal to these beliefs in order to justify my belief that Recent Envatment is false. I counter by saying that I can make the same objection to the Embodiment argument that I made to the Recent Envatment argument; there are empirical beliefs to which I can appeal to justify my belief that Embodiment is false. The sceptic is simply attempting to save one unsound argument by running another unsound argument.

At this point someone might object that when the sceptic presents a closure argument based on a non-radical sceptical scenario that only threatens some of my empirical beliefs, she need not in the meantime concede that the rest of my empirical beliefs are justified. If she does not, my reliance on the rest of my empirical beliefs to rule out the non-radical sceptical scenario is illegitimate.²¹ However, this objection can only be made if we forget the type of sceptic we are dealing with. In §1 I drew a distinction between two sorts of sceptic. The first kind claims that I am not entitled to think that I have justified empirical beliefs until I have argued from non-empirical premises that I do. The second concedes that I am entitled to think that many of my empirical beliefs are justified until I am given a good reason to think otherwise, and then attempts to give me that reason. In this paper I am concerned only with the second kind of sceptic. If such a sceptic attempts to attack only some of my empirical beliefs with one argument, she must concede that the rest of my empirical beliefs are still justified, for no one has yet presented me with a good reason to think that they are not.

Let us turn to the second way in which the arguments can be run 'at the same time'. The sceptic may reply that my counter missed her point. She did not mean that she could run both arguments, one after the other. What she meant was that I am not justified in believing that a scenario that combines the past envatment of Recent Envatment and the future envatment of Embodiment is false. This Combination Scenario can be substituted for SS, resulting in the following closure argument:

- (CS₁) If I am not justified in believing that the Combination Scenario is false, then I am not justified in believing that *p*, but

²¹ My thanks to an anonymous Mind referee for raising this objection.

- (CS₂) I am not justified in believing that the Combination Scenario is false; so
(CS₃) I am not justified in believing that *p*.

Call this the Combination Scenario argument. The conclusion of this argument implies that all the empirical beliefs that I could otherwise have appealed to in order to justify my belief that the Combination Scenario is false are not justified. Thus the Combination Scenario argument does not succumb to the objection that we have made to both the Recent Envatment argument and the Embodiment argument.

However, the Combination Scenario argument is vulnerable to the vat argument. The only difference between the Combination Scenario and Lifelong Envatment is that, according to the former, I was taken out of my vat and placed in a body before being re-envatted last year. This strange process happened in the blink of an eye and afforded me no opportunity to acquire the concept of a brain, so there is no more reason to think that I would have the concept of a brain in the Combination Scenario than there is to think that I would have the concept of a brain in Lifelong Envatment. It follows that if the Combination Scenario is true, then I cannot entertain the Combination Scenario. Once this claim is in place, we can run the vat argument to show that the Combination Scenario is false.

5. Radical scenarios and non-radical scenarios

At this point I take it that I have replied to the usual argument that the vat argument cannot be used as a defence against closure scepticism. That is, I have replied to the version of this argument that takes Recent Envatment and Embodiment (or slight variations thereof) as examples of non-radical scenarios. That this argument has turned out to be unsound should alone be enough to prompt interest in the vat argument as a defence against closure scepticism. In this section I will try to do more to show that the vat argument provides a defence against closure scepticism, if it works. I will do this by showing that there is good reason to think that no variation on the objection to the vat argument (as a response to closure scepticism) considered in the last section will succeed.

§3 and §4 effectively present a dilemma for the sceptic: either she makes use of Lifelong Envatment, in which case I can use the vat argument to show that Lifelong Envatment is false; or she makes use of Recent Envatment and Embodiment, and so enables me to

appeal to my empirical beliefs to show that these scenarios are false. This suggests that the sceptic may face a more general dilemma: either she uses a radical sceptical scenario (like Lifelong Envatment), in which case I can use the vat argument to show that I am not in that scenario; or, she uses a non-radical sceptical scenario (like Recent Envatment or Embodiment), in which case I can appeal to my empirical beliefs to show that I am not in that scenario. All sceptical scenarios that might otherwise have been of use in a closure argument for scepticism are either radical or non-radical, and so there is no sound closure argument for scepticism.

If the sceptic really does face the general dilemma, the vat argument is a very strong defence against closure scepticism indeed, and I think that there is good reason to think that the sceptic does in fact face it. The second horn is supported by the fact that so long as a sceptical scenario is non-radical there will be large swathes of my empirical beliefs that remain untouched. There is no principled reason these beliefs, in conjunction with my sensory evidence, will not indicate that the scenario is false, and it is not at all obvious that there is a non-radical scenario where this is not so. Even if the sceptic succeeded in describing such a scenario and employing it in a closure argument she would not have reached the strong sceptical conclusion that none (or hardly any) of my empirical beliefs are justified; empirical beliefs I have that are compatible with the non-radical scenario would be outside the scope of the closure argument. If the sceptic cannot reach this strong conclusion we are still making progress against scepticism. In order to mop up the remaining justification the sceptic must come up with another non-radical scenario that I am not justified in believing to be false that is incompatible with the empirical beliefs that the first sceptical scenario was compatible with. This second scenario would then be used in a second closure argument that mops up the justification left over from the first closure argument. The prospects for this strategy are dim, to say the least.

The sceptic would do better to attempt to grasp the first horn of the dilemma and argue that there is a radical sceptical scenario that the vat argument does not rule out. The sceptic may be encouraged here by the fact that the causal constraint on concept possession described in §2 is vague, so it is not obvious that it will rule out all radical sceptical scenarios.

However, this point cuts both ways: as stated, it is not clear which scenarios the causal constraint will rule out, but it is not clear which scenarios it will rule in either. If the sceptic wants to show that there is

a radical sceptical scenario that the causal constraint does not rule out she will have to sharpen up the constraint. This will not be easy, because the correct way of sharpening up the constraint depends upon the correct account of the conditions under which a subject has a particular concept. However, it is controversial what the correct account of concept possession is. If the sceptic must make claims about the correct account of concept possession before she can make her sceptical argument, the dialectical force of scepticism will be considerably weaker than that of the sceptical argument considered in §1, which relied only on a few simple and (*prima facie*, at least) plausible epistemic principles.

My point here can be illustrated by considering an example of how the sceptic might attempt to grasp the first horn of the dilemma. Button describes the following scenario:

The Vat Earth Scenario: Earth has a distant neighbour, Vat Earth. This is a planet whose only inhabitants are eternally envatted brains. There is no relevant causal link between Earth and Vat Earth, but it so happens that, for every brain on Earth, there is a brain on Vat Earth in exactly the same state, and vice versa.²²

Because there is no relevant causal link between Vat Earth and Earth, there is no more reason to think that the Vat Earth brains will have the concepts necessary to describe their scenario than I would in Lifelong Envatment.

However, Button asks us to alter the scenario, moving Vat Earth closer to Earth, and allowing Earthlings to visit Vat Earth. Earthlings look at the brains on Vat Earth. We are then to suppose that some Vat Earthlings and an equal number of visitors from Earth simultaneously think the following words: ‘if I were envatted and someone were looking at me right now thought the glass walls of my vat, I would want to be judged according to the language of [those outside the vat], in thinking “I am envatted”’.²³ It may seem more plausible that, if I were a Vat Earthling, I could possess the concept of a brain and the concept of a vat in the altered scenario than in Lifelong Envatment. If I could, then perhaps the sceptic has found a radical scenario that is not vulnerable to the vat argument.

²² The description of this scenario is taken with some very minor stylistic alterations from Button (2013, p. 156).

²³ I am here directly quoting a bit of speech that Button (2013, p. 157) puts into the mouth (as it were) of a Vat Earthling.

But could I have these concepts in the altered scenario? Button (2013, p. 157) reports that he ‘boggle[s] at the attempt to *start* answering that question’, (emphasis in original). He ultimately wants to claim that there is no clear cut answer to this question. This idea is interesting, but for present purposes we only need the weaker claim that it is not obvious that the answer to the question is yes. If a sceptic wants to employ the Vat Earth Scenario in a closure argument she will have to provide an *argument* that I would indeed have the relevant concepts in the altered scenario.

In order to provide such an argument she will have to make strong and controversial claims about concept possession. It is not clear that she will ever be able to produce convincing arguments for these claims. It is even less clear that these arguments could ever be convincing enough to make us take seriously the sceptical conclusion that they are meant, in conjunction with the rest of the sceptical argument, to imply. At any rate, the burden is now on the sceptic to provide such arguments before we start worrying that the sceptical conclusion is true. If I were a sceptic I would not want to be laden with such a burden.

If the sceptic wishes to avoid entangling herself in a debate about concept possession one natural strategy is to describe a radical sceptical scenario in the abstract, thereby avoiding making use of concepts, such as the concept of a brain, that I would not have were I to be in that scenario. This is the sort of strategy that Wright (1992, p. 91) pursues when he argues that the vat argument cannot be used to rule out the ‘bare’ possibility that I am in a sceptical scenario.²⁴ However, this strategy does not succeed.

Let us imagine that the sceptic gives an abstract definition of a sceptical scenario by saying only that it is a scenario that meets the criteria for a sceptical scenario that can be substituted for SS resulting in a plausible closure argument, as these conditions were described in §1. That is, the scenario is described as follows:

Abstract Scenario: My sensory evidence is the same as the sensory evidence that I currently have. However, all my empirical beliefs are false.

To be clear: this is *all* the description of this scenario that we are given.

²⁴ It should be noted that when Wright considers this strategy he is no longer considering a sceptical scenario of the sort that could be employed in a closure argument. Button (2013) also considers this strategy.

Consider the concepts required in order to entertain Abstract Scenario: the concept of belief, of sensory evidence, of a scenario, etc. The sceptic claims that the causal constraint would not prevent me from having these abstract concepts if I were in Abstract Scenario. For one thing, the causal constraint is most plausible when it is applied to concepts of medium-sized dry goods, such as brains; it may not apply to abstract concepts, such as the concept of belief, at all. For another, even if the causal constraint does apply to abstract concepts there seems to be no reason to think that I could not have had the kind of causal contact with, for example, beliefs required to have the concept of a belief if Abstract Scenario is true. Thus we shall admit, at least for the sake of argument, that there is no vat argument that I would not be able to entertain Abstract Scenario if that scenario obtained.

However, this retreat into abstraction will not save the sceptic. As Cross (2010) has argued, sceptical scenarios can be viewed as hypotheses that attempt to explain why I have the sensory evidence that I do. If the sceptical hypothesis is just as good an explanation as any non-sceptical hypothesis according to which my empirical beliefs are largely correct, then I am not justified in believing one of the non-sceptical hypotheses to be true. If however the sceptical hypothesis is not as good as one of the non-sceptical hypotheses, then I am justified in thinking it to be false. What makes for a good explanation is not entirely clear. However, it is clear that it is not a good explanation of why it is the case that p to say that p . This is not an explanation at all, it is simply a restatement of the fact that needs explaining. So, as Cross points out, whatever scenario the sceptic employs had better do more than simply restate that I have the sensory evidence that I do.

Now consider the statement that I have the sensory evidence that I do because I am in Abstract Scenario. Abstract Scenario is described only as a scenario in which I have the sensory evidence that I currently have, and yet my empirical beliefs are false. The first conjunct simply re-states that I have the sensory evidence that I am trying to explain, and so it cannot be an explanation of that evidence. The second conjunct adds that one possible explanation of why I have this sensory evidence – the explanation that states that most of my empirical beliefs are true – is unavailable. Stating that one possible explanation of my sensory experiences is not true is hardly sufficient to make Abstract Scenario an explanation of those experiences. So Abstract Scenario is not equal as an explanation of my sensory experiences to the non-sceptical explanation that is provided by my empirical beliefs about

the external world. Thus, my belief that Abstract Scenario is false is justified, and the sceptic cannot make use of Abstract Scenario in a closure argument.

An analogy may help to illustrate the point. Imagine that a scientist carefully hypothesises about some data, and is then confronted by someone who does not reject the data, but claims that this hypothesis is wrong. Suppose that the scientist asks why the hypothesis is wrong. If no answer is forthcoming, then it does not seem that any argument has been provided that the scientist's belief in her hypothesis is not justified. This is because no alternative hypothesis has been introduced that can compete with the scientist's hypothesis as an explanation of the data. However, things are different if the answer consists of a detailed hypothesis that explains the data and is logically incompatible with the scientist's own hypothesis. If this new hypothesis is at least as good an explanation of the data as the scientist's hypothesis, then we may conclude that the scientist is not justified in believing the new hypothesis to be false. We may go on to conclude that the scientist's belief in her own hypothesis is not justified.

When the sceptic claims that my belief that Abstract Scenario is false is not justified, this is analogous to the scientist's confrontation with someone who claims only that her hypothesis is wrong. In both cases, a good alternative explanation of the evidence has not been provided, and so the rival hypotheses can be dismissed. The sceptic may attempt to fill in the details of her story by talking, for example, about BIVs. By doing so she hopes to put herself on the same footing as someone who presents the scientist with a detailed and convincing alternative explanation of the data. However, once she retreats from abstraction the sceptic is back with the general dilemma described at the beginning of this section: either she talks about a radical sceptical scenario, in which case I can use the vat argument to rule it out; or she talks about a non-radical sceptical scenario, in which case I can appeal to my empirical beliefs to rule it out.

6. Conclusion

I have argued that, if it works, the vat argument is of use as a defence against closure scepticism. First I showed that the usual argument that the vat argument cannot be used as a defence against closure scepticism fails. Because closure arguments that make use of Recent Envatment and Embodiment do not threaten the justification of all my empirical beliefs at once, I am able to appeal to my empirical

beliefs to show that I am not in those scenarios. I then argued that no variation on the usual argument works either. The sceptic faces the general dilemma described at the beginning of the last section: either she uses a radical sceptical scenario, in which case I can use the vat argument to show that I am not in that scenario; or the sceptic uses a non-radical sceptical scenario, in which case I can appeal to my empirical beliefs to show that I am not in that scenario.

The conclusion of this paper is conditional: *if the vat argument works*, it can be used as a defence against closure scepticism. I have not argued at any length that the vat argument is sound. Interestingly, very few people have claimed that it is not.²⁵ Perhaps epistemologists have thought that it is not worth discussing whether or not the vat argument is sound because they think that the sceptic only needs to make use of non-radical sceptical scenarios. However, if I am right that the sceptic cannot make use of such scenarios then it would be well worth considering whether or not the vat argument works. If it does, it is a powerful weapon when directed against closure scepticism.²⁶

References

- Brueckner, Anthony. 2010. *Essays on Scepticism*. Oxford: Oxford University Press.
- 2006. 'Johnsen On Brains In Vats'. *Philosophical Studies* (129): pp. 435–40.
- 1999. 'Transcendental Arguments from Content Externalism'. In R. Stern (ed.), *Transcendental Arguments: Problems and Prospects*. Oxford: Clarendon Press.
- Button, Tim. 2013. *The Limits of Realism*. Oxford: Oxford University Press.
- Byrne, Alex. 2004. 'How Hard are the Sceptical Paradoxes?' *Nous* 38 (2): pp. 299–325.
- Christensen, David. 1993. 'Skeptical Problems, Semantical Solutions'. *Philosophy and Phenomenological Research* 53 (2): 301–21.
- Cross, Troy. 2010. 'Sceptical Success'. In *Oxford Studies in Epistemology*, pp. 35–62. Oxford: Oxford University Press.
- Dretske, Fred. 1992. 'Epistemic Operators'. *Journal of Philosophy* 67: pp. 1007–23.

²⁵ One exception is Johnsen (2003). I deal with Johnsen's objection in §2.

²⁶ I would like to thank Tim Button, Adrian Haddock, Bernhard Salow, Peter Sullivan, Crispin Wright, and two anonymous referees for *Mind* for comments and discussion.

- Farrell, Frank. 1986. 'Putnam and the Vat-People'. *Philosophia* 16 (2): pp. 147–60.
- Fogelin, Robert. 1994. *Pyrrhonian Reflections on Knowledge and Justification*. Oxford: Oxford University Press.
- Forbes, Graeme. 1995. 'Realism and Skepticism: Brains in a Vat Revisited'. *The Journal of Philosophy* 92 (4): pp. 205–22.
- Glymour, Clark. 1982. 'Conceptual Scheming or Confessions of a Metaphysical Realist'. *Synthese* 51 (2): pp. 169–80.
- Jackson, Frank. 1998. 'Reference and Description Revisited'. *Philosophical Perspectives* 12: pp. 201–18.
- Johnsen, Bredo. 2003. 'Of Brains in Vats, Whatever Brains in Vats May Be'. *Philosophical Studies* 112 (3): pp. 225–49.
- Lewis, David. 1984. 'Putnam's Paradox'. *Australasian Journal of Philosophy* 62 (3): pp. 221–36.
- Moore, Adrian. 1996. 'Solipsism and Subjectivity'. *European Journal of Philosophy* 4 (2): pp. 220–33.
- 2011. 'Vats, Sets and Tits'. In P. Sullivan and J. Smith (eds.), *Transcendental Philosophy and Naturalism*.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge Mass: Harvard University Press.
- Pryor, James. 2000. 'The Sceptic and the Dogmatist'. *Nous* 34 (4): pp. 517–49.
- Putnam, Hilary. 1973. 'Meaning and Reference'. *The Journal of Philosophy* 70 (19): pp. 699–711.
- 1981. *Reason Truth and History*. Cambridge: Cambridge University Press.
- 1975. The Meaning of 'Meaning'. In K. Gunderson (ed.), *Language, Mind, and Knowledge*.
- Smith, Peter. 1984. 'Could we be Brains in a Vat?' *Canadian Journal of Philosophy* 14 (1): pp. 115–23.
- Sosa, Ernest. 1997. 'Reflective Knowledge in the Best Epistemic Circles'. *The Journal of Philosophy* 94 (8): pp. 410–30.
- Tymoczko, Thomas. 1989. 'In Defense of Putnam's Brains'. *Philosophical Studies* 57 (3): pp. 281–97.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wright, Crispin. 1992. 'On Putnam's Proof that We are not Brains-in-a-Vat'. *Proceedings of the Aristotelian Society* 92 (1): pp. 67–94.
- 1991. 'Scepticism and Dreaming: Imploding the Demon'. *Mind* 100: pp. 87–116.