

Penultimate draft.

Final version in: *Abstracta* 4:1, 2008; pp. 5-26.

Skepticism, Externalism and Inference to the Best Explanation

Jochen Briesen
Humboldt-University Berlin

Abstract

This paper focuses on a combination of the antiskeptical strategies offered by semantic externalism and the inference to the best explanation. I argue that the most difficult problems of the two strategies can be solved, if the strategies are combined: The strategy offered by semantic externalism is successful against standard skeptical brain-in-a-vat arguments. But the strategy is ineffective, if the skeptical argument is referring to the recent-envatment scenario. However, by focusing on the scenario of recent envatment the most difficult problems of the antiskeptical strategy posed by the inference to the best explanation can be solved. The most difficult problems with this strategy are: (1) Why is an explanation of our experience offered by the skeptical hypothesis more complex than our standard explanation? (2) Why is the more complex explanation less likely to be true? By focussing on the recent envatment hypothesis both questions can be answered satisfactorily. Therefore, the combination of semantic externalism and the inference to the best explanation yields a powerful antiskeptical argument.

1. Introduction

Philosophical skepticism casts doubt on the possibility of human knowledge and/or justification. An interesting and well-known response to this doubt is offered by semantic externalism. But even if we assume the validity of semantic externalism as a theory, is it really able to offer a successful strategy against skepticism? Or is externalism at least able to clear the way for alternative solutions to skeptical problems? Before these questions are addressed some preliminary remarks on the skeptical argument are provided in section 2.

Section 3 concerns the antiskeptical strategy outlined by externalism. The strategy traces back to Hilary Putnam and has been recently defended and improved by

Olaf Müller.¹ The conclusion of section 3 is this: If semantic externalism is right, then it forces the skeptic to reformulate her argument. Against the reformulated argument, however, externalism is ineffective.

Finally, section 4 investigates whether the reformulated skeptical argument forced by semantic externalism can be refuted. Are there any strategies that exploit the antiskeptical thrust of semantic externalism? I will argue that the most promising candidate for such a freeloading strategy is based on the inference to the best explanation: By focussing on the reformulated argument, the most difficult problems with the antiskeptical strategy posed by the inference to the best explanation can be solved.

2. Skepticism

There are a number of skeptical arguments that vary in structure and scope. This paper focusses on a specific kind of skepticism that is directed against our knowledge of the external world. Before turning to the argument, I will clarify its most essential component – the skeptical hypothesis.

2.1 The Skeptical Hypothesis

The so called brain-in-a-vat-scenario is probably the most popular skeptical hypothesis in the recent literature:

SH: Assume the brain of person X is detached from its skull by hostile aliens right at the beginning of X's existence. The brain is kept in a vat filled with a liquid of nutrients to keep it alive. The world as we think we know it has never existed and the vat is placed on a deserted planet. Moreover the brain of X is stimulated by electrodes and a giant computer to simulate X's experience. This simulation is so perfect that X lacks introspective criteria to recognise whether his or her experience is based on a simulation or not. Suppose we are all in the same situation as X.

Some of the most important features of SH are:

1. At least *prima facie* we cannot rule out SH by means of empirical evidence because all the evidence accessible to us could be based on the exact same computer simulation which the evidence tries to rule out. This feature distinguishes

¹ See Müller (2003).

SH from other more limited skeptical hypotheses which obviously could be rejected by further empirical investigation.²

2. SH is incompatible with (almost) all our beliefs about the external world (hereafter, „mundane beliefs“). This property distinguishes SH from the dream-hypothesis, which is consistent with the truth of our mundane beliefs. Even if our beliefs were formed on the basis of a dream experience, they could be true – one can dream what is actually taking place.
3. However, SH is not incompatible with all but only most of our beliefs with respect to the world. This feature distinguishes SH from Descartes’ demon-hypothesis, which states that we could be systematically deceived by an evil demon. Descartes formulates his hypothesis such that it is incompatible with all our mundane beliefs. Contrary, SH is consistent with some of these beliefs – for example: *My brain contains countless nerve pathways*. To some extent SH is therefore less radical than Descartes’ scenario. But this diminished radicality of SH is compensated for by its increased intelligibility. The demon-scenario refers to purely mental entities and their abilities – a scenario which is hard to imagine and impossible to describe in further detail.
4. SH is not only metaphysically but also scientifically possible. This is another feature which distinguishes SH from the demon-hypothesis. The demon-hypothesis assumes the existence of mental entities without any material bearer and is therefore hardly in accordance with our commonly accepted scientific beliefs. However, the skeptical scenario is just one component of the skeptical argument – it is by no means solely capable of undermining our commonly accepted scientific beliefs. This is why the skeptic will strengthen her position if she uses SH instead of the demon-hypothesis in her argument. She will invite more thinkers to follow her reasoning.

2.2 The Skeptical Argument

How does the mere possibility of SH undermine our knowledge-claims? The following argument answers this question:

² An example of this kind of limited skeptical hypothesis is Fred Dretske’s disguised mule scenario. See Dretske (1970: 1016).

- (i) If S knows that p, then S is in a position to know that \sim SH.
- (ii) S is not in a position to know that \sim SH.
- (iii) Hence, S does not know that p.

The skeptical argument is obviously valid, but whether it proves the conclusion (iii) depends on the plausibility of the premises (i) and (ii).

Premise (i)

Premise (i) depends on the acceptance of the closure principle of knowledge:

If S knows that p and S knows that p entails q then S is in position to know that q.

I cannot discuss this principle here in detail, so in the following I will presuppose its correctness and will accept the justification of (i) which rests on the principle.³ But it should be pointed out that because (i) is based on the closure principle, the skeptical argument proves only that propositions incompatible with SH cannot be known.

Premise (ii)

A way to justify premise (ii) that is *prima facie* appealing is to show that S is not in a position to know that \sim SH, either a posteriori or a priori.

Why is S not in a position to know a posteriori that \sim SH? A posteriori knowledge evidentially rests on empirical evidence. But in respect to \sim SH this seems impossible. S cannot rule out that she is a brain-in-a-vat whose experience is simulated by citing empirical evidence, since all the cited evidence could be part of the exact same simulation which it tries to rule out. Thus we might say: \sim SH cannot be justified by empirical evidence, so S is not in a position to know that \sim SH a posteriori.⁴

Why is S not in a position to know a priori that \sim SH? A priori knowledge can be justified independently of empirical evidence. But a priori knowledge of \sim SH also seems impossible. The proposition \sim SH concerns the position of an epistemic subject in the world. And to justify such a proposition you'll need empirical evidence – you have

³ John Hawthorne (2003) argues that only a reformulated and somehow weaker version of the closure principle can be defended successfully. If this was true, the skeptical premises (i) and (ii) would have to be adjusted to this reformulation.

⁴ Even though this line of reasoning is problematic, I will accept it in this paper. Prominent ways to question it would be: (1) The evidence externalism of Timothy Williamson (2000) and other forms of epistemic externalism. (2) The Dogmatism of James Pryor (2000). The influential antiskeptical strategies provided by these theories are not in the focus of this paper. I will concentrate on the strategy offered by semantic externalism and its potential freeloaders.

to take a look at the world to locate your position in it. Thus we might say: \sim SH cannot be justified by non-empirical evidence, so S is not in a position to know a priori that \sim SH.

The consideration that makes (ii) plausible can be summed up thus: (a) S is not in a position to know a posteriori that \sim SH. (b) S is not in a position to know a priori that \sim SH. (c) All cases of knowledge which are not a posteriori are a priori. Hence, (d) S is not in a position to know that \sim SH.⁵

Conclusion (iii)

The skeptical argument is valid and its premises are (at least *prima facie*) plausible. So it seems that we have to accept the conclusion that we know (almost) nothing about the external world. But there are strong reasons to resist this conclusion. First of all, it is absurd to deny this kind of knowledge generally. I do know that I am sitting in this chair right now, that my mother's name is »Ursula«, etc.⁶ Second, it would be very strange if the concept »knowledge«, which functions quite well in ordinary speech acts, was wrongly applied most of the time. In order to explain this, one would have to assume that we generally use the term »knowledge« in a way that does not correspond to its proper meaning. But the meaning of a word is at least partly determined by its daily use.

In the light of these considerations, the skeptical argument presents us with a paradox. A paradox is a set of plausible but inconsistent propositions and the lines (i), (ii) and the negation of (iii) constitute such a paradox. The argument of knowledge skepticism constitutes a paradox because it draws a very implausible conclusion via plausible premises. And the argument poses a threat because it proves via plausible premises that at least some of our most plausible candidates for knowledge cannot be cases of knowledge at all.

3. The Antiskeptical Strategy of Semantic Externalism

The antiskeptical strategy provided by externalism is best seen as an attempt to block the skeptical argument by refuting premise (ii) „S is not in position to know that \sim SH“. Section 2.1 justifies this premise by explaining why S is neither in a position to know a

⁵ See Weatherson (2007).

⁶ In this context David Lewis speaks of „Moorean Facts“: „It is one of those things we know better than we know the premises of any philosophical argument to the contrary.“ (1996: 418).

priori nor a posteriori that \sim SH. By contrast, some externalists argue that we are in fact in a position to know a priori that \sim SH. They think that certain semantic considerations, which can be made from the philosophical armchair, can prove that we cannot be the brain-in-a-vat described in SH: Even if the skeptical premise (ii) seems very plausible at first sight, the justification of (ii) is false – complex semantic reflections can provide a priori reasons that support \sim SH.

3.1 A Tentative Version of the Antiskeptical Argument

A variety of ways have been proposed to derive this antiskeptical conclusion from a semantic externalist standpoint.⁷ A simple argument can illustrate the core theses of semantic externalism as well as important problems with its antiskeptical strategy:

- (a) The term »tiger« used by a brain-in-a-vat is not referring to tigers.
- (b) My term »tiger« is referring to tigers.
- (c) Hence, I am not a brain-in-a-vat.

This argument is valid and it could be reproduced by any epistemic subject. Therefore we might infer from line (c) that S is in a position to know that \sim SH. In other words: If the argument is sound, it proves that the skeptical premise (ii) is wrong and therefore blocks the skeptical argument. But is this a satisfactory way to reject skepticism? The answer to this question depends on the plausibility of the premises (a) and (b). The following comment on (a) introduces the main features of semantic externalism. The comment on (b) provides some considerations regarding the a priori character of the antiskeptical argument and addresses its most difficult problem.

Premise (a)

Premise (a) „The term »tiger« used by a brain-in-a-vat is not referring to tigers“ is based on key theses of semantic externalism. By raising the famous twin-earth thought experiment Hilary Putnam concludes:

- (1) The reference of a natural kind term is not determined by the internal states of a speaker.

⁷ I restrict myself to the antiskeptical strategy which traces back to the semantic considerations of Hilary Putnam (1975, 1981). Concerning this strategy, different arguments for the antiskeptical conclusion are proposed: See Brueckner (2004); Wright (1991); and others. Donald Davidson's externalist reflections also seem to have antiskeptical consequences, which cannot be discussed in this paper.

- (2) To refer successfully to the natural kind X, the causal history of the speaker must contain appropriate causal relationships to samples of X.

The theses (1) and (2) explain why premise (a) is correct: Regarding its intrinsic features the brain-in-a-vat might be equivalent to a normal brain in that they both share the same neuronal and narrow psychological processes. But this does not mean that both of them are capable of referring to the same things. The precondition for referring successfully to natural kinds is an appropriate causal connection to them. But the brain-in-a-vat lacks this appropriate causal contact to natural kinds and to the external world in general. Therefore the brain-in-a-vat's term »tiger« does not refer to tigers. But what does its term »tiger« refer to? Presumably it refers to the computer program's binary configurations that are causally responsible for the tigerish experience – bit-tigers, so to speak. So the consequences of accepting externalism are twofold. First, if we accept externalism we have to agree on premise (a). Second, externalism reveals that most of the beliefs entertained by a brain-in-a-vat could be true. The belief „There is a (bit-) tiger in front of me“ for example is true if the corresponding bit configuration occurs.⁸

In the following I will take for granted the validity of the justification of (a), which is based on the key theses of externalism, without discussing these theses in detail. It lies outside the scope of the present paper to investigate whether externalism is right. Our present concern is the question of whether externalism, given its correctness, provides us with a solution to the skeptical paradox.

Premise (b)

In our dialectical situation premise (b) „My term »tiger« is referring to tigers“ seems more problematic. The following objection might be raised: Premise (b) presupposes the existence of tigers. In order for my term »tiger« to refer to tigers, tigers have to exist, since otherwise there would not be any referent for my term. So in order to justify that my term »tiger« in fact refers to tigers, I have to assume the existence of tigers. In relation to the antiskeptical argument this assumption is problematic on two grounds. First, the argument falls into epistemic circularity, because the justification of one of its

⁸ This consequence already has antiskeptical potential: We are frightened by the skeptical hypothesis, because we assume that the majority of our beliefs would be false if the hypothesis was realised. But externalism proves that even if we were brains-in-a-vat, our beliefs could be true. This might be quite reassuring, but it cannot block the skeptical argument.

premises presupposes the correctness of its conclusion: To justify that my term »tiger« refers to tigers, I have to assume that I am not a brain-in-a-vat in a tigerless universe. And second, the justification of (b) must be question-begging, because it depends on world-knowledge which we are not allowed to hold in our current dialectical situation.

3.2 An Improved Version of the Antiskeptical Argument

Olaf Müller suggests an improvement of the argument based on two modifications: A disjunctive premise (o) is added and all instances of »tiger« are substituted by »brain«.⁹

(o) Either brains exist or brains do not exist.

Case 1: Brains exist.

(a)* The term »brain« used by a brain-in-a-vat is not referring to brains.

(b)* My term »brain« refers to brains.

(c)* Hence, I am not a brain-in-a-vat.

Case 2: Brains do not exist.

(d)* Hence, I am not a brain-in-a-vat.

Line (o) can be known a priori, because it is based on the Law of the Excluded Middle. Line (a)* is backed by the same externalist assumptions which justified line (a), as long as one accepts that externalist theses can be applied to the term »brain« as well as to the term »tiger«.¹⁰ And on the condition that there are brains (case 1), line (b)* rests completely on the rules of disquotation.¹¹ From (a)* and (b)* we conclude line (c)* which states that I am not a brain-in-a-vat. And line (d)*, which states the same, follows solely from the description of case 2.

This improved argument is neither question-begging nor epistemically circular. It proves convincingly that the skeptical scenario must be wrong and thereby refutes the skeptical premise (ii) „S is not in a position to know that \sim SH“. The skeptical argument is blocked.

⁹ See Müller (2003: 167-183). Notice that Müller formulates the improvement a little differently. He argues that even if premise (b) was analyzed not to presuppose the existence of tigers, we would need an improvement of the antiskeptical argument along the indicated lines.

¹⁰ Everybody who denies this on the basis that »brain« is not a natural kind term, should replace »brain« by »human brain«, which clearly is such a term. The given argument would not be compromised by that replacement.

¹¹ The following formulation of premise (b)* can illustrate that the premise is independent of world knowledge: „My term »brain« refers to brains (whatever that is).“

An Objection

The given reasoning invites the following objection: A brain-in-a-vat (or a community of brains-in-a-vat) could reason in precisely this manner and thereby convince itself that it is not a brain-in-a-vat. But in this case the conclusion would be false. After all, the conclusion had been drawn by a brain which is actually in a vat. So the proposed antiskeptical argument must be defective.

If you accept the core theses of externalism, this objection can be easily dismissed. According to externalism, the language of brains-in-a-vat would be different from ours. The brains-in-a-vat are not capable of referring to the things we do. Therefore, if the envatted brain followed the proof, we would have to translate our antiskeptical proof into the language of the brains-in-a-vat. Let us call this language Bit-English. And in Bit-English the conclusion of the antiskeptical argument would be: I am not a bit-brain-in-a-bit-tank. And this conclusion would be true even if it was drawn by a brain which is actually envatted. Thus, the objection is dismissed and the antiskeptical proof is rehabilitated: The proposed antiskeptical argument, which is valid for every subject regardless of its language, proves that a given subject cannot itself be what it refers to as a »brain-in-a-vat«. This is enough to block the skeptical line of reasoning.

3.3 The Skeptic's Resort

The improved antiskeptical argument is not obviously defective and it successfully blocks the skeptical line of reasoning. To avoid the antiskeptical conclusion the skeptic has to reformulate her argument to immunise it against externalist attack. The skeptic has to modify the skeptical scenario such that it assures that brains-in-a-vat are capable of referring to the same things we do. And she has to adjust the skeptical argument to that modified hypothesis. A hypothesis SH2 along these lines is easily formulated.

SH2: This hypothesis is identical to the former hypothesis SH, but with two important exceptions: The brain of X was envatted only five minutes ago and the earth was destroyed afterwards.

Why is a recently envatted brain, even in the light of externalism, able to refer to the same things we refer to? Before the envatment the brain could have held all the necessary causal relations to brains, tigers, trees etc. Therefore the belief „There is a brain (tiger) in front of me“ entertained by a recently envatted brain, would be false.

»Brain« (»tiger«) would refer to brains (tigers), and in the vicinity of the recently envatted brain there are no brains (tigers) but only bit-brains (bit-tigers).¹²

Therefore, to immunise the skeptical argument against the externalist attack, the original hypothesis has to be substituted by SH2.

- (i)* If S knows that p, then S is in a position to know that \sim SH2.
- (ii)* S is not in a position to know that \sim SH2.
- (iii)* Hence, S does not know that p.

Both premises (i)* and (ii)* are (at least *prima facie*) justified in the same way that the original premises (i) and (ii) were justified. But contrary to the justification of (ii), the justification of (ii)* cannot be rejected by means of externalist considerations. So even if we accept externalism, the modified argument proves that we know almost nothing about the world we live in.

An Important Difference

There is a difference between the original and the modified skeptical argument that should be clearly stated. The scope of the modified argument is smaller. Like premise (i), premise (i)* is justified by the closure principle. But this justification presupposes that p entails \sim SH2. And there is a whole range of instances for p which entail \sim SH but do not entail \sim SH2. For all these instances, the original argument can prove their unknowability but the modified argument cannot. Take a look at the following example:

z: Last week the tiger at the Zoo Berlin was petted by its keeper.

The proposition z entails \sim SH but does not entail \sim SH2. Even if S was a recently envatted brain, z could still be true. The skeptic cannot prove via her modified argument that z is unknowable because z cannot be inserted in the modified argument – the premise (i)*_z „If S knows that z, then S is in position to know that \sim SH2“ is clearly false. There is a general lesson to be learned from this example: The modified skeptical

¹² Regarding its semantic features the position of a recently envatted brain is equivalent to the position of a person S, who has been placed on Putnam's twin-earth. Imagine twin-earth actually exists and that, without being aware of it, S was brought there (see Boghossian 1989 for this kind of thought experiments). Most externalists agree that if S were to remain on twin-earth long enough, S would eventually acquire the concepts appropriate to that situation. But nobody assumes that S would acquire twin-earthian concepts immediately. Similarly I agree that the envatted brain would eventually acquire bit-concepts, but this would not take place in the first few minutes.

argument cannot prove that we do not have knowledge of the past. Could this restriction be used against the skeptic? This question will be answered in the next section.

4. Exploiting the Antiskeptical Thrust of Externalism

Externalism forces a reformulation of the skeptical argument. To immunise the skeptical reasoning against externalistic attack, the skeptic has to substitute the original hypothesis SH by the scenario of recent envatment SH2. But is the modified argument weakened in a sense that we have thus far overlooked? Is the recent envatment hypothesis SH2 easier to rule out than SH? Three attacks against the modified skeptical argument are possible and all of them will be discussed in the remainder of the paper.

4.1 Modification and Skeptical Threat

The first attack against the modified argument is:

The substitution of SH by SH2 involves a restriction of the modified skeptical argument. In fact, the scope of the modified argument is so restricted that the argument no longer poses a real threat. Therefore, the modified argument requires no substantial refutation.

This suggestion is not very convincing. The scope of a skeptical argument is not very closely connected with its property of being a real philosophical threat. Skeptical arguments pose a threat because they prove via plausible premises that at least some of our most plausible candidates for knowledge cannot be cases of knowledge at all. And the modified skeptical argument is obviously in a position to do this. It cannot cast doubt on our knowledge of the past, but it still concludes that we cannot know anything about our present surroundings. And the beliefs referring to middle-sized objects in our present environment are surely some of the most plausible candidates for world-knowledge we can think of. So the modified skeptical argument remains highly problematic.

4.2 Recollection and Exclusion

The second attempt to refute the modified skeptical argument is:

The argument cannot cast doubt on our knowledge of the past. So this knowledge of the past can be used to rule out the recent envatment hypothesis SH2. Therefore premise (ii)* „S is not in a position to know that \sim SH2“ is wrong.

How is knowledge of the past supposed to rule out SH2? The suggestion to be discussed in this section answers this question along the following lines: Some of my beliefs formed on the basis of my past experience count as empirical reasons which justify my belief that I am not a recently envatted brain a posteriori. Take for example the belief that r , which I formed on the basis of a lively conversation with a couple of honest and highly qualified physicians yesterday.

r : Yesterday, the point where brains can be kept alive in vats was still more than 50 years away.

The belief that r was not formed on the basis of a simulated experience, because according to SH2, I was only envatted 5 minutes ago. Therefore the belief is a valid a posteriori reason for \sim SH2 – after all, the belief is based on empirical evidence which cannot be part of the exact same simulation the belief tries to rule out. Additionally, the envisaged a posteriori justification of \sim SH2 is not question-begging against the skeptic because the skeptical argument cannot prove that r cannot be known. Therefore, we can rule out SH2 via that empirical belief without presupposing anything the skeptic casts doubt on. So the skeptical premise (ii)* „S is not in a position to know that \sim SH2“ can be refuted a posteriori and the skeptic is caught in a dilemma. Either she is using SH in her argument and can be refuted a priori but not a posteriori. Or she uses SH2 and can be refuted a posteriori but not a priori. The antiskeptical mission is accomplished.

Unfortunately, the given considerations are defective. It is true that the suggested strategy is not question-begging. But it is not true that the aforementioned a posteriori reason cannot be based on the simulation SH2 refers to – it could be true that we have that reason because we are in fact in the simulation: Our recollections could be simulated and our beliefs referring to the past could therefore be formed on the basis of the simulation. Therefore, these beliefs cannot refute SH2, because we have already accepted that we cannot dismiss a simulation hypothesis by citing evidence which could be based on the exact same simulation the hypothesis assumes.

Is there any objection to the view that our recollection could be simulated by means of the simulation SH2 refers to? I do not think so. In our dialectical context, there are two ways to make the simulation of recollection plausible.

1. We have already accepted that neurology and simulation technology are capable of simulating all present experiences. Under these conditions it is hardly conceivable that the same technology is not also able to overwrite or replace existing memories.
2. In some respects a dream is comparable to SH2. Suppose I started dreaming five minutes ago, I can certainly dream that there is an apple lying in front of me without an apple actually present. Similarly, I can dream that I had an informative conversation with scientists even though this conversation never actually took place. Although I started dreaming only recently, my recollection could be based entirely on the dream experience. If a dream can overwrite recollections, a computer simulation can probably do so as well.

Although the modified skeptical argument cannot prove that our beliefs based on recollection are no cases of knowledge, the modified hypothesis SH2 cannot be refuted by means of that recollection.

4.3 Inference to the Best Explanation

The last and most promising attempt to refute the reformulated skeptical argument is:

Contrary to the standard explanation of our perceptual experience, the skeptic's explanation provided by SH2 must assume a double-mechanism: The experience before the envatment must be explained differently from the experience after the envatment. This double-mechanism makes the latter explanation clearly more complex than the standard one. Therefore an almost forgotten antiskeptical strategy can be reanimated: The inference to the best explanation.

The general antiskeptical idea behind this suggestion can be characterised in the following manner: There are different criteria for a good explanation of a phenomenon. A commonly accepted criterion for a good explanation is simplicity. Other things being equal, if the explanation E1 is simpler than the explanation E2, then one is justified in preferring E1 to E2. The following hypothesis W is a good starting point for an explanation of our perceptual experience.

W: We live (and always have lived) in an external world with all its objects. These objects causally interact with our senses. This causal relation with objects is responsible for our perceptual experience.

In contrast to W, a skeptical hypothesis provides a totally different explanation of our experience. Even if we grant the skeptic that her hypothesis has the same explanatory strength as W, her explanation is more complex. Therefore, we have reason to dismiss the skeptical hypothesis.

Traditionally, the most obvious problems with this antiskeptical strategy are:

- (1) Why is the explanation provided by a skeptical hypothesis more complex than the standard explanation provided by W?
- (2) Why is it epistemically justified to prefer the simpler explanation?

Both questions will be addressed separately.

Question (1): Why is the explanation provided by the skeptical hypothesis more complex?

Generally it is hard to clarify why an explanation of our experience provided by a skeptical hypothesis should be more complex than the standard explanation. Thus far no clarification of this matter has been widely accepted.¹³

David Chalmers argues that with respect to the modified hypothesis SH2 on the other hand, there is a clear sense in which the skeptical explanation is more complex.¹⁴ The explanation of our experience provided by SH2 involves a double-mechanism. Through the assumption that we have only recently been envatted, the skeptic offers different explanations of our experiences before and after the envatment. Therefore the skeptic's explanation is more complex than the standard explanation, which is not committed to this kind of double-mechanism.

An Objection

Is it really true that the standard explanation of our experience based on W is not committed to a double-mechanism? To follow the outlined antiskeptical idea, we have to understand the term »perceptual experience« in a wide sense, such that it is not true by definition that every perceptual experience has an object in the external world. Only if we understand perceptual experience in this wide sense, is it right to say that the brain-in-a-vat has experiences at all. Regarding this sense of »perceptual experience«

¹³ Nonetheless there are some interesting suggestions in this direction: see Vogel (1990: 658-666).

¹⁴ See Chalmers (2005: 35-36).

hallucinatory experiences count as perceptual experience as well. But hallucinations – in contrast to normal experiences – cannot be explained by means of our causal contact with objects and therefore the standard explanation of our experience also involves a double-mechanism. Thus, the outlined antiskeptical strategy is ineffective, because the standard explanation is committed just as much to a double-mechanism as the skeptical explanation is.

This objection can be handled. Even if the given considerations are right and our standard explanation is really committed to a double-mechanism, all that follows is that the explanation provided by SH2 is minimally committed to a triple-mechanism. The skeptic has to explain all the experiences before envatment: normal experiences and hallucinations. But additionally she has to explain the experiences after the envatment. And for all three kinds of experience a different story has to be told. Thus, the sceptic's explanation provided by SH2 is still more complex.

The standard explanation of our experience along the lines of hypothesis W – which I will call $E(W)$ – involves two subexplanations. One of them contains a set of assumptions [A] to explain the normal perceptual experience.¹⁵ The other contains a set of assumptions [B] to explain hallucinations:

$E(W)$: [A] & [B]

The explanation provided by SH2 – which I will call $E(SH2)$ – explains all experiences before envatment in accordance with $E(W)$. So $E(SH2)$ is bound to the same assumptions [A] and [B]. But $E(SH2)$ also has to explain experiences after envatment and is therefore bound to another set of assumptions [C] (e.g. there is a giant super computer, etc.). Thus, $E(SH2)$ involves:

$E(SH2)$: [A] & [B] & [C]

Irrespective of how complex the standard explanation $E(W)$ turns out to be, the skeptical explanation $E(SH2)$ of our experience will be theoretically more complex. It will be more

¹⁵ What kind of assumptions the set [A] involves depends on the account of explanation one prefers. For example, according to the deductive-nomological model, [A] contains minimally one law of nature (see Hempel 1965: 335-76). According to the causal model of explanation on the other hand, laws are not themselves part of the explanatory information provided by [A] (see Lewis 1986: 239).

complex because it will involve everything $E(W)$ involves but it will be committed to a further set of assumptions as well.

A Further Objection

This picture cannot be right. The skeptical hypothesis SH2 is not a special case of the real world hypothesis W – rather, the two hypotheses are incompatible. Therefore the explanations $E(SH2)$ and $E(W)$ are incompatible explanations of our perceptual experience. But according to the given characterisations of $E(W)$ and $E(SH2)$, the two explanations are not incompatible, so the picture must be wrong.¹⁶

This objection can be countered by adding a further assumption to $E(W)$ and $E(SH2)$ which renders the explanations clearly incompatible. This further assumption states that neither more nor less than the listed sets of assumptions are needed to explain our experience satisfactorily. With respect to $E(W)$ the envisaged assumption is:

- e: To explain our perceptual experience satisfactorily, neither more nor less than [A] & [B] is needed.

And with respect to $E(SH2)$ the additional assumption is:

- e*: To explain our perceptual experience satisfactorily, neither more nor less than [A] & [B] & [C] is needed.

According to this suggestion we get:

$E(W)$: [A] & [B] & e

and

$E(SH2)$: [A] & [B] & [C] & e*.

Obviously, these two explanations are incompatible and the latter is theoretically more complex – it is committed to more assumptions.

Regarding question (1) we can conclude: There is a clear sense in which the explanation provided by SH2 is more complex than the standard explanation of our experience.

¹⁶ Thanks to David Chalmers for highlighting this problem in correspondence.

Question (2): Why is it epistemically justified to prefer the simpler explanation?

The antiskeptical strategy provided by the inference to the best explanation depends on a strong version of the simplicity principle:

Other things being equal, if the explanation E1 is simpler than the explanation E2, then E1 is a better explanation than E2 and it is therefore epistemically justified to prefer E1 to E2. It is justified to prefer E1 to E2, because the simpler explanation is more likely to be true.¹⁷

Is this principle plausible? Is there a way to defend this principle in the light of skeptical reasoning? Even if we could defend the principle on a posteriori grounds, this a posteriori defence would not help in our dialectical situation. Like the a posteriori strategy discussed in section 4.2 (Recollection and Exclusion) the strategy posed by the inference to the best explanation would be doomed to fail. But an a priori defence of the simplicity principle seems inaccessible as well. It would be an empirical fact (if it is a fact at all) that simpler explanations are more likely to be true – on which a priori ground should we think that the world is simple instead of complex?

Admittedly, there is no way to justify the simplicity principle in its general form on a priori grounds. But I think that a special case of the principle can obviously be justified a priori. And the justification of this special case is enough to defend the aforementioned antiskeptical idea.

The A Priori Defence of a Special Case

Let E1 be an explanation committed to the assumptions: a&b, and let E2 be an explanation committed to the assumptions: a&b&c. Let both E1 and E2 be equally powerful and lovely – both explain the same events (or type of events) equally well, both are empirically adequate and both provide the same amount of understanding. Obviously E2 is theoretically more complex than E1. And E2 contains all the assumptions of E1 (plus one more).

¹⁷ I agree with Peter Lipton (1991: 61) that we have to distinguish two senses in which something may be the best explanation. An explanation can be the „loveliest“ such that it would, if correct, be the most explanatory or provide the most understanding. Or an explanation can be the „likeliest“ such that it is the most warranted. In this terms the simplicity principle states: Given E1 and E2 are equally lovely, if E1 is simpler than E2, then E1 is more likely to be true.

The probabilities of E_1 and E_2 are determined by a multiplication of the probabilities of their constituents: $P(E_1)=P(a) \times P(b)$; $P(E_2)=P(a) \times P(b) \times P(c)$. Therefore, if $P(c) < 1$, then $P(E_2) < P(E_1)$.¹⁸ In this case it would be justifiable to prefer E_1 to E_2 .

This reasoning cannot prove that the simplicity principle is true in general but for this special case of increased complexity it can prove a priori that the more complex explanation E_2 is less probable than E_1 . This special case can be characterised as follows: Two explanations E_1 and E_2 are equally powerful and lovely. One of the compared explanations is one conjunct richer than the other. The probability of the additional conjunct is < 1 . Both explanations would be equally probable, if the additional conjunct was omitted.

The Transfer of the Defence to $E_{(SH2)}$ and $E_{(W)}$

The given a priori reasoning can be transferred to the explanations $E_{(W)}$ and $E_{(SH2)}$. Let us first take a look at the two explanations without the assumptions e and e^* :

$E_{(W)}$: [A] & [B];
 $E_{(SH2)}$: [A] & [B] & [C].

Thus far $E_{(SH2)}$ is clearly less probable than $E_{(W)}$. $E_{(SH2)}$ has (at least) one conjunct more than $E_{(W)}$ and without the additional conjunct [C] the probabilities of the two explanations would be identical. Moreover, the probability of [C] is < 1 because the skeptic's additional assumptions are neither scientifically nor metaphysically necessary. Therefore, the more complex explanation $E_{(SH2)}$ is less probable than the simpler explanation $E_{(W)}$.

To render the standard explanation of experience incompatible with the skeptic's explanation we have to add the assumptions e and e^* :

$E_{(W)}$: [A] & [B] & e ;
 $E_{(SH2)}$: [A] & [B] & [C] & e^* .

But this addition would compromise the given reasoning only if the probability of e^* were higher than the probability of e . And this is not a case to worry about.

Therefore, regarding question (2) „Why is it rational to prefer the simpler explanation?“ we can conclude: $E_{(SH2)}$ is more complex than $E_{(W)}$ in a special way. And

¹⁸ I presuppose here and in the following that the probabilities I refer to are independent from each other.

for this special kind of increased complexity a priori reasons justify that the more complex explanation is less probable.

In summary, we can formulate an antiskeptical argument:

- (a) The explanations $E(\text{SH2})$ and $E(W)$ are only distinguished in two respects:
 - In contrast to $E(W)$, $E(\text{SH2})$ contains the additional conjunct $[C]$.
 - In contrast to $E(W)$, $E(\text{SH2})$ involves e^* instead of e .
- (b) Irrespective of the assumptions e , e^* and the additional conjunct $[C]$, the probability of $E(\text{SH2})$ equals the probability of $E(W)$.
- (c) The probability of the additional conjunct $[C]$ is < 1 .
- (d) The probability of e^* is not higher than the probability of e .
- (e) Hence, $E(\text{SH2})$ is less probable than $E(W)$.

Line (a) states that $E(\text{SH2})$ and $E(W)$ are only distinguished in the two mentioned aspects. So we concede to the skeptic that her explanation is as lovely and powerful as the standard explanation of experience.¹⁹ If this was contested, it would be even worse for the skeptic. The antiskeptical conclusion will not be averted this way. Line (b) is true because – irrespective of e , e^* and $[C]$ – the two explanations $E(\text{SH2})$ and $E(W)$ are identical. Line (c) is true because the skeptic's additional assumptions are neither scientifically nor metaphysically necessary. And line (d) is true because if both explanations were empirically adequate, equally lovely and powerful, then there is no reason why the probability of e^* should be higher than the probability of e . Therefore, we are allowed to draw the conclusion (e). It is justified a priori that $E(\text{SH2})$ is less probable than $E(W)$. Thus, we are justified in dismissing the skeptical hypothesis SH2 .

Of course the given reasoning does not entail $\sim\text{SH2}$. But such an entailment is not necessary to block the skeptical argument. Knowledge is fallible in the sense that S can know p on the basis of r where r only makes p probable. In order to know p on the basis of r it is not required that r entails p . Therefore S is in a position to know $\sim\text{SH2}$ on the basis of our reasoning even though this reasoning does not entail $\sim\text{SH2}$. The given reasoning blocks the skeptical argument by means of refuting the skeptical premise (ii)* „ S is not in a position to know that $\sim\text{SH2}$ “.

Furthermore, even if the skeptic argued on the basis of other considerations that our given antiskeptical reasoning is not sufficient to prove that the skeptical premise (ii)* is false, she would at least have to concede that we undermined the justification

¹⁹ That is, we concede that $E(\text{SH2})$ and $E(W)$ both explain the same events (or type of events) equally well, both are empirically adequate and both provide the same amount of understanding.

for her premise (ii)*. We have undermined her justification, because the skeptic justifies (ii)* by the assumption that neither an a priori nor an a posteriori reasoning is available for \sim SH2 (see sections 2.2 and 3.3). And we have just proved that this assumption is false. We have provided an argument which was able to justify that \sim SH2 a priori. Therefore, even if we are not able to prove that premise (ii)* is false, we have blocked the skeptical line of reasoning, because undermining the justification of a premise is also a way of blocking an argument.

5. Concluding Remarks

Semantic externalism forces a modification of the skeptical argument. Regarding this modified argument, the most promising candidate for a successful refutation is the antiskeptical strategy provided by the inference to the best explanation. This antiskeptical strategy is so promising because it exploits the forced modification of the skeptical hypothesis. By focussing on the scenario of recent envatment SH2, the most difficult problems with the strategy posed by the inference to the best explanation can be solved:

- There is a clear sense in which the explanation of our experience provided by SH2 is more complex than the standard explanation.
- It is justified a priori that the more complex explanation SH2 is less probable than the standard explanation.

Therefore, by a combination of the strategies offered by semantic externalism and the inference to the best explanation, the skeptical argument can be blocked successfully.

Notice that this antiskeptical suggestion depends on the following thesis concerning the theory of explanation for which I haven't argued in this paper:

If a general explanation of a phenomenon is divided in subexplanations: SE1-SEn, because the phenomenon is inhomogeneous, then E is a conjunction of SE1-SEn.

If this thesis were wrong, then the proposed antiskeptical strategy would not get off the ground.

References

- Boghossian, P. (1989), 'Content and Self-Knowledge', *Philosophical Topics* 17, 5-26.
- Brueckner, T. (2004), 'Brains in a Vat', IN E.Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Winter 2004 Edition.
<http://plato.stanford.edu/archives/win2004/entries/brain-vat/>.
- Chalmers, D. (2005), 'The Matrix as Metaphysics', IN C.Grau (ed.), *Philosophers Explore the Matrix*, Oxford: Oxford University Press. (See also:
<http://consc.net/papers.html>)
- Dretske, F. (1970), 'Epistemic Operators', *The Journal of Philosophy* 67, 1007-23.
- Hempel, C (1965), *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, New York: Free Press.
- Hawthorne, J. (2003), 'The Case for Closure', IN M.Steup and E.Sosa (eds), *Contemporary Debates in Epistemology*, Oxford: Blackwell, pp. 26-43.
- Lewis, D. (1996), 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-567.
- Lewis, D. (1986), 'Causal Explanation', IN D.Lewis, *Philosophical Papers, Volume II*, Oxford: Oxford University Press, pp. 214-40.
- Lipton, P. (1991), *Inference to the best Explanation*, London/New York: Routledge.
- Müller, O. (2003), *Wirklichkeit ohne Illusionen, Bd. I*, Paderborn: Mentis.
- Pryor, J. (2000), 'The Skeptic and the Dogmatist', *Nous* 34, 517-549.

Putnam, H. (1975), 'The Meaning of Meaning', IN H.Putnam, *Mind, Language and Reality, Philosophical Papers Volume II*, Cambridge: Cambridge University Press, pp. 215-71.

Putnam, H. (1981), *Reason, Truth and History*, Cambridge: Cambridge University Press.

Vogel, J. (1990), 'Cartesian Skepticism and Inference to the best Explanation', *Journal of Philosophy* 87, 658-666.

Weatherson, B. (2007), 'The Bayesian and the Dogmatist', *Proceeding of the Aristotelian Society* 107, 169-185.

Williamson, T. (2000), *Knowledge and its Limits*, Oxford: Oxford University Press.

Wright, C. (1991), 'On Putnam's Proof that we are not Brain-in-a-Vat', *Proceedings of the Aristotelian Society 1991*: 67-94.