



This is a repository copy of *Transcendental Arguments: A Plea for Modesty*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/81346/>

Version: Accepted Version

Article:

Stern, R. (2007) *Transcendental Arguments: A Plea for Modesty*. *Grazer Philosophische Studien*, 74. 143 - 161. ISSN 0165-9227

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

TRANSCENDENTAL ARGUMENTS: A PLEA FOR MODESTY

Robert Stern

Abstract: A modest transcendental argument is one that sets out merely to establish how things must appear to us or how we must believe them to be, rather than how things are. Stroud's claim to have established that all transcendental arguments must be modest in this way is criticised and rejected. However, a different case for why we should abandon ambitious transcendental arguments is presented: namely, that when it comes to establishing claims about how things are, there is no reason to prefer transcendental arguments to arguments that rely on the evidence of the senses, making the former redundant in a way that modest transcendental arguments, which have a different kind of sceptical target, are not.

‘[If] any man were found of so strange a turn as not to believe his own eyes, to put no trust in his senses, nor have the least regard to their testimony, would any man think it worth while to reason gravely with such a person, and, by argument, to convince him of his error? Surely no wise man would’ (Reid 1863: 230)

Although it has never been an issue quite at the centre of recent epistemology, the promise and potential of transcendental arguments has received a good deal of discussion.¹ Up until now, the main focus of that discussion has been whether such arguments work, and in particular whether they can be used successfully to establish certain facts about the world in a sceptic-proof manner, where Barry Stroud's influential paper from 1968 has persuaded many that they cannot.² I want to suggest here, however, that the argument of this paper is not as persuasive as is widely supposed. Nonetheless, I will claim, there is another way of criticising transcendental arguments to the same effect, which is to show that our hopes for such arguments should be modest and not ambitious. I will begin by sketching what I take

¹ For a bibliography, see Stern 1999: 307-21.

² I will refer to the reprinted version in Stroud 2000: 9-25. Stroud's was not the only critical voice from this period: another significant critic of transcendental arguments was Stephan Körner; see for example Körner 1967.

transcendental arguments to be, and then consider Stroud's critique of them, before criticising this and offering an argument for modesty of my own. That argument hinges on whether there is any reason to prefer a transcendental argument as a response to scepticism over some other sort of response, for example, one that relies on the evidence of our senses? I will argue that in the ambitious way these arguments are commonly conceived, against the sceptic who is commonly taken to be their target, there is no reason to so prefer them; but if our conception of these arguments is made more modest, there is room to think they offer us something additional to other anti-sceptical manoeuvres, and so that it is here that their main value should be seen to lie.

I

While the exact nature of transcendental arguments is far from unproblematic, they are generally taken to have the following features: They begin from some sort of self-evident starting point concerning our nature as subjects (for example, that we have experiences of a certain kind, or beliefs of a certain kind, or make utterances of a certain kind) which the sceptic can be expected to accept, and then proceed to show that this starting point has certain metaphysically necessary conditions, where in establishing that these conditions obtain, the sceptic is thereby refuted. So, in the face of the sceptical suggestion that we do not know that there is an external world, or other minds, or the past, a transcendental argument might be offered to provide deductive support for these claims from certain facts about our nature as subjects, based on the premise that the former are necessary conditions for the latter, where the form of the argument is: we have certain experiences etc; a necessary condition for us having these experiences etc is the truth of S; therefore S.

However, it is now widely held that this kind of argument is more problematic than it may at first appear. A highly influential source of this suspicion is Barry Stroud, who in his article 'Transcendental Arguments' suggested that for any claim concerning the necessary condition S, 'the sceptic can always very plausibly insist that it is enough [that] we believe that S is true, or [that] it looks for all the world as if

it is, but that S needn't actually be true' (Stroud 2000: 24).³ So, in the case of the problem of the external world, for example, the concern is that no argument can be constructed to show that there must actually be an external world, but just that there must appear to us to be one, or that we must believe there to be one.

In subsequent work, Stroud has gone on to explain why the sceptic can 'very plausibly' weaken the necessary condition for experience etc from S to 'we must believe S' or 'S must appear to be true'. For, while he allows that we might reasonably be able to make modal claims about 'how our thinking in certain ways necessarily requires that we also think in certain other ways', he thinks it is puzzling 'how...truths about the world which appear to say or imply nothing about human thought or experience' (for example, that things exist outside us in space and time, or that there are other minds) '[can] be shown to be genuinely necessary conditions of such psychological facts as that we think and experience things in certain ways, from which the proofs begin'. Stroud goes on: 'It would seem that we must find, and cross, a bridge of necessity from the one to the other. That would be a truly remarkable feat, and some convincing explanation would surely be needed of how the whole thing is possible' (Stroud 2000: 158-9).⁴ Thus, Stroud is prepared to allow (and indeed exploits our capacity himself, in his own arguments against the sceptic)⁵ 'that we can come to see how our thinking in certain ways necessarily requires that we also think in certain other ways, and so perhaps in certain other ways as well, and we can appreciate how rich and complicated the relations between those ways of thinking must be' (Stroud 2000: 158-9); but he believes that anything more than this, which asserts that 'non-psychological facts' about the world outside us constitute necessary conditions for our thinking, is problematic.

Faced with Stroud's challenge, it has appeared that there are three ways to go. First, one can opt for idealism, which sees no gap to bridge between how we think and

³ In this 'Transcendental Arguments' paper, the starting point for which S is meant to be a necessary condition is language, where if the transcendental claim could be established, S could be shown to be true from the fact that what the sceptic says makes sense. But, as many transcendental arguments have been proposed that are not just focussed on the conditions for language, I take it that Stroud's worry here cannot just apply to transcendental arguments of this sort, but also those that focus on conditions for experience, self-consciousness etc.

⁴ Cf. also Stroud 2000: p. 212: 'All this would be so on the assumption that transcendental arguments deduce the truth of certain conclusions about the world from our thinking or experiencing things in certain ways. That strong condition of success is what I continue to see as the stumbling-block for such ambitious transcendental arguments. Can we ever really reach such conclusions from such beginnings?... [The most troubling danger is] that of not being able to reach substantive, non-psychological truths from premises only about our thinking or experiencing things in certain ways'.

⁵ See, for example, Stroud 2000: 165 ff.

how things are insofar as the former determines the latter, and then try to qualify this position in such a way as to render it somehow plausible. Second, one can opt for verificationism, which stipulates that some of what we believe about the world must be true. But both these options are seen as problematic in themselves, and as sufficiently anti-sceptical on their own to make any appeal to transcendental arguments redundant. A third possibility, however, is to opt for what can be seen as a more ‘modest’ approach. On this view, it has been suggested that we accept that transcendental arguments should not attempt to cross Stroud’s ‘bridge of necessity’ at all; instead, we should allow that the only necessary conditions that we can establish concern how we must think or how things must appear to us, thus avoiding Stroud’s call for an explanation of how we can get from the ‘psychological’ to the ‘non-psychological’ by remaining within the former, and eschewing claims about the latter. Stroud’s own position *viz-à-viz* the sceptic involves this sort of modest approach, although there are other ways to take it.

However, before adopting any such ‘modest’ strategy, the question should now be asked: how powerful is Stroud’s position here? I think it is less compelling than is generally supposed.⁶ For, according to Stroud, there is something inherently problematic in making a modal claim about how the world must be as a condition for our thought or experience, but there is not anything particularly problematic about making a claim about our thought or experience being a condition for some other aspect of our thought or experience. But why should it be somehow easier to make modal claims between ways of thinking or types of experience, than ways of thinking or types of experience and the world? Why are such ‘bridges’ or modal connections easier to make ‘within thought’ than between how we think and how the world must be to make that thought possible? Of course, one might take this symmetry between the two to be reason to be suspicious of modal claims of this sort at any level: but as we have seen, Stroud himself seems to think they are viable between ‘psychological facts’.⁷ If so, I believe, he needs to give us some account of why they are less

⁶ Whilst it is commonly taken as a starting point, Stroud’s position has of course not gone totally uncriticised: see e.g. Glock 2003: 37-9. Glock cites an anticipation of Stroud’s position from the lectures of C. D. Broad: ‘What Kant claims to prove by his transcendental arguments is that certain propositions, such as the law of causation and the persistence of substance, are true with the interpretation and within the range he gives them. But it is doubtful whether his arguments could prove more than that all human beings must believe them to be true, or must act as if they believed them to be true’ (Broad 1978: 15). Broad gives no reason to substantiate his doubt here.

⁷ Cf. Stroud 2000: 224-44, where Stroud defends the possibility of modal knowledge but still only considers modal claims involving ‘psychological facts’, such as that ‘thinking of the world as

problematic here than between our thought and the world; but as far as I can see he just takes it to be obvious, and so provides no such account.

Perhaps, however, it could be said that Stroud is right to take this to be obvious: for, if we are dealing with modal connections between us and the world, and the world is conceived in a realist manner, as independent of us, then of course we know less about how we depend on the world than on how we depend on other facts about us – this is just a feature of the mind-independence of the world, which makes it more opaque to us in this way, so that our modal claims concerning it are correspondingly more problematic, than they are concerning connections between our ways of thinking.

Now, of course, it is often taken to be the case that self-knowledge (knowledge of our ‘inner states’) is less problematic than worldly knowledge, on these sorts of grounds, and Stroud’s position may be taken to be trading on that intuition. But if so, I think it is mistaken. For, even if we grant that certain sorts of self-knowledge are unproblematic in this way, on grounds of their immediacy or self-evidence or infallibility, it seems unlikely that the modal claims involved in transcendental arguments would have these features, even if they involve merely ‘psychological facts’ rather than claims about the world. If we make some sort of claim regarding the dependence of one aspect of our thought or experience on some other aspect of our thought or experience, what reason have we for thinking this is somehow self-evident or immediate in the way that (perhaps) ‘I am in pain’ is self-evident and immediate – even though the former as well as the latter involve ‘facts about us’ rather than the world outside us? It seems implausible that in the transcendental case, we have some sort of privileged ‘first person access’ of the sort that might be used to establish a relevant epistemic difference in the non-transcendental case between ‘I am in pain’ and ‘You are in pain’, as what is involved in the transcendental case is not introspection but the use of modal intuition. Once again, therefore, there seems no reason to support Stroud’s view concerning the asymmetry between ambitious and modest transcendental arguments, and so his claims for the greater viability of the latter over the former.

It appears, then, that we have little reason to accept Stroud’s concerns that a special ‘bridge of necessity’ is needed if we are to make transcendental claims that

contingently containing subjects of experience would require thinking of it as contingently containing objective particulars independent of experience as well’ (p. 236).

take us outside ‘psychological facts’. If Stroud’s position is to be rejected in this way, however, does that mean we should go back to conceiving of transcendental arguments in more ambitious terms? I do not think so, as I believe there is another reason that can be given for modesty, which is more compelling than the one offered by Stroud.

II

This argument for modesty is really rather simple, and relates to the dialectic with the sceptic. The central thought is this: It is implausible to think that the sceptic would be satisfied with the use of any transcendental argument against him, given what has driven him to be a sceptic in the first place. We must either prevent him being driven to scepticism at some earlier point, in which case an appeal to transcendental arguments will be unnecessary; or we must accept that we cannot so prevent this, but by then a transcendental argument will come too late.

In order to see this, let me begin by characterising the sort of sceptic that an ambitious transcendental argument is supposed to convince, where I will focus on the problem of the external world. This sceptic is someone who has her reasons for doubting the truth of what most (or maybe all) of us believe we know: namely, that there is an external world of material objects outside us in space and time. This doubt is based on the thought that the kind of evidence we would present in favour of this knowledge is inadequate, where I take it that a large part of that evidence is perceptual. There are of course a number of arguments that the sceptic thinks she can give to show that this evidence is inadequate in this way, but most hinge on some sort of argument from error: we just cannot be sure that this evidence is sufficient to support our belief, given the compatibility of that evidence with various so-called ‘sceptical scenarios’, such as Descartes’ evil demon or the brain-in-a-vat hypothesis.

Now, it can be tempting to think that a transcendental argument is just what we need here, because it can seem that our difficulty is a fallibilist one: we must admit to the sceptic that our perceptual experience is fallible, so we cannot rule out the sceptical scenario on the basis of how things appear to us, as that appearance could be radically misleading. If the transcendental argument can be made to work, however, it

has the advantage of being a deductive strategy, and thus error-proof – where it can be taken for granted that a sceptic who questions our reliance on the laws of logic has taken a step too far, even for a sceptic.

The difficulty, however, is that while a transcendental argument is indeed a deductive argument, it relies on certain premises which involve modal claims, about the world having to be a certain way in order for our experiences or thoughts to be possible. The question is, therefore, how can we satisfy the sceptic of our entitlement to make such modal claims?

Now, as much of the recent literature in this area suggests, our capacity to make such claims is somewhat mysterious – so mysterious, in fact, that even some who are not in any way sceptics would argue that we should eschew them.⁸ However, my point here is not that general: let us assume that we do make such claims in a way that can be satisfactorily understood, so that we should have no doubts about their legitimacy on that score. Nonetheless, my worry is that if transcendental arguments rely on these modal claims, it is hard to see how they could then put us in a dialectically advantageous position with respect to the sceptic: for if the sceptic thinks she can question our perceptual evidence for the existence of an external world, can't she on very similar grounds question our modal intuitions for such claims of necessity – viz. that they are equally prone to error? For if (as seems plausible) we rely on criteria like conceivability or imaginability to test such modal claims, can't the sceptic plausibly say that our capacities here can go wrong, to the same degree as in the perceptual case – so how can the use of such claims make us better off?

One way to respond to a worry of this kind has been suggested recently by Thomas Grundmann and Catrin Misselhorn. They have argued that although a transcendental argument relies on our modal intuitions, this is not problematic, because the sceptic relies on such intuitions as well, in claiming that her sceptical scenario is metaphysically possible. They write: 'If the sceptic claims that modal intuitions are unreliable, sceptical hypotheses could not be justified. For this reason, the sceptic must grant the reliability of modal intuitions as a method of justification' (Grundmann and Misselhorn 2003: 211). Thus, Grundmann and Misselhorn think they can get our modal intuitions to speak in favour of a claim like 'Necessarily, perceptual beliefs about the external world are largely true', which can then be used

⁸ The literature in this area is growing steadily. For a useful collection with a good selection of articles, see Gendler and Hawthorne 2002.

as a premise of a transcendental argument to the effect that we have perceptual knowledge of the external world, in a way that refutes the sceptic in an ambitious manner (cf. Grundmann and Misselhorn 2003: 207).

This is an interesting approach; I am not sure, however, that it properly does justice to the dialectic with the sceptic. It is certainly true that in appealing to sceptical scenarios, the sceptic tries to exploit the fact that he can make them seem metaphysically possible to us, where he uses our modal intuitions to do so. Thus, when our intuitions are going his way, so to speak, he is happy to exploit them. But suppose that Grundmann and Misselhorn are right, and that our intuitions can be made to go the other way, in support of a claim like ‘Necessarily, perceptual beliefs about the external world are largely true’. What is to prevent the sceptic now changing his tune, and questioning our reliance on these intuitions? The sceptic, after all, is not someone with a settled position of his own to support: he will use whatever means are at his disposal to generate doubt – and if it turns out that using our modal intuitions is not an effective way to do so, because in the end they do not support his sceptical scenarios, why shouldn’t he abandon them?

Of course, if he gives up appealing to his sceptical scenarios, the sceptic will still need something to base his doubt upon, and thus will now need some reason to question Grundmann and Misselhorn’s transcendental argument and the modal intuitions it relies upon. But we have already seen what the ground of that doubt could be: namely, the claim that there is a possibility of error in the intuitions they use in support of their modal claim, such that the proof must remain open to question.

Now, Grundmann and Misselhorn might reasonably respond to this by saying: all this sceptical argument amounts to is an argument from fallibilism, and that ‘[i]t is generally accepted that fallibilism is not sufficient to generate scepticism’ (Grundmann and Misselhorn 2003: 210), so the mere fact that our modal intuitions about this modal claim might be wrong is not a reason to doubt it – the sceptic has to have ‘proper modal intuitions speaking against it, and, as we have argued, the sceptic has not provided convincing modal evidence for his claim, so far’ (Grundmann and Misselhorn: 218).

However, if we adopt this sort of strategy in defence of the use of transcendental arguments, what is to prevent us adopting it from the beginning of our debate with the sceptic, and apply it to the perceptual case: that is, why can’t we dismiss merely fallibilistic arguments against our perceptual evidence for the

existence of the external world, and ask the sceptic to come up with ‘proper evidence speaking against it’ – for example, that there are brains-in-vats experiments going on, that scientists are available to conduct such experiments, and so on? In the absence of such evidence, why won’t the ‘straight’ perceptual evidence that we have for the existence of an external world do – making any appeal to a transcendental argument redundant in our response to the external world sceptic? To put this in the Reidian terms of my epigraph, by trying to use a transcendental argument against the sceptic, aren’t we doing what no sane man should?

I think similar considerations would tell against another way of using transcendental arguments ambitiously. This would be to use them not in support of claims about the world in the face of arguments from error, but in support of claims about the reliability of our belief-forming methods, in the face of arguments from circularity – namely, we must presuppose such methods in order to establish their reliability.⁹ Either we should block such scepticism from the beginning,¹⁰ or the transcendental argument comes too late: for if we offer a transcendental argument to the effect that (for example) perceptual experience is reliable, as this reliability is a condition for something else, the question remains, how can we establish the reliability of the modal intuitions we employ in constructing the transcendental argument? Once again, therefore, it would seem that an appeal to transcendental arguments will not help in this situation.

If I am right, then, it turns out that there is no reason to think that the use of transcendental arguments gives us any advantage over the sceptic as he has been conceived so far: either he can be defeated some other way,¹¹ or if not, transcendental arguments gives us no additional advantage against him.

⁹ For a useful discussion of the issues, see Alston 1993.

¹⁰ For an attempt to do so, see Stern 2003: 229-32.

¹¹ Stroud, however, is not in a position to accept the argument against ambitious transcendental arguments sketched in this section, which is perhaps why he resorts to the argument against them outlined in section I. For, while he accepts that it is a central feature of the sceptic’s position that he raises doubts based on possibilities for which no ground can be given, Stroud thinks we cannot reject scepticism simply on that score, as we could in the case of ‘ordinary’ inquiries – where the sceptic’s ‘extraordinariness’ is on Stroud’s view a corollary of the ‘extraordinariness’ of the epistemological project itself. Ultimately, then, Stroud counsels that we should not try to answer the sceptic, but question the epistemological project that makes scepticism possible. See, for example, ‘Taking Scepticism Seriously’ and ‘Understanding Human Knowledge in General’, both reprinted in Stroud 2000.

III

It might be argued, however, that weaknesses in the kind of non-transcendental, fallibilistic response I have given to the sceptic will lead to ambitious transcendental arguments being required after all. For, it could be said, these fallibilistic responses work by claiming that because the sceptic cannot give us any grounds for actually taking his sceptical scenario seriously and thinking that we are brains in vats etc., these scenarios pose no epistemic threat to our ordinary beliefs, as if they are mere logical possibilities, they do not provide us with sufficient reasons for doubting those beliefs. However, the sceptic could respond by arguing that even if he cannot give us any reasons for thinking that his sceptical scenarios are actually the case, they are still a threat, despite the fact that he cannot provide any evidence in their favour. I will here consider three sceptical strategies that might seem to undercut the fallibilist response to scepticism in this way, and whether or not these strategies require us to appeal to ambitious transcendental arguments if we are to deal with them.

A first sceptical strategy that seems to require no positive evidence in favour of the sceptical scenario I will call the simple tracking argument. On this strategy, it is claimed that the sceptical scenarios show that a belief like ‘I have two hands’ fails to meet a fundamental tracking or sensitivity requirement on knowledge: viz. that if p was false, A would not believe that p . To show that I would fail to meet this requirement for knowing ‘I have two hands’, the sceptic does not have to show that it is false or that I have good reason to think it is because I have good reason to think I am a brain in a vat: he just has to show that if I were a brain in a vat, it would be false, but I would continue to believe it, thereby showing (he argues) that my belief ‘I have two hands’ violates the tracking requirement for knowledge.

In the face of this sceptical position, it may seem tempting and indeed obligatory to return to some sort of ambitious transcendental argument. For, in response to the tracking problem, it could be claimed on the basis of such an argument (of the sort suggested by Putnam, for example),¹² that the sceptic is wrong to suggest that in the sceptical scenario my belief ‘I have two hands’ would fail to track the truth,

¹² See Putnam 1981: Chap 1. Putnam himself takes as his target the very coherence of the sceptical hypothesis; but I am here suggesting that some of his claims about reference could also be used to resolve the tracking problem, on the grounds that if I were a brain in a vat my belief ‘I have hands’ would have a different ‘vatted’ meaning, and so not fail to track how things are, because how things are would shape the content of my belief.

because I could not be a brain in a vat while still believing that I have hands, as my belief would have a different reference, so that the tracking condition on this belief (and others like it) can be guaranteed not to fail in this manner.

In fact, however, I think the fallibilist has a perfectly adequate response to the sceptic here, without recourse to any ambitious transcendental argument of this kind. Of course, if the fallibilist is not to avail himself of an argument like Putnam's, he must allow that if he were a brain in a vat, he would continue to believe that he has hands, and thus that the tracking requirement would fail in this respect. But the question here is whether it is a plausible condition on knowledge that if p were false, A would not believe that p in all situations in which p might be false? What makes it implausible that this condition holds is precisely what makes infallibilism implausible: namely, that we can know things, even when our grounds for knowing them or our methods for knowing them are prone to lead to error in some circumstances. What matters, of course, is in what circumstances we are prone to error: if we are error-prone in circumstances that would require a lot of manipulation¹³ in order for us to make the error (such as would be needed to make the brain in a vat scenario work) then the fact that we could not track the truth in these circumstances arguably does not count against our knowing the truth in more normal ones. Rather than telling against fallibilism, therefore, all the tracking objection reveals is one of its consequences, namely that knowledge does not require a logical entailment between not- p and A not believing p :¹⁴ it all depends on how and why that relation breaks down, where it can be claimed that in the case of the sceptical scenario, that breakdown would not be enough to show that our capacity to track 'I have two hands' is inadequate for knowledge.

It appears, then, that the simple tracking argument can be defeated without any need to appeal to an ambitious transcendental argument. However, this strategy focused on our belief 'I have two hands', and directly claimed (implausibly as it turned out) that that belief fails to meet the tracking requirement, because it would not

¹³ There are different ways this idea could be worked out. One could be in terms of possible worlds, namely, that is only in worlds some distance from the actual one that we will be fooled. Another way might be in terms of normal functioning, namely, our cognitive mechanisms would require serious distortion for us to be misled. And of course, depending how 'normal functioning' is spelt out, these approaches may end up converging.

¹⁴ Cf. Nozick 1981: 199, who observes that the tracking condition we have been considering 'does not say that in all possible situations in which not- p holds, S doesn't believe p . To say there is no possible situation in which not- p yet S believes p , would be to say that not- p entails not- $(S$ believes $p)$, or logically implies it',

do so in a sceptical scenario. However, the sceptic might now offer a more complex tracking argument, which I will call the tracking plus closure argument. Rather than focusing on the belief 'I have two hands', this argument focuses on the belief 'I am not a brain in a vat' and claims that this fails the tracking argument; it then follows from the closure principle that 'I have two hands' is not known.

The first step in this strategy, then, is to claim that I do not know I am not a brain in a vat because this belief fails to track: if I were a brain in a vat, I would still believe I am not. But why can't we respond to this sceptical point as before: namely, why can't the belief 'I am not a brain in a vat' be said to only fail to track that it is false in exceptional circumstances, namely when we are envatted, just as the belief 'I have two hands' only fails to track the truth in such circumstances? If this failure doesn't prevent us knowing we have hands in the latter case, how does it prevent us knowing that we are not brains in vats in the former? However, I think there is something different in the brain in a vat case which blocks this sort of response: for, while there are circumstances in which I would track the falsity of 'I have two hands' (for example, if they were chopped off in an accident, or as a result of disastrously incompetent surgery) but just fail to do so in extreme circumstances like being envatted, the sceptic can claim that there are no circumstances in which I would track the falsity of 'I am not a brain in a vat' – because if it were false, and so I was a brain in a vat, then *ex hypothesi* I would never pick this up and so never change my belief accordingly.¹⁵ Thus, while an externalist might say that in both cases my belief only fails to track in a remote possible world, I would be prepared to grant the sceptic that these beliefs are not on a par, and that a successful tracking argument can be made against my belief 'I am not a brain in a vat'.

The second step of the sceptical argument is then to go from this admission to the conclusion that I do not know I have two hands, via the closure principle: if A knows p, and A knows that p entails q, then A knows q. Using *modus tollens* on this principle, it can be argued that as we don't know we are not brains in vats, then we don't know we have two hands either (as it seems right to say that if we have hands, this entails we are not brains in vats). It would seem, then, that even without giving us

¹⁵ So, it will not do here to say that I might come to pick up the falsity of 'I am not a brain in a vat' in some circumstances, for example if the evil scientist were not very competent and sent me some information that tipped me off. I am taking it that it is part of the sceptic's conception of what it is to be a brain in a vat that the scientists concerned never make such errors. If such malevolent perfection seems implausible to attribute to human scientists, substitute evil demon scientists instead.

any reason to think we actually are brains in vats, the sceptic can use his scenario to undermine our ordinary beliefs, such as the belief that we have hands.

Now, of course, there are a variety of responses to this problem in the literature which do not employ any sort of transcendental argument strategy, such as approaches that deny the closure principle, or that reject the tracking requirement on knowledge, replacing it with a weaker requirement that our belief ‘I am not a brain in a vat’ can be said to meet.¹⁶ The fallibilist can therefore explore ways out of this difficulty without being obliged to adopt a transcendental argument. But still, it might be suggested, if a transcendental argument can be used against the sceptic here without requiring any such manoeuvres, that might seem to show that transcendental arguments have a significant role to play against the sceptic in enabling us to answer him without modifying what appear to many to be plausible epistemic principles.

I would once again claim, however, that the promise of transcendental arguments in this respect is once again illusory. It may seem that the way in which such an argument could be used is in relation to the tracking issue. For, it could be said, arguments such as Putnam’s show that the tracking requirement for ‘I am not a brain in a vat’ can be met, by showing that this is something we could not falsely believe: for, if I were a brain in a vat, thinking ‘I am not a brain in vat’ would not refer to brains in vats but something else (perhaps vat images).¹⁷ The difficulty is, however, that although this sort of transcendental argument meets the tracking requirement, it can nonetheless be plausibly claimed that there are further conditions it does not meet. For, it doesn’t show what also seems to be needed, which is that the grounds on which we form the belief ‘I am not a brain in a vat’ do not prevent us from believing it when it is false; the transcendental argument just shows that semantic externalist conditions on reference make it impossible to believe ‘I am not a brain in a vat’ falsely, much as physical conditions on belief make it impossible to believe ‘I am alive’ or ‘There is oxygen in the room’ falsely, without in itself showing that the grounds we have for that belief are what make us sensitive to its truth and falsity – and this seems to be what is required for knowledge. If I am right, we will therefore be obliged to find other ways of responding to the sceptical challenge such as those mentioned above, that do not employ transcendental arguments of this sort.

¹⁶ For the former approach, see e.g. Nozick 1981: 197-247, and Dretske 1970. For the latter approach, see e.g. Sosa 1999.

¹⁷ Putnam 1981: 14-15.

The same sort of difficulty applies to another way of using a transcendental argument in relation to the tracking issue, which is even stronger than the semantic externalist one we have just considered. This would be to try to meet the tracking problem by using a transcendental argument to show that ‘I am not a brain in a vat’ is akin to a necessarily true mathematical proposition:¹⁸ for, it could be said that the sceptical scenario is metaphysically impossible, because we could not be brains in vats, as being envatted would prevent us from being believers at all. If this is right, then it would seem that the sceptical scenario is necessarily false, as it could never be actualised; so just as ‘if $2 + 2 = 4$ were false’ is a necessarily false supposition, doesn’t this show that the same is true of supposing the falsity of ‘I am not a brain in vat’?

The difficulty here, however, is that all this argument would show is not that the tracking requirement has been met by my belief ‘I am not a brain in a vat’, but rather that the antecedent of the tracking conditional is necessarily false, as is also the case for necessary truths like $2 + 2 = 4$; but then, as Nozick has argued, it seems that it is best to say not that the tracking requirement has been satisfied and so that our belief constitutes knowledge in this respect, but that it is not a requirement at all (Nozick 1981: 186-7). Now of course, even if ‘I am not a brain in a vat’ cannot turn out to be falsely believed by me, this does not in itself show that I know it, any more than the fact that ‘ $2 + 2 = 4$ ’ cannot be falsely believed by me shows I know it either: I still need some adequate grounds for believing each of them, where it is precisely those grounds that the sceptic questions, in the transcendental argument case as much as in the ordinary perceptual one. We therefore still need to be told why the sceptic should take the transcendentalist’s reasons for believing he is not a brain a vat more seriously than the non-transcendentalist’s.

Finally, we can consider a third sceptical strategy that again treats the sceptical scenario as a mere possibility, which might be called the priority argument. Here, the sceptic claims that there must be a certain order to our knowledge, such that to know familiar things like ‘I have two hands’ I must already know things like ‘I am not a brain in a vat’. But then, the sceptic argues, on what grounds could I know that I am not a brain in a vat, if all the empirical premises from which I might infer this can only be accepted after the hypothesis has been refuted? Unless we are prepared to

¹⁸ Of course, it can’t be exactly the same, as I do not exist in all possible worlds, whereas numbers (arguably) do.

follow G. E. Moore, and argue from these empirical premises directly against the sceptical hypothesis, it seems like we are required to argue against it while being deprived of any basis on which to do so, where again the sceptic here does not need to provide any positive evidence in its favour in order to make his case.

Once again, in the face of this difficulty, it may seem that only something like an ambitious transcendental argument can give us what is needed: for, it gives us an argument against the sceptical scenario which is a priori, and which can therefore be used to first show that we are not brains in a vat, in a way that seemed to be required before we could lay claim to any of our ordinary empirical knowledge. We might therefore grant the sceptical suggestion that we must know that we are not brains in vats before we can know that we have hands etc., but employ a transcendental argument to show that this requirement can be met.

However, the problem with this way of using transcendental arguments is that the sceptic's priority argument is less than compelling, despite being plausible on the surface. For, the crucial move is to say that we cannot know an empirical proposition like 'I have two hands' unless we already know that the sceptical hypothesis is false and can rule it out. But why should we accept this move? Consider the following propositions:

a: my copy of War and Peace is in my study

and

b: my copy of War and Peace has not been stolen.

Do I have to have evidence for b before I can come to know a? The answer would seem to depend on the circumstances. If I already know (or have reason to believe) that a lot of stealing of Russian classics has been going on, then even if I clearly remember putting War and Peace in my study this morning, have a generally good memory etc., that may not be sufficient grounds for knowing a, unless I am in a position to rule out b. But in different circumstances, where as far as I know stealing of this sort never or very rarely happens, why must I establish b before my evidence for a can be accepted? It seems to me this is not required; and of course, the fallibilist argues that we find ourselves in this latter situation when it comes to the sceptical scenario, where we have no grounds for believing this scenario to hold. If this is the right approach to take (as I believe it is) it undercuts the need to argue against the sceptical scenario in an a priori manner using a transcendental argument as a first step: we can just claim (as before) that the lack of evidence for the sceptical scenario

is sufficient to justify us in accepting our ordinary empirical beliefs for the usual reasons, such as perceptual evidence, memory, testimony and so on.

IV

It appears, then, for reasons rather different from those presented by Stroud, but more to do with the dialectical situation involved in our debate with the sceptic, that ambitious transcendental arguments have little work to do.

Does that mean that all forms of transcendental argument have little work to do however – modest ones included? I would suggest not.¹⁹ For, while it turns out that transcendental arguments have little to add in the battle against the external world sceptic who thinks our perceptual evidence is insufficient because error-prone, this is not the only way to be an external world sceptic. For example, one can argue that the perceptual content of our experience does not tell us anything about an external world as we believe it to be, but that we get such beliefs by inferring from a more impoverished perceptual content (as on Hume's view that 'Tis commonly allow'd by philosophers, that all bodies, which discover themselves to the eye, appear as if painted on a plain surface, and that their different degrees of remoteness from ourselves are discover'd more by reason than the senses' [Hume 1978: 56]), where the issue then is how such inferences can be justified. In this situation, I have argued elsewhere,²⁰ a transcendental argument that concerns merely the perceptual content of our experience (how things appear to us) can be useful, in making our beliefs concerning the external world direct and perceptual rather than indirect and inferential.

Now, if we adopt a target of this kind for our transcendental argument, I think we can avoid the dialectical difficulty we faced previously. Previously, it seemed dialectically inappropriate to use modal claims against the sceptic who argues from error, when the grounds for making such claims seem at least as vulnerable to error as the empirical grounds (such as perceptual experience) that these claims are supposed to replace; but there is less dialectical incongruity in making transcendental claims against the kind of sceptic I have just sketched. For in the latter case, we really are in

¹⁹ For further discussion see Stern 2000.

²⁰ See Stern 2000: Chap 4.

the sort of situation viz-à-viz the sceptic that Grundmann and Misselhorn hoped we were in previously, as the sceptic is interested in doing more than offer a mere argument from error and fallibility to question our belief; she is trying to show not just that our evidence is fallible, but that the belief is not justified in our own terms, given the problematic nature of the inference from how she takes things to appear to us to how we think they are, where the ‘veil of appearances’ seems to mean that ‘reason’ could not have enough to go on in making any such inference. Similarly, therefore, the sceptic must do more than just offer an argument from fallibility against a transcendental claim that how things appear to us must be rich enough to support our belief in the external world on direct perceptual grounds (for example, contra Hume, that the world is immediately presented to us in three dimensions, so that this is not ‘discover’d more by reason than the senses’); instead, he must give us grounds for thinking that this claim is not properly supported by our modal intuitions, where then he must show that experience on his impoverished model would be sufficient to be a condition for us to be the kind of conscious creatures we are, contra our transcendental argument against him. The sceptic is thus not in a position to make a general argument against our reliance on our modal intuitions to support our transcendental claim about how our perceptual experience must be, so if our intuitions can be made to speak in favour of this claim, he cannot shrug them off as he did previously, but must show why in this case our intuitions are in fact mistaken, or can be made to go the other way. The dialectic of this situation, then, gives the transcendental argument some genuine work to do, with some prospect of success – but only a transcendental argument of a modest kind.

If I am right, however, that a transcendental argument can only be used successfully in this modest way, against this form of sceptic, what about the sceptic of the more radical kind, who argues from various sceptical scenarios, and against whom it seemed that an ambitious transcendental argument might be needed? If we have abandoned the latter form of argument, does this sceptic therefore win the day? And if so, if we use a transcendental argument to defeat the more modest sceptic, can’t she remain a sceptic by turning more radical?

In abandoning the use of an ambitious transcendental argument against the radical sceptic, however, all that was claimed was that there is no reason to think these arguments should be any more successful than approaches that do not use such

arguments, some of which have been mentioned above.²¹ Thankfully, therefore, if I am right to say that we must live without ambitious transcendental arguments at this level, and be content to settle for modesty, I think it is also right to say that we can do so without conceding defeat to the sceptic, of either the more modest or the more radical kind.²²

REFERENCES

- Alston, William P. (1993), *The Reliability of Sense Perception* (Ithaca and London: Cornell University Press).
- Broad, C. D. (1978), *Kant: An Introduction* (Cambridge: Cambridge University Press).
- Dretske, Fred (1970), 'Epistemic Operators', *Journal of Philosophy* 67, pp. 1007-23
- Gendler, Tamar Szabó and John Hawthorne (eds.) (2002), *Conceivability and Possibility* (Oxford: Oxford University Press).
- Glock, Hans-Johann (2003), 'Strawson and Analytic Kantianism', in Hans-Johann Glock (ed), *Strawson and Kant* (Oxford: Oxford University Press), pp. 15-42.
- Grundmann, Thomas and Catrin Misselhorn (2003), 'Transcendental Arguments and Realism', in Hans-Johann Glock (ed.), *Strawson and Kant* (Oxford: Oxford University Press), pp. 205-218.
- Hume, David (1978), *A Treatise of Human Nature*, ed. L. A. Selby-Bigge, 2nd edn., rev. P. H. Nidditch (Oxford: Oxford University Press).
- Körner, Stephan (1967), 'The Impossibility of Transcendental Deductions', *Monist* 61, pp. 317-31.

²¹ Of these, the one I would favour would be of a quasi-pragmatist kind, hinted at above: The sceptic cannot give us any positive reason to think that the sceptical scenarios actually hold, and as such cannot use them to generate any real doubt, of the sort that might provide a genuine threat to our beliefs about the world.

²² I would like to thank Alex Burri and Christian Beyer for inviting me to the 'Philosophical Knowledge' conference, for which this paper was originally written; and I am grateful to participants at the conference for comments on that occasion, and particularly to Ernest Sosa for correspondence on it afterwards. I am also grateful for comments from Paul Faulkner, Bob Hale and Christopher Hookway, and to participants at the departmental seminar of the University of York, where a previous version was also delivered.

- Nozick, Robert (1981), *Philosophical Explanations* (Oxford: Oxford University Press).
- Putnam, Hilary (1981), *Reason, Truth and History* (Cambridge: Cambridge University Press).
- Reid, Thomas (1863), *Essays on the Intellectual Powers of Man*, in William Hamilton (ed), *The Works of Thomas Reid*, 2 vols, 6th edn (Edinburgh: Maclachlan and Stewart), vol 1, pp. 215-508.
- Sosa, Ernest, 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13 (1999), pp. 141-53.
- Stern, Robert (ed) (1999), *Transcendental Arguments: Problems and Prospects* (Oxford: Oxford University Press).
- Stern, Robert (2000), *Transcendental Arguments and Scepticism* (Oxford: Oxford University Press).
- Stern, Robert (2003), 'On Strawson's Naturalistic Turn', in Hans-Johann Glock (ed.), *Strawson and Kant* (Oxford: Oxford University Press), pp. 219-234.
- Stroud, Barry (2000), *Understanding Human Knowledge* (Oxford: Oxford University Press, 2000).