

BRAINS IN VATS? DON'T BOTHER!

PETER BAUMANN

pbaumann@swarthmore.edu

ABSTRACT

Contemporary discussions of epistemological skepticism – the view that we do not and cannot know anything about the world around us – focus very much on a certain kind of skeptical argument involving a skeptical scenario (a situation familiar from Descartes' *First Meditation*). According to the argument, knowing some ordinary proposition about the world (one we usually take ourselves to know) requires knowing we are not in some such skeptical scenario SK; however, since we cannot know that we are not in SK we also cannot know any ordinary proposition. One of the most prominent skeptical scenarios is the brain-in-the-vat-scenario: An evil scientist has operated on an unsuspecting subject, removed the subject's brain and put it in a vat where it is kept functioning and is connected to some computer which feeds the brain the illusion that everything is "normal". This paper looks at one aspect of this scenario after another – envatment, disembodiment, weird cognitive processes, lack of the right kind of epistemic standing, and systematic deception. The conclusion is that none of these aspects (in isolation or in combination) is of any relevance for a would-be skeptical argument; the brain-in-the-vat-scenario is irrelevant to and useless for skeptical purposes. Given that related scenarios (e.g., involving evil demons) share the defects of the brain-in-the-vat-scenario, the skeptic should not put any hopes on Cartesian topoi.

PROLOGUE

S, someone like us, doesn't know whether Goldbach's conjecture is true or false. S believes that it is true but doesn't know it. S does, however, know that 2 is not greater than 3. S's lack of knowledge whether Goldbach's conjecture is true or false poses no threat to and is simply irrelevant to the possibility of S's knowledge that 2 is not greater than 3. Otherwise S, or anyone like us, would have little if any mathematical knowledge. I will argue here that a (somewhat) parallel claim can be made for our lack of knowledge whether we're brains in vats (which we believe we aren't) and our knowledge that we have hands (or anything like that). The first poses no threat to and is simply irrelevant to the second.

I. INTRODUCTION

Contemporary discussions of epistemological skepticism, the view that we do not and cannot know anything about the world around us, focus very much on a certain type of skeptical argument. This type of argument uses the idea of a skeptical scenario: a situation such

that when one is in it one cannot tell that one is in it while when one is not in it one cannot tell that one is not in it. According to the argument, we do not and cannot know that we're not in such a skeptical scenario, even when we're not in one. But knowing this is required for knowing ordinary propositions about the world. Hence, we do not and cannot know any ordinary propositions about the world around us (see for many DeRose 1995: 1, *passim*).

One of the recently most discussed such skeptical scenarios is the brain-in-a-vat-scenario. Here is Putnam's classic passage (1981: 5–6; see also the precursor in Nozick 1974: 43 or, before that, Harrison 1966/1967):

“[I]magine that a human being (you can imagine this to be yourself) has been subjected to an operation by an evil scientist. The person's brain (your brain) has been removed from the body and placed in a vat of nutrients which keeps the brain alive. The nerve endings have been connected to a super-scientific computer which causes the person whose brain it is to have the illusion that everything is perfectly normal. There seem to be people, objects, the sky, etc.; but really all the person (you) is experiencing is the result of electronic impulses travelling from the computer to the nerve endings. The computer is so clever that if the person tries to raise his hand, the feedback from the computer will cause him to 'see' and 'feel' the hand being raised. Moreover, by varying the program, the evil scientist can cause the victim to 'experience' (or hallucinate) any situation or environment the evil scientist wishes. He can also obliterate the memory of the brain operation, so that the victim will seem to himself to have always been in this environment. It can even seem to the victim that he is sitting and reading these very words about the amusing but quite absurd supposition that there is an evil scientist who removes people's brains from their bodies and places them in a vat of nutrients which keep the brains alive. The nerve endings are supposed to be connected to a super-scientific computer which causes the person whose brain it is to have the illusion that ...”

The skeptical argument working with this particular skeptical scenario would then go like this: One does not and cannot know that one is not a brain in a vat (biv) suffering from certain systematic illusions. Since this kind of knowledge is required¹ for knowing any ordinary proposition about the world (like *There is another person in the room with me right now*), one does not and cannot know any ordinary proposition about the world.²

What work is the idea of a biv-scenario doing here? What about it is relevant to the problem of epistemological skepticism and why? Is it relevant at all to any interesting question about knowledge? The answer to the latter question will be negative. Here is how I will proceed. I will take one aspect of the original biv-scenario off after another – until it is so much thinned out that it is not a biv-scenario any more. I am not making any claims about the relative importance of the different aspects. I will refer to the different resulting variations on the original biv-scenario as different variations of “our scenario”. In a first step (Sections 2–5), I will argue that we should discard three salient but irrelevant aspects of the biv-scenario; I also explain (Section 5) what remains as the bare bones of the

1 This assumption is usually underwritten by a principle of epistemic closure (see the overview in Baumann 2011) or by an underdetermination principle (see for many: Brueckner 1994 or Cohen 1998); for this, see below (Sections 6 and 7).

2 It does not matter for the argument presented here whether we consider a skeptical strategy that offers one skeptical scenario for all ordinary propositions or rather a skeptical strategy that offers for each such proposition one skeptical scenario (though not necessarily the same for all).

scenario. In a second step (Sections 6–8) I will argue that we should discard what remains of the scenario (the possibility of being wrong and the possibility of being unreliable) as unconvincing skeptical challenges. I will also return (in Section 8) to the question whether the skeptical argument under consideration has any force. Finally (Sections 9–10), I will make a few remarks about why we might be impressed and misled by irrelevant scenarios; I will end with the recommendation that skeptics should not look at Cartesian scenarios for help. – Let us now take a closer look at the biv-scenario.

2. NO OPERATION NEEDED

First of all, we need not assume that brains in vats result from operations on normal human beings. This assumption creates unnecessary complications and red herrings. Suppose, for instance, that S's brain has been envatted. Is S, the human person, then identical with the brain that used to be located in his skull and is now located in a vat? This is a difficult question but it is not relevant to the skeptical problem. Or consider this question: Can't S (or S's brain or the brain in the vat, etc.) still know a lot about the world when the envatment has happened recently? This question may be (given additional substantial assumptions about mental content for a recently envatted brain) somewhat related to the problem of skepticism but recent envatment or even envatment in general (as we will see) is still not essential or relevant (and neither necessary nor sufficient) for a skeptical scenario involving bivs. We do not need more for our skeptical argument than the possibility that someone is not a normal human being but a biv. Come to think of it: Do we even need that for our scenario to work?³

3. DON'T BOTHER WITH MISSING BODIES

Second, there is certainly some massive disembodiment involved in envatment. Or, if we skip talk about the process of envatment, we should rather say that bivs are lacking a body. Since there is still a brain to be found in the vat and the lack of a body is thus not complete (how could it ever be?), we should be more precise and say that bivs are missing the rest of a body (e.g., a human body).

John Searle once quipped that we normal human beings are (or have) brains in vats, namely in skulls. Searle has a point insofar as it cannot be relevant here that the vat in the typical biv-scenario is not a normal human skull. Nothing relevant changes if we imagine a brain in a vat which happens to be a human skull without the rest of a body (though this scenario might make us feel somewhat uneasy and imagine heads being cut off or evil scientists purchasing human skulls from headhunters in order to use them (the skulls) as vats for brains).

Does it matter that the rest of a normal (e.g., human) body is missing? How could it? We could, for instance, give hands to the biv so that it could clap its hands (for instance when applauding a talk by an anti-skeptical philosopher). Sure, the added pair of hands

3 It should go without saying that whether the scientist in Putnam's scenario is evil or not (and, e.g., just indifferent) is completely beside the point and irrelevant (see also Putnam 1981: 6).

complicates things a bit because now the biv might be able to think a true thought when it thinks *I have hands*. Perhaps there could even be knowledge of hand ownership. But this complication is not really a new one in kind since even without hands the biv is able to think the true thought (and perhaps even know) that it has a brain (or, alternatively, is a brain). Whether it can, in addition, also think a few other true thoughts about other body parts does not make more than a very minor difference to the skeptical challenge.

Hence, if nothing relevant to our skeptical scenario changes if we give the biv a pair of hands, nothing relevant changes if we also give it a pair of feet. Nothing relevant changes if we imagine that the biv has a whole body (e.g., a human body). If you give a biv a hand, it will snatch the whole body (see also Roush 2010: 248–50 on this point).

The idea of massive disembodiment or of lack of a normal body is a red herring. It is spooky, especially when one is under the impression that one has a body. But the spookiness does not do any relevant work here (nor does the idea that the scientist taking care of the brain is evil). Rather, the spookiness might do some other, more damaging, kind of work: occlude the fact that many aspects of the spooky scenario are irrelevant to our skeptical scenario (even if some skeptical scenarios might also be spooky in a very special epistemic sense); the spookiness even helps create the illusion that the idea of a brain without a body is relevant to our skeptical scenario (necessary or sufficient for it). It isn't. S's knowledge that there is a tree in front of him is compatible with huge weight loss – even if the weight loss is due to having one's body reduced to a brain. It (that is, the knowledge and not just the fact known) is compatible with having no body apart from the brain (and with ignorance of the fact that one misses other body parts). If connected in the right way with the environment via the computer, a disembodied brain can receive information and gain knowledge about its environment (see also the end of Section 4 below). On the other hand, if the biv doesn't know that there is a scientist in the room, then the addition of things like hands and feet won't better its epistemic standing.

4. NO PLACE FOR WEIRD WIRES

Where does all this leave us then? If we strip the biv-scenario of at least some of its inessential and irrelevant aspects we get a scenario involving someone who might just be like a normal human being or a normal subject except that his brain states, events and processes and his mental states, events and processes (whatever the relation between the two types of states, events and processes) are being caused in a unusual way: for instance, by a computer used by another human being and connected to the subject's brain. It is completely irrelevant to the skeptical scenario whether there is some other agent, like an evil (or not so evil?) scientist or Cartesian demon, in the background running the process. These assumed agents are bogeymen (see Dennett 1984: Ch. 1 who also applies this general idea to certain conceptions of free will) whose main function here is to induce some epistemic panic in us and make us more susceptible to the idea that biv-scenarios really capture something important and relevant. All that remains here then is the idea of a subject whose process of belief acquisition is unusual in a certain way, including wires. Let us call such a subject a "wired subject" (wis). Every biv is a wis but not every wis is a biv. Since wis-scenarios capture everything in a biv-scenario that is relevant to epistemological skepticism, we can

thus, in the context of the discussion of epistemological skepticism, replace talk about a biv by talk about a wis.⁴

But is even the less phantastic possibility of being a wis relevant to the problem of epistemological skepticism? Let us remind ourselves here that some people think that unusual belief generation is incompatible with knowledge. Many would, for instance, think here of Keith Lehrer's famous Truetemp-example (Lehrer 1990: 163–64):

Suppose a person, whom we shall name Mr. Truetemp, undergoes brain surgery by an experimental surgeon who invents a small device which is both a very accurate thermometer and a computational device capable of generating thoughts. The device, call it a tempucomp, is implanted in Truetemp's head so that the very tip of the device, no larger than the head of pin, sits unnoticed on his scalp and acts as a sensor to transmit information about the temperature to the computational system in his brain. This device, in turn, sends a message to his brain causing him to think of the temperature recorded by the external sensor. Assume that the tempucomp is very reliable, and so his thoughts are correct temperature thoughts. All told, this is a reliable belief-forming process. Now imagine, finally, that he has no idea that the tempucomp has been inserted into his brain, is only slightly puzzled about why he thinks so obsessively about the temperature, but never checks a thermometer to determine whether these thoughts about the temperature are correct. He accepts them unreflectively, another effect of the tempucomp. Thus, he thinks and accepts that the temperature is 104 degrees. It is. Does he know that it is? Surely not.

One point which is interesting here is that, in contrast to much of the later discussion of Lehrer's Truetemp case (see for a more recent example Pritchard 2010), Lehrer himself does not use the weirdness or the unusual nature of the process of belief acquisition as such against the possibility of coming to know the temperature.⁵ And right he is! In the biv- or wis-scenario, too, the unusual nature of the process of acquiring mental states and beliefs in particular seems irrelevant to the skeptical problem and even potentially misleading. Why should it matter whether we are wired or how we are wired? Why should the unusual nature of some cognitive process as such matter for skepticism? Some extraordinary scientists might have unusual cognitive capacities but that as such does not entail at all that they don't or can't know what they come to believe on the basis of the exercise of those unusual cognitive capacities. The use of new technologies can enhance our cognitive abilities but why should one think that this enhancement destroys knowledge?

Similarly, wired subjects are weird subjects and perhaps also spooky. But this does not entail anything about the possibility of getting to know things about the world: Weird wired subjects can know things, too. Let us suppose that the biv or the wis has the good fortune to be connected to a computer used by a very knowledgeable, competent

4 A wis does have a cognitively relevant connection to an environment. The question of how to individuate this environment (as more proximal or more distal, as involving just a computer or the whole world), can be left open here because this won't affect my argument. We can also leave aside questions concerning the reference of the wis' thoughts (or sentences). These are the kinds of questions which occupied Putnam (1981) but here it does not matter whether we are dealing with a wis' thought about bits and bytes in some computer or about trees and coffee mugs. All that remains as potentially relevant to the skeptical argument at this point is not affected by whether one adheres to semantic externalism or not (on the relations between epistemic internalism or externalism and semantic internalism or externalism, see also Carter *et al.* 2014).

5 The situation is different if the subject has reason to think that something strange is going on with her cognition. However, we can leave this aspect aside here because it plays no role in biv- or wis-scenarios.

and benevolent scientist (or a benign Cartesian genius or some God) whose main interest is to have the subject come up with true beliefs. In that case, it is conceivable that the process of belief acquisition is extremely (though perhaps not perfectly) reliable. We can easily imagine many cases where the biv or the wis acquires a true and reliable belief, say about there being a tree next to the house (by the way, “reliable belief” is just short here for “belief resulting from a reliable process”). Why should we then not credit the subject with knowledge that there is a tree next to the house? Analogous points can be made about other theories of knowledge than reliabilism.

Hence, we can strip our skeptical scenario of the idea of a wis, too; neither wired nor weird stuff is required or sufficient for our skeptical scenario.⁶

5. WHAT REMAINS THEN?

Has our idea of a skeptical scenario disappeared into thin air then? Have we taken away both the blade and the handle from the skeptical knife?

Lehrer thinks that Truetemp in the example quoted above lacks knowledge about the temperature because he has no idea about the reliability of the process leading to his temperature beliefs and does not know that his beliefs are true. If this type of meta-cognition is required for knowledge, then the gates to skepticism might seem wide open: Who can tell that his perceptual abilities are reliable and deliver true beliefs about the environment? However, the idea that knowledge requires such meta-beliefs constitutes a substantial assumption of a particular epistemological theory.⁷ If skepticism is wedded to a particular epistemological theory, then it loses much of its initial appeal: the appeal of posing some very basic, radical and natural threat to the very idea of the possibility of knowledge, no matter how one thinks about knowledge in further detail (see Stroud 1984).⁸ One might rather want to rethink one's epistemological theory here than accept its (potential) skeptical consequences.⁹

Even though Lehrer did not bring his Truetemp example up in a discussion of skepticism, the two problems he mentions point to what remains as factors of potential relevance for epistemological skepticism in biv- and wis-scenarios: whether the subject's process of belief acquisition is reliable (or shorter: whether the subject's beliefs are reliable) and whether the subject's beliefs are true. All that is left of our skeptical scenario is the possibility that the subject's beliefs are not true (or false¹⁰) or not reliable.¹¹ Neither false nor

6 It would be interesting in its own right to give a comparative discussion of the skeptical potential of Lehrer's Truetemp case and the biv-case. Pursuing this would, however, lead too far away from the main goals of this paper.

7 But see Sections 6 and 7 for less “loaded” skeptical arguments using Lehrer's points.

8 A claim or threat is more or less natural; it is natural to the degree to which it does not depend on substantial philosophical assumptions. I don't have to say more about “naturalness” here; this would be part of the skeptic's burden.

9 I don't take this concern to constitute an anti-skeptical argument in its own right. Rather, this point adds to the case made here against biv-skepticism.

10 For the sake of simplicity, we can assume here that a belief is not true iff it is false (not false iff it is true). We can also assume for simplicity's sake that someone who knows that *p* is true (false) *ipso facto* knows that it is not false (not true).

11 Nothing depends on the idea that knowledge requires reliability. One can easily replace reliability by some other (presumed) condition of knowledge or type of good epistemic standing, and the same general argument will go through (see Section 7). For the sake of simplicity I will go with reliability here

unreliable true beliefs can constitute knowledge. Does any of this suffice to make our reduced skeptical scenario convincing and useful for a skeptical problem and challenge?

6. FALLIBILITY IS BESIDE THE POINT

Let us go back to the general form of the skeptical argument under discussion here:

- (1) We do not and cannot know that we're not in a skeptical scenario;
- (2) Knowing we're not in a skeptical scenario is required for knowing anything about the world;
- (C) Hence, we do not and cannot know anything about the world.

If we apply this general template to what remains of the biv- or the wis- or, more generally, our skeptical scenario we get two arguments, one relating to the possibility of a lack of truth, the other relating to the possibility of a lack of reliability. Let us first look at the argument relating to the establishment of truth:

- (T)
- (1t) We do not and cannot know of any belief about the world that it is not false (true);
- (2t) Knowing of such a belief that it is not false (true) is required for it to amount to knowledge;
- (Ct) Hence, we do not and cannot know anything about the world.¹²

If one holds the (independently questionable) view that knowing that the belief that *p* is not false (true) just is knowing that *p*, then the above argument is utterly trivial and question-begging in an uninteresting sense:

- (1t') We do not and cannot know for any *p* about the world that *p*;
- (2t') Knowing for any given *p* that *p* is required for knowing that *p*;
- (Ct') Hence, we do not and cannot know for any *p* about the world that *p*.

and treat it as a representative of all such conditions (and spare the reader going through tedious repetitions of the argument for other types of positive epistemic status).

12 In more detail: First we get from the above general template to a biv-argument (using "having hands" as an example):

- (1') We do not and cannot know that we're not a handless (etc.) biv;
- (2') Knowing that we're not a handless (etc.) biv is required for knowing that we have hands (etc.);
- (C') Hence, we do not and cannot know that we have hands (etc.).

In a second step, we can replace reference to the biv-scenario in general by the remaining first option: that our beliefs are false. If we choose a de-re version of this claim, we get (T). – The same holds, *mutatis mutandis*, for (R) below (see Section 7). – The difference between (1t) and (Ct) is one of levels. (1t) is of the second order and has the following form (with "b" ranging over beliefs, "◇" for the possibility operator and "K" for the knowledge operator: $\forall b \text{ not } K(b \text{ is not false})$) (or, alternatively and stronger: $\forall b \text{ not } \diamond K(b \text{ is not false})$). (Ct), in contrast, is a first order claim of the form: $\forall p \text{ not } Kp$; or, alternatively and stronger: $\forall p \text{ not } \diamond Kp$. The skeptical worry behind (T) is based on a second-order worry that one's beliefs might be false; the skeptic is trying to derive a first order skeptical conclusion here. – Thanks to a referee here.

But even if we put this extreme variant aside and revert to (1t) and (2t) it is very implausible that either of them is true. Why does knowledge that there is a tree next to the house require that one determines that this is true? One might lack the concept of truth¹³ but still know that there is a tree next to the house. To be sure, under specific circumstances (for instance when serious doubts have been raised whether it really “is true” that there is a tree next to the house) a sophisticated subject who possesses the concept of truth might discuss the question in terms of the truth of that proposition. But even such a sophisticated subject need not do this and could rather discuss this in terms of the first-order question whether *p* (rather than the higher order question whether *p* is true).

What if we assume that the subject has the concept of truth? There are two very common ways of defending (2t). The first one uses the principle of epistemic closure, a very basic version of which (sufficient for our purposes here) has it that (pertaining to all subjects and propositions)

(Closure) If S knows that *p* and if S knows that *p* entails *q*, then S knows that *q*.

Suppose I know there is a tree in front of me now. Suppose I also know that this entails that my belief that this is so is true. Then, according to (Closure) I also know that my belief is true. Assuming that a subject knows the entailment from *p* to *q*, (2t) stands defended. And it does not seem to presuppose much of a particular, substantial epistemological theory – though some find (Closure) problematic (see Dretske 1970 or Nozick 1981: Ch. 3). However, (2t) is restricted to subjects who know the relevant entailment and possess the relevant concepts (*truth*, for instance).

An alternative defense of (2t) uses a principle of underdetermination (again, pertaining to all subjects and propositions):

(Underdetermination) If S's epistemic position with respect to *p* (S's evidence, reliability, etc.) is not better than with respect to some incompatible proposition *q*, then S does not know *p*.

Suppose I believe there is a tree in front of me now. This is incompatible with the object being a cleverly designed fake tree. According to (Underdetermination), I don't know this object is a tree if my epistemic position with respect to the proposition that the object is a fake tree is not weaker than with respect to the proposition that it is a genuine tree. So, I do not know it's a tree if I do not know it's not a fake tree. This easily generalizes and, again, (2t) is vindicated. However, there are good reasons to see (Underdetermination) as a substantial epistemic claim rather than an uncontroversial epistemic principle. Do I really not know I'm talking to you if I don't know I'm not talking to your identical twin?¹⁴

So, (2t) can be defended but the defense comes with a price: either the making of a substantial and pretty controversial assumption (an underdetermination principle) or by using a somewhat restricted (to subjects having the relevant concepts and knowing the relevant entailments) epistemic principle (closure) which to some seems a bit less controversial.

13 A subject that has beliefs thereby also has some understanding of truth (even if very implicit or “behavioral” ones: it can distinguish between truths and falsehoods). However, this does not mean that the subject can also form de dicto beliefs the content of which contains the concept of belief.

14 Cases like this one have also been used against (Closure).

However, the main problem with (T) lies in (1T): Why should it be so difficult to determine that it is true that there is a tree next to the house (given that one has the concept of truth, and given that one make certain basic inferences)? Figuring out that there is a tree next to the house is in the same ballpark of epistemic difficulty (for subjects possessing the relevant notions and capable of making certain basic inferences) as figuring out that it is true that there is a tree next to the house. Why then should one think that the latter is impossible to figure out?¹⁵ The skeptic needs an argument for this claim and no such argument seems in sight. If one wants to argue on the basis of some particular epistemic theory that knowledge does indeed require that one establishes that one's belief is true, and also claim that one cannot establish that, then one's skepticism depends on some particular substantial epistemological theory which might be much less plausible than the denial of skepticism.

But couldn't it be the case that we only know very little about the world around us? Perhaps but it is not easy to see why one should characterize the claim that this is indeed so as a "skeptical" one. Also, one would need an argument for this claim of the paucity of knowledge. General fallibility does not help at all here.

(T) thus turns out to be very unconvincing and an utter failure as a skeptical argument. Fallibility as such is irrelevant to our or any serious skeptical scenario and argument. We can thus leave (T) aside and not worry or even think much about it.¹⁶

7. LAST RESORT: POTENTIAL UNRELIABILITY

This leaves us with one potential skeptical scenario only, one relating to the establishment of reliability. Couldn't our beliefs just be "accidentally" true? Couldn't they be unreliable (taking reliability as representative for a whole set of epistemic statuses: see fn. 11 above).¹⁷ Here is the corresponding skeptical argument:

(R)

(1R) We do not and cannot know of any belief about the world that it is reliable;

(2R) Knowing of such a belief that it is reliable is required for it to amount to knowledge;

(CR) Hence, we do not and cannot know anything about the world.

This argument also does not appear exciting or even very interesting if one looks at it in its naked form and stripped of all inessential wild stories about brains in vats, evil scientists,

15 That knowing that *p* is in the same ballpark of epistemic difficulty as knowing that *p* is true does by no means entail or support the very different claim that knowing that *p* is in the same ballpark as knowing that one is not wrong about *p* because one is a biv. I don't take the first claim to have much relevance at all for the possibility of knowing that one is not a biv. Thanks to a referee here.

16 Is fallibility necessary for our or any skeptical scenario? Not if, e.g., veridical hallucination or veridical dreaming by an infallible subject is possible and counts as a skeptical scenario. Is fallibility sufficient for our or any skeptical scenario? No, this would be an uninterestingly cheap form of skepticism. There is also no reason to assume that (T) in combination with the argument discussed in the next section, (R) would have any relevance for skepticism.

17 See also Winters (1981) who argues that the potential falsehood of our beliefs is a red herring in the discussion of skeptical scenarios and that it is rather the lack of knowledge that matters. See also, in a similar vein, Kraft (2013).

wired subjects, weird processes, and so on. One would at least need an argument for (2r) and for (1r). Why does knowledge that there is a tree over there require knowledge that my vision is reliable? Again, like in the case of (T) above, one might want to use a principle of closure or a principle of underdetermination to argue for (2r). The doubts about (Underdetermination) mentioned above apply here, too. However, there are additional doubts concerning the applicability of (Closure) as support for (2r). Does knowledge that my car still has some gasoline in its tank also allow me to come to know (by some basic inferences) that my gas gauge is working properly? Can I really not know that my car is not out of gas if I don't know that my gas gauge is working properly? This leads to topics currently discussed as issues of "epistemic bootstrapping" and "easy knowledge" (see, e.g. Cohen 2002). Whatever stance one takes on these issues, they will be part of some substantial and not uncontroversial epistemological theory.

Apart from all that: Why can we not establish about at least some of our beliefs about the world that they are reliable (e.g., by using other types of beliefs)? Even if arguments can be supplied here, they would, again, be based on a very specific view of knowledge (specific in the same way in which, e.g., Lehrer's internalism is specific). Again, skepticism would lose its appeal as a very basic, radical and natural threat to the very idea of the possibility of knowledge, – a threat which does not depend on prior acceptance of any particular substantial epistemological theory. Again, one might rather want to rethink one's epistemological theory and commitment to (1r) and (2r) here than accept the skeptical consequences.

So, (R) is also of little use as a skeptical argument. The possibility of unreliability might worry some of us at some time but it does not rescue our or any skeptical scenario.¹⁸ Nothing then remains.¹⁹ Upon closer inspection, the biv-scenario collapses and leaves nothing behind (apart from, perhaps, some hot air). As far as epistemological skepticism is concerned, we should simply forget about brains in vats and all similar and related scenarios.

8. BIV REDUX: IGNORING THE ARGUMENT?

Have I spoken too soon? Before we say bye-bye to the biv-argument, it might be worth taking a brief look back at it:

(BIV)

(1b) I do not and cannot know that I'm not a handless biv;

(2b) Knowing that I'm not a handless biv is required for knowing that I have hands;

(Cb) Hence, I do not and cannot know that I have hands.

Again, (2b) can be supported to some degree by using (Closure) or (Underdetermination). The main problem with (BIV), however, is with (1b). To be sure: I do not and cannot

18 Is possible unreliability necessary for our or any skeptical scenario? Even if it should be it is certainly not sufficient.

19 Could one argue that a set of individually irrelevant aspects are relevant if conjoined? One could. But one would have to offer an argument that shows that and how the conjunction of irrelevant conjuncts can be relevant. There does not seem to be any such argument available.

know that I'm not a biv (but cf. Jackson 2015 against this claim). However, I can and do know that I'm not a handless biv (but cf. the very brief remarks in Jackson 2015: 2802, 2817). How's that? Well, the latter proposition constitutes a conjunction: that it is not the case that I am both handless and a biv. Neither one of the conjuncts entails the other (I could be a handless but otherwise normal human being; I could be a handed biv). And I can know the falsehood of a conjunction as soon as I know the falsehood of one of the conjuncts (and can make some basic inferences) – no matter whether I am in principle ignorant about the truth or falsehood of the other conjunct (also see Atkins and Nance 2014 here). Assume for illustration (one can choose other examples if one so prefers) that we do not and cannot know whether Goldbach's conjecture is true or false. Assume further that we still believe that it is true. We certainly also want to claim knowledge that $2 + 2 = 4$. Hence, we certainly also want to claim the following: We know that it is not the case that both $2 + 2$ does not equal 4 and that Goldbach's conjecture is false. But if we can know the falsehood of this conjunction, then we can also know the falsehood of the proposition that we're handless bivs. There are no relevant structural or formal differences between these two cases.

But (one might want to retort), am I not begging the question here against biv-skepticism? Am I entitled to claim knowledge that I have hands, even in the face of the biv-scenario? Yes, I am. In this paper we stripped the biv-scenario down to its "bare bones". Since neither an operation nor lack of a body is relevant to skepticism, we turned the biv-scenario into a wis-scenario. Since the weird wires also don't matter and have no relevance for our topic, we ended up with nothing but (T) and (R). However, these arguments are, at best, shallow skeptical arguments (shallow in the sense in which repeated "why"-questions by kids can be shallow), not substantial, basic and serious skeptical arguments. Hence, our second conjunct in the proposition that I'm not a handless biv can be replaced without any loss for our discussion by some other conjunct expressing a proposition that we believe to be false but are also tempted to hold we don't know to be false: for instance, that Goldbach's conjecture is false. But how interesting, relevant or worrisome for the epistemic skeptic could the following proposition be: that it is not the case that both I am handless and Goldbach's conjecture is false? Given the parallel with the proposition that I'm not a handless biv, we can now say bye-bye to biv-scenarios as far as epistemic skepticism is concerned.²⁰

I have treated the biv-skeptic as someone who is offering a skeptical argument based on certain basic principles or claims (see, e.g., (BIV) above). But what if the skeptic replied to all of the above that they don't want to make any assumptions and only rely on some brute and plausible intuition that brains in vats cannot have knowledge about the world (and that we need to be able to rule out such skeptical possibilities)? This would not be a very strong objection to my argument. First, skepticism is not intuitively plausible. Second, in order not just to defend but even to motivate such a view, the skeptic needs and cannot avoid arguments based on certain basic claims. It is no surprise that,

20 So, I am not arguing, in a Moorean way, that one can infer the denial of some skeptical hypothesis from knowledge that one has hands. The discussion of Moorean strategies against skepticism is orthogonal to this paper.

e.g., Cartesian skepticism is always formulated as involving and being based on such an argument. Skepticism without arguments and assumptions lacks rational force.²¹

9. BUT CAN ALL THAT BE REALLY TRUE? YES, AND PLEASE DO RESIST RED HERRINGS AND WILD PHANTASIES!

To be sure, we might not be in a position to know that we are not bivs or wiss. But that is not very interesting as far as epistemological skepticism is concerned. It does not pose a threat to the possibility of knowledge about the world (see above). Perhaps we are also not in a position to know whether we are this kind of cognizer (using the fast and less carefully deliberative system 1 rather than the slow and more deliberative system 2 for tasks of type N; see, for instance the recent Kahneman 2011) or that kind of cognizer (using system 2 rather than system 1 for tasks of type N). That would certainly be interesting in many ways but still of no relevance to questions about epistemological skepticism. Similarly, the question whether we are brains in vats (or wired subjects) or not has little if any epistemological relevance (though it might be of great relevance for ethical and other questions). Perhaps I am a biv; perhaps I cannot come to know whether I'm one or not. Perhaps it rained here 10 million years ago; perhaps I cannot come to know whether it has or not. But so what? The former limitation is as irrelevant to the possibility of knowledge as the latter. The illusion of philosophical importance might well be due to the wild bits of the story about brains in vats or weird subjects which fire up the imagination and pump up certain misleading intuitions and red herrings.

Here is a glimpse of an error-theory. Brain-in-the-vat-scenarios contain epistemically irrelevant aspects (the spooky and wild stuff) and epistemically relevant aspects (fallibility, limited reliability) which are, however, not also relevant to epistemic skepticism. The irrelevant aspects are the ones beyond our ken, not the relevant ones. If we don't distinguish between the two clearly, we risk falling for the illusion that there is a basic threat to the possibility of knowledge. We first mistake the relevant bits for bits relevant to epistemic skepticism in particular, and then we misjudge that we cannot rule out these "relevant" bits because we feel we cannot rule out the irrelevant bits. Apart from all this, our strong practical (negative) attitude towards a possible envatted life (see Nozick 1974: 43) might trigger an epistemic panic and skeptical thoughts.

10. NO CARTESIAN HOME FOR THE SKEPTIC

One might wonder whether the same things can be said, *mutatis mutandis*, about more traditional skeptical scenarios of a similar make, like the evil-demon-scenario or the dreaming-scenario (both discussed, e.g., in Descartes' *First Meditation* (1907/1913)). It very much seems so but I won't go into this here. It is certainly not easy to see how evil-demon-scenarios or dreaming-scenarios could have relevance for epistemological skepticism when biv-scenarios don't.

²¹ Thanks to a referee here. – Williams (1996) goes even further when he argues that the skeptic is forced to make substantial and controversial assumptions, not just plausible ones.

So, where can the skeptic go then? What remains of epistemological skepticism? No reason to panic for the skeptic: There is still Agrippa's trilemma or the problem of induction – just to mention two options. And then there are, in addition, more recent forms of skepticism, like lottery-skepticism (see Harman 1973: 61). Ironically, Descartes' *First Meditation* offers no home at all for the skeptic. It has very successfully deceived many philosophers into thinking it contains food for skeptical thought. However, neither it nor what has followed in its train is of any use to the skeptic.²²

REFERENCES

- Atkins, P. and Nance, I. 2014. 'A Problem for the Closure Argument.' *International Journal for the Study of Skepticism*, 4: 36–49.
- Baumann, P. 2011. 'Epistemic Closure.' In S. Bernecker and D. Pritchard (eds), *The Routledge Companion to Epistemology*, pp. 597–608. London: Routledge.
- Brueckner, A. L. 1994. 'The Structure of the Skeptical Argument.' *Philosophy and Phenomenological Research*, 54: 827–35.
- Carter, J. A., Kallestrup, J., Orestis Palermos, S. and Pritchard, D. 2014. 'Varieties of Externalism.' *Philosophical Issues*, 24: 63–109.
- Cohen, S. 1998. 'Two Kinds of Skeptical Argument.' *Philosophy and Phenomenological Research*, 58: 143–59.
- . 2002. 'Basic Knowledge and the Problem of Easy Knowledge.' *Philosophy and Phenomenological Research*, 65: 309–29.
- Dennett, D. C. 1984. *Elbow Room. The Varieties of Free Will Worth Wanting*. Cambridge, MA: MIT Press.
- Descartes, R. 1907–1913. *Meditationes de prima philosophia, René Descartes, OEuvres de Descartes* (C. Adam and P. Tannery, eds), Paris: Cerf, 1907–1913/ Paris: Librairie philosophique J. Vrin 1964–1976 (reprint), vol. VII.
- DeRose, K. 1995. 'Solving the Skeptical Problem.' *Philosophical Review*, 104: 1–52.
- Dretske, F. I. 1970. 'Epistemic Operators.' *Journal of Philosophy*, 69: 1007–23.
- Harman, G. 1973. *Thought*. Princeton, NJ: Princeton University Press.
- Harrison, J. 1966/1967. 'A Philosopher's Nightmare or the Ghost not Laid.' *Proceedings of the Aristotelian Society*, 67: 179–88.
- Jackson, A. 2015. 'How You Know You are not a Brain in a Vat.' *Philosophical Studies*, 172: 2799–822.
- Kahneman, D. 2011. *Thinking, Fast and Slow*. New York, NY: Farrar, Straus and Giroux.
- Kraft, T. 2013. 'Sceptical Scenarios are not Error-Possibilities.' *Erkenntnis*, 78: 59–72.
- Lehrer, K. 1990. *Theory of Knowledge*. London: Routledge.
- Nozick, R. 1974. *Anarchy, State, and Utopia*. New York, NY: Basic Books.
- . 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Pritchard, D. 2010. 'Cognitive Ability and the Extended Cognition Thesis.' *Synthese*, 175: 133–51.
- Putnam, H. 1981. 'Brains in a Vat.' In *Reason, Truth and History*, 1–21. Cambridge: Cambridge University Press.
- Roush, S. 2010. 'Closure on Skepticism.' *Journal of Philosophy*, 107: 243–56.
- Stroud, B. 1984. *The Significance of Philosophical Scepticism*. Oxford: Clarendon.
- Williams, M. 1996. *Unnatural Doubts: Epistemological Realism and the Basis of Skepticism*. Princeton, NJ: Princeton University Press.
- Winters, B. 1981. 'Sceptical Counterpossibilities.' *Pacific Philosophical Quarterly*, 62: 30–8.

22 For comments and conversations I am grateful to Philip Atkins, Sven Bernecker, Stewart Cohen, Richard Fumerton, Suck Young Won, Xiaoxing Zhang, anonymous referees and audiences in Bled, Paris, Shanghai, Bogotá and Healdsburg.

PETER BAUMANN is Professor of Philosophy at Swarthmore College. He works mainly on topics in epistemology (e.g., contextualism about “knowledge”) but he also has interests in other areas of philosophy such as, for example, human decision making and practical reason.
