BOGHOSSIAN ON INFERENTIAL KNOWLEDGE[*]

José L. Zalabardo
University College London

1.   *Content and Justification* (Boghossian 2008) collects fourteen essays by Paul Boghossian, published over the last twenty years or so. These essays contain seminal contributions to a very wide array of philosophical issues, including the theory of content, realism, semantic normativity, self-knowledge, rule-following, analyticity, a priori knowledge and the nature of colour. Perhaps the magnitude of the intellectual achievement that these essays represent can be best appreciated by those of us who have spent a good part of the last two decades trying to come to terms with the consequences of Boghossian's arguments.

   Arguments: that's, above all, what these essays offer. There can't be many works of philosophy with a higher argument-per-page ratio. And their quality matches the quantity. They are consistently good arguments for important conclusions. Many strike me as conclusive, including some that seem to refute views that I find very appealing. There wouldn't be much point in summarising these here, as I cannot hope to improve on the clarity of Boghossian's writing. Instead I'd like to discuss a couple of rare cases where Boghossian's reasoning seems to me to contain some gaps.

2.   I want to concentrate on Boghossian's views on inferential knowledge. Some of our knowledge is best described as inferential. *S* knows *p* inferentially when *S*'s belief that *p* owes its epistemic status to its connection to other items of knowledge, in a way that makes it natural to say that *S*'s belief that *p* inherits the status of knowledge from those other beliefs. Assuming that *S* knows $q_1,..., q_n$, how would her belief that *p* have to be connected to her beliefs in $q_1,..., q_n$ in order for the former to obtain the status of knowledge from this link?

One condition is fairly uncontroversial. In order for S to obtain knowledge of p in this way, the truth of $q_1, ..., q_n$ has to make the truth of q sufficiently likely. The level of likelihood required, and the meaning of 'likelihood' in this context are open to debate. But on one privileged kind of case there is widespread agreement. If $q_1, ..., q_n$ *entail p*, then this condition for inferential knowledge of p is satisfied.

If we restrict ourselves, then, to deductive cases, we have two necessary conditions for S to have (deductive) inferential knowledge of p based on $q_1, ..., q_n$: S needs to know $q_1, ..., q_n$ and $q_1, ..., q_n$ have to entail p. Clearly, in addition, S will need to register in some way the link between $q_1, ..., q_n$ and p. A natural way to achieve this is an act of inference—S inferring p from $q_1, ..., q_n$. So let's suppose that S knows $q_1, ..., q_n$, that $q_1, ..., q_n$ entail p and that S infers p from $q_1, ..., q_n$. Are these conditions jointly sufficient for S to have inferential knowledge of p? The view that they are, for which Boghossian uses the label *Simple Inferential Externalism*, is expressed with the following principle:

(SIE) S's deductive inference of p from $q_1, ..., q_n$ will enable her to know p if and only if:

  (a) S knows $q_1, ..., q_n$ (independently of knowing p) and

  (b) $q_1, ..., q_n$ entail p.

Boghossian argues persuasively that (SIE) is unacceptable. Suppose S knows the following about four numbers, e.g. the solutions to four equations:

  x, y, z and n are whole numbers and n is greater than 2

And suppose that she infers from this the following conclusion:

  $x^n + y^n$ is not equal to z^n

We know now that the premise entails the conclusion. However, it would be wrong to suggest that anyone who drew this inference knowing the premise would acquire as a result

knowledge of the conclusion "whether or not they knew anything about Andrew Wiles' proof of Fermat's last theorem" (257).[1] In order for inferential knowledge to result, some additional conditions would need to be satisfied.

One natural line at this point is to introduce an additional necessary condition for inferential knowledge to the effect that S needs to know that the premises entail the conclusion. Boghossian uses the label *Simple Inferential Internalism* for the resulting view:

(SII) S's deductive inference of p from $q_1, ..., q_n$ will enable her to know p if and only if:

   (a) S knows $q_1, ..., q_n$ (independently of knowing p),

   (b) $q_1, ..., q_n$ entail p and

   (c) S knows that $q_1, ..., q_n$ entail p

Boghossian's recent work contains a battery of important arguments against this position.

3.   In "How are Objective Epistemic Reasons Possible", Boghossian starts his attack on (SII) with an argument to which he doesn't devote much attention. It is based on the contention that condition (c) is "far too sophisticated a requirement" (255). Children who can obtain knowledge by deductive inference cannot be expected to have the beliefs that (c) calls for.

This point doesn't seem to me to pose an insurmountable obstacle for the advocate if (SII). The objection will succeed only if there are cases in which it is correct to describe a child, or some other intellectually unsophisticated subject, as inclined to infer q from p, but not as believing that p entails q. It seems to me that an advocate of (SII) could plausibly contend that there are no cases of this kind. A subject's inclination to infer q from p should normally be taken as sufficient grounds for ascribing to him a belief in the existence of at least a probabilistic link between the truth of p and the truth of q. And if this inclination to infer is

---

[1] Numbers in brackets are page numbers in (Boghossian 2008).

found to be unwavering, impervious to any changes in background beliefs, it would provide

adequate grounds for ascribing the belief that p entails q.[2] Needless to say, we shouldn't

expect a child to be able to entertain in consciousness the proposition that p entails q, but on

any plausible account of belief this conscious episode is not a necessary condition for belief.[3]

4. I want to turn next to a line of reasoning against (SII) that Boghossian discusses in much

more detail. He presents it as an application of Lewis Carroll's argument in "What the

Tortoise Said to Achilles" (Carroll 1895). Boghossian has offered three different versions of

the argument.

The first version appears in "How Are Objective Epistemic Reasons Possible".

Boghossian concentrates on an inference from the particular proposition p and the particular

proposition *if p then q* to the proposition q. Suppose then that *S* infers q from p and *if p then

q*. In order for condition (c) to be satisfied, the subject will have to know this:

(N) p and *if p then q* entail q

Boghossian argues that this demand leads to an infinite regress. The problem arises when we

pose the following question:

> […] how does [knowledge of (N)] help justify the thinker in drawing the
> conclusion q from the premises with which he began? (256)

He goes on to argue that in order for knowledge of (N) to play this role, the subject will have

to perform an inference from (N), p and *if p then q* to q. But if knowledge of q is to result

from this inference, it will also have to satisfy (SII). I.e. the subject will need to know this:

---

[2] Ascribing a belief in a probabilistic link on the basis of inferential inclinations might seem
more natural that ascribing a belief in an entailment on this basis. An advocate of (SII) who
wants to take this point on board could weaken condition (c) by requiring only knowledge of
the existence of a probabilistic link.

[3] Another possibility at this point, explored in my (Zalabardo 2005), is to formulate (SII) not
in terms of knowledge, but in terms of warrant, conceived as the property that turns true
beliefs into knowledge. Having warrant for p doesn't in principle require believing p.

(N')   *p*, *if p then q* and (N) entail *q*.

But with respect to (N') we can ask the same question that we asked about (N). This will call

for another inference, and then (SII) will require a further item of knowledge. Clearly we are

off on an infinite regress.[4]

   If (SII) launches this regress, we seem to have no option but to join Boghossian in

rejecting (SII). But does (SII) have this unfortunate consequence? The regress is supposed to

arise when we ask the question: 'How does knowing (N) help justify the thinker in

concluding *q* from *p* and *if p then q*?' On the face of it, the advocate of (SII) would seem to

be entitled to a very simple answer to this question: 'Knowing (N) helps justify the thinker in

drawing the conclusion by satisfying a necessary condition for inferential knowledge

(/justification)'. According to this response, in order for knowledge of (N) to have this effect,

none of the further inferences involved in the regress would have to be drawn. So long as the

thinker knows (N), (SII) will be satisfied.[5]

   I think that this answer to Boghossian's question is perfectly legitimate, and hence that

(SII) does not launch the regress that Boghossian describes. It will help to appreciate the

legitimacy of this response if we reflect on the basic nature of the analytic task that an

account such as (SIE) or (SII) is supposed to discharge. People often draw inferences from

premises they believe to other propositions that they already believe or that they come to

believe as a result of the inference. In some, but not all, of these cases, we want to say that the

inference bestows the status of knowledge on the subject's belief in the conclusion. The

analytic question concerns the circumstances under which this epistemic assessment is

---

[4] Boghossian adds that the inference from (N), *p* and *if p then q* would have to invoke modus
ponens, just as the original inference from *p* and *if p then q*, but if my construal of the
argument is along the right lines, this observation is irrelevant at this point.

[5] Crispin Wright has discussed a response along these lines to Boghossian's argument. See
(Wright 2001: 41-85).

correct. We look for necessary conditions for the inference to have this power—conditions in whose absence the inference would not produce knowledge—and for sufficient conditions for this—conditions that guarantee that knowledge results from the inference.

(SII) has to be understood as a contribution to this enterprise. Its advocates maintain that when someone draws a deductive inference without knowing that the premises entail the conclusion, the inference does not enable him to know the conclusion. In the case of the inference from *p* and *if p then q* to *q*, their claim is that if someone draws this inference without knowing (N), the inference won't produce knowledge of *q*. Now, Boghossian's Carroll-inspired argument against (SII) doesn't challenge the knowability of (N). What the argument purports to show is that, even if we grant that the subject knows (N), by reflecting on how knowledge of (N) would help to produce inferential knowledge of *q* we discover that knowledge of (N) wouldn't have this effect unless further impossibly regressive conditions were satisfied. I am suggesting that the proponent of (SII) shouldn't accept this. She should insist that someone who knows (N), but hasn't drawn any of the additional inferences that figure in Boghossian's regress, will have satisfied (SII). (SII) doesn't entail that the impossible infinite chain of inferences and items of knowledge is a necessary condition for knowledge. This argument doesn't show that there's anything wrong with treating (SII) as a necessary condition for inferential knowledge.

Sometimes Boghossian gives the impression that the problem has to do, not with condition (c) of (SII), but with another necessary condition for inferential knowledge that he discusses in "How Are Objective Epistemic Reasons Possible":

(D)   In order for *S* to have deductive inferential knowledge of *p* based on $q_1, \ldots, q_n$, *S* needs to infer *p* from $q_1, \ldots, q_n$ because of his belief that $q_1, \ldots, q_n$ entail *p*.[6]

In the case of *S*'s inference form *p* and *if p then q* to *q*, this is the claim that in order for the inference to produce knowledge of *q*, *S* needs to draw the inference *because* of his belief in (N). The thought is that *S*'s belief in (N) should play a role in the inference if the inference is to produce knowledge of *q*, and one could argue that this involvement in the original inference would require a further inference in which (N) figures as a premise, thus launching the infinite regress.

Could the regress be reinstated through this route? I want to suggest that an advocate of (D) can easily reject this outcome, so long as she understands (D) as the demand that belief in (N) should play a *causal* role in the inference. When the condition is understood in this way, it has no tendency to generate an infinite regress. In any case, the issue that I am pursuing here is whether (SII) is threatened by Boghossian's arguments, and even if (D) had to be rejected, there is no reason to think that (SII) would fall with it.

5.   Boghossian presents two further versions of the Carroll-inspired regress argument in "Blind Reasoning". They target the suggestion that the knowledge called for by (SII) could arise from a faculty for rational insight.

The first of these versions of the argument is based on the following thought:

> For obvious reasons, it's not plausible to think of this capacity for rational insight as operating on individual inferences one by one, generating for each of them the insight that if its premises are true, then so is its conclusion. Rather, we suppose that rational insight equips the thinker to arrive at the wholly general insight that MPP is valid […] (273).

---

[6] Boghossian presents the account of inferential knowledge that he discusses in "How Are Objective Epistemic Reasons Possible" as a version, with small modifications, of an account considered by Alston (254). However, Boghossian's condition (D) is largely unrelated to Alston's correlate. See (Alston 1986: 11).

According to this thought, knowledge of (N) would have to be explained as arising from knowledge of the following principle:

(MPV) For all propositions ϕ, ψ, ϕ and *if ϕ then ψ* entail ψ.

But now a problem arises in the transition from knowledge of (MPV) to knowledge of (N):

> Rational insight, we are conceding, gets us as far as the general propositional knowledge that all arguments of MPP form are valid. However, to bring this knowledge to bear on the justifiability of any particular inference will require the thinker justifiably to infer the validity of that particular inference from the validity of all arguments of MPP form. And this will require him to be able to reason according to MPP justifiably. (273-74)

This passage makes two important points. The first is that knowledge of (N) will have to be construed as inferential, arising from an inference with (MPV) as a premise. The second is that this inference would have to invoke MPP as a rule of inference. Boghossian goes on to argue that the second of these points generates a fatal circularity, arising from the fact that MPP is the very same rule of inference that needs to be invoked in *S*'s original inference from *p* and *if p then q* to *q*.

I agree with Boghossian that, subject to his assumptions about the universal nature of rational insight and the ensuing inferential character of knowledge of (N), (SII) renders inferential knowledge impossible. It seems to me, however, that the difficulty can be brought out more simply. A problem arises even before we consider which premises (N) would have to be inferred from, or which rules would have to be invoked to underwrite the inference. In order to show that (SII) renders inferential knowledge impossible, it will suffice to assume that knowledge concerning entailment has to be inferential:

(NI)   *S* can know that $q_1,\ldots, q_n$ entail *p* only as a result of a deductive inference

It is easy to see how (SII) and (NI) generate an infinite regress from the assumption that *S* knows a proposition *q* as a result of a deductive inference:

(1)   *S* knows *q* as a result of a deductive inference I (Assumption)

(2)   *S* knows that the premises of I entail *q* (from 1, by (SII))

(3)   *S* knows that the premises of I entail *q* as a result of a deductive inference I' (from 2, by (NI))

(4)   *S* knows that the premises of I' entail its conclusion (i.e. that the premises of I entail *q*) (from 3, by (SII))

(5)   *S* knows that the premises of I' entail its conclusion as a result of a deductive inference I'' (from 4, by (NI))

(6)   *S* knows that the premises of I'' entail its conclusion (from 5, by (SII))

        …

It is plain that this regress would render deductive inferential knowledge impossible.

   Clearly, the threat that this regress poses for (SII) is conditional on the acceptance of (NI). If the advocate of (SII) could reject (NI), she would be able to sidestep the difficulty. Notice that this is also the case for Boghossian's vicious-circle argument, since (NI) is clearly one of its premises. I want to suggest that knowledge of the validity of particular arguments can be plausibly construed as non-inferential, and hence that (SII) doesn't fall prey either to Boghossian's vicious circle or to my infinite regress.

   My proposal is that the advocate of (SII) should construe knowledge of entailment along the same lines as the knowledge that we obtain from unproblematic cases of feature or pattern recognition. We have the ability to recognize colours, shapes, human faces, voices, accents, melodies, chord progressions, pictorial styles, grape varieties, grammatical sentences, friendly strangers, dangerous situations … Clearly not all of us have all of these abilities, and

those who have them don't have them to the same extent. Furthermore, we are not infallible about our possession of these abilities: someone might be very good at recognizing one of these items without realising that she is, or be very bad at it, although she thinks that she is very good. Nevertheless, there are many cases in which a subject has the ability to recognise one of these items with an arbitrarily high level of success, and in these cases it is perfectly natural to say that the ability is a source of knowledge. My ability to recognize faces enables me to know that the person across the road is my wife, and my ability to recognize grape varieties enables me to know that the wine I'm drinking tastes of cabernet sauvignon.

Now, it is an uncontroversial empirical fact that these recognitional capacities are not cognitively fundamental. Our recognition of a complex taste has to be achieved in terms of the input-processing make-up of our gustatory devices, and we are not endowed with cabernet sauvignon-detecting taste buds. It follows that from a physiological point of view my recognition of cabernet sauvignon has to be explained as the recognition of a combination of a few basic tastes. However, it would be implausible to suggest that my knowledge that the wine I'm drinking tastes of cabernet sauvignon is inferential knowledge, using beliefs about more basic tastes as premises. From an epistemological point of view, my recognitional knowledge that the wine tastes of cabernet sauvignon is fundamental.

I want to suggest that this model can be naturally applied to knowledge of the validity of specific arguments. All of us, even the logically illiterate, have the ability to recognise valid arguments, to a greater or lesser extent. Some of us are very good at it. At least in the simplest cases, and outside our logic class, judgments of validity can be perfectly immediate. Those with logical training will sometimes be able to identify the inferential rules that explain the validity of an argument, and will be able to use this knowledge to acquire inferential knowledge of the validity of certain arguments, with the universal rules of inference acting as premises. However, in the multitude of everyday cases in which we recognize valid

arguments without formulating the underlying rules of inference, sometimes even without being capable of formulating them, it is natural to say that our ability to recognize valid arguments gives us knowledge, and that this knowledge is non-inferential. And this would be so even if the sub-personal cognitive devices that sustain this ability operate by recognising certain universal inference patterns.[7]

If, as I am suggesting, this account of our knowledge of logical validity is along the right lines, it follows that we often know non-inferentially that the premises of a particular argument entail its conclusion. Hence (NI) has to be rejected and (SII) is not threatened either by Boghossian's circle or by my regress.

6.   I turn now to the other version of the Carroll-inspired argument that Boghossian presents in "Blind Reasoning". Here Boghossian concedes for the sake of the argument that we can have non-inferential knowledge of the validity of particular arguments, but argues that the difficulty survives this concession. Here, as in the version of the argument in "How Are Objective Epistemic Reasons Possible", the problem arises when we ask the question, how does my knowledge of the validity of an argument bear on my warrant to infer its conclusion, e.g., how does my knowledge of (N) help justify me in inferring q from p and *if p then q*? However, in "Blind Reasoning" the question is taken to concern the bearing of knowledge of (N) on the epistemic status, not of my belief in q, but of my second-order, epistemic belief that I am justified in inferring q. Boghossian expresses some reluctance to running the argument in these terms, but he defends the move as raising "fewer distracting objections" (275).

---

[7] It is interesting to reflect in this connection that the fact that there is a recursive set of rules of inference that can explain all instances of first-order logical validity, far from being a platitude, was only established with Gödel's proof of the completeness of a first-order logical calculus.

An important point to bear in mind is that even if this version of the argument succeeded in showing that (SII) makes it impossible to know that an argument gives me inferential knowledge, it wouldn't follow directly that (SII) makes inferential knowledge impossible, so long as we are prepared to accept the plausible thought that it is possible to have knowledge that you don't know you have. Boghossian doesn't discuss the issue explicitly, but his argument doesn't rely on the rejection of this point. If (SII) had the consequence that I can't know that an inference gives me knowledge, we would have compelling grounds to reject the principle, even if it didn't threaten our ability to obtain knowledge from inferences.

And it seems that the shift of target does improve the prospects of the argument. When I discussed the version of the argument in "How Are Objective Epistemic Reasons Possible", I suggested that (N) should not be treated as a premise in an argument for *q*, but simply as a proposition that the subject needs to know in order to satisfy a necessary condition for inferential knowledge postulated by (SII). In the new setting, by contrast, the status of (N) as a premise seems inescapable.

What we are concerned with now is the epistemic status of my belief in the following proposition:

(K)   My inference from *p* and *if p then q* gives me knowledge of *q*.

It is hard to see how my knowledge of this proposition could fail to be inferential, and how (N) could fail to figure as a premise in the requisite inference. Knowledge of (K) would result, as Boghossian suggests, from an inference along the following lines:

i.   *p* and *if p then q* entail *q* (N)

ii.   If *p* and *if p then q* entail *q*, then anyone who knows *p* and *if p then q* and knows that *p* and *if p then q* entail *q* will obtain knowledge of *q* from an inference from *p* and *if p then q*.

iii.   Anyone who knows *p*, *if p then q* and that *p* and *if p then q* entail *q* will obtain

knowledge of *q* from an inference from *p* and *if p then q*.

iv.   I know *p*, *if p then q* and that *p* and *if p then q* entail *q*.

Therefore:

(K)   My inference from *p* and *if p then q* gives me knowledge of *q*.

Clearly, this argument needs to invoke MPP as a rule of inference, for example in the transition from i and ii to iii. And this circumstance, Boghossian suggests, reinstates the threat of circularity (275). How is the vicious circle supposed to arise? Presumably the problem is that (K) asserts the knowledge-producing power of an MPP inference, but since knowledge of (K) is supposed to arise from an MPP inference, its knowledge-producing power will have to be presupposed.

The first point to notice is that if there is a problem here it's not clear how it's supposed to bear on (SII). That (K) can only be known as a result of an inference along these lines is not a consequence of (SII) that we can reject if we abandon (SII). That (K) can only be known in this way seems to me independently plausible. If there is a problem with the power of this inference to produce knowledge of (K), it's not a problem that can be solved by abandoning (SII).[8]

In any case, I don't think there is a real problem with the inference. It is important to bear in mind that at this stage in the dialectic Boghossian is conceding to his opponent that we can know the validity of specific MPP inferences individually, without an inference from the universal claim (MPV) that all MPP inferences are valid. Consider now the proposition that

---

[8] Needless to say, abandoning (SII) would force us to modify premise ii, but if there is a problem with the original argument it won't be removed by this modification.

expresses the knowledge-producing power of the inference from which knowledge of (K) is expected to result:

(K') My inference from i and ii gives me knowledge of (K).

The first point to notice is that K and K' are different propositions, since each of them ascribes knowledge-producing power to a different inference. The argument from i and ii to (K) does not produce a vicious circle, since what is presupposed is different from what is concluded.

One could argue on Boghossian's behalf that the only effect of this point is to move us from a vicious circle to an infinite regress. If knowledge of (K) presupposes knowledge of (K'), we need to ask how knowledge of (K') could be obtained. Presumably this would require an inference, whose knowledge-producing power would have to be presupposed, etc.

However, a second important point invalidates this line of reasoning. Knowledge of (K) does not presuppose knowledge of (K'). All that's presupposed by knowledge of (K) is the *truth* of (K'). So long as (K') is true, whether I know it or not, the inference from i and ii will give me knowledge of (K). The impression that a regress is generated is entirely spurious.

I conclude that the third version of the Carroll-inspired argument is no more successful than its predecessors in undermining (SII).

7. In "How Are Objective Epistemic Reasons Possible", Boghossian presents his attack on (SII) as a vindication of the knowledge-producing power of rule-circular arguments. An argument for the conclusion that a rule of inference R is valid is rule-circular when R is among the rules of inference employed by the argument. In premise-circular arguments, by contrast, the conclusion is among the premises. Boghossian sees (SII) as a major obstacle to the acceptance of rule-circular arguments as sources of knowledge:

> If C [essentially, condition (c) of (SII)] were a correct necessary condition on warrant transfer, then it would follow immediately that there could be no such thing as rule-circular justification. For C requires that, in order to use an argument employing a given rule to support the claim that the rule is truth-preserving, one already has to know that that rule is truth-preserving. And that would make the rule-circular justification otiose: the knowledge arrived at would already be presupposed. (255)

As we have seen this is not, strictly speaking, correct. (SII) does not require that we know that the rules employed in a knowledge-producing inference are known to be truth-preserving. All it requires is that the particular inference is known to be truth-preserving—that its premises are known to entail its conclusion. As far as (SII) is concerned, I could come to know the universal validity of MPP with an inference A that employs MPP. (SII) would require that I know A to be valid, but if this knowledge is in place, and other conditions for inferential knowledge are satisfied, A would enable me to obtain knowledge of the universal validity of MPP.

Hence accepting (SII) is no obstacle to joining Boghossian in ascribing to some rule-circular arguments the power to produce knowledge. On the contrary, (SII) would seem to force us to sanction some rule-circular arguments, since, as we have seen, an inference based on a rule-circular argument can in principle satisfy (SII). If we wanted avoid treating rule-circular arguments as sources of knowledge, we would have to replace (SII) with a stronger account of inferential knowledge—one that ruled out some of the inferences that (SII) condones.

8. Boghossian defends rule-circular arguments as offering our only chance of explaining our knowledge of the validity of rules of inference. Is he right about this? He argues for this claim in two stages. On the one hand, he argues that it wouldn't be possible to have inferential knowledge of the validity of rules of inference unless rule-circular arguments had the power to produce knowledge. On the other hand, he argues that it is not possible to have non-inferential knowledge of the validity of rules of inference. The first point strikes me as

incontestable. Every deductive argument must employ rules of inference, and this goes, in particular, for arguments for the validity of rules of inference. So if we take the inferential route, rule-circularity, or some closely related phenomenon, will be ultimately unavoidable.

I want to argue that the second point is less compelling, that it is possible to find a plausible non-inferential account of how we know that rules of inference are valid. I agree with Boghossian that the forms of non-inferential knowledge that he considers cannot be invoked to explain knowledge of the validity of rules of inference. Let's concentrate on the claim that every instance of MPP is truth-preserving, i.e. (MPV), above. As Boghossian argues, knowledge of (MPV) can't be construed as observational, or as arising from the fact that belief in (MPV) is self-fulfilling, self-evident, necessarily presupposed by any justification of (MPV) that the thinker might give or a necessary condition for possessing the concept of conditional (239-43). How can (MPV) then be known?

I want to suggest that knowledge of (MPV) can be explained in terms of the same recognitional capacity that I used to explain knowledge of its substitution instances. This capacity can be applied, not only to particular inferences, but also to types of inference, and it can produce a belief in (MPV) as well as a belief in (N). When this recognitional capacity is sufficiently reliable, and other conditions on non-inferential knowledge are satisfied, we say that it is a source of knowledge—that the true (particular or universal) beliefs that we form as a result of its operation have the status of knowledge.

Notice that this ability to operate at different levels of abstraction is uncontroversially present in other recognitional capacities. The very same ability that enables any English speaker to recognise the sentence *Mary arrived the house* as ungrammatical will enable speakers with the requisite concepts to acknowledge the truth of the proposition that the verb *to arrive* never takes a direct object. In this way, the capacity to recognise grammatical sentences can be a source of knowledge, not only of particular propositions, but also of

universal generalisations. I am suggesting that the same goes for our ability to recognise valid arguments. It can be used to recognise the validity of specific instances of MPP as well as the truth of (MPV).

If our knowledge of the validity of rules of inference can be explained along these lines, we have at our disposal an alternative to the account in terms of rule-circular arguments that Boghossian defends. This outcome has important consequences for the plausibility of Boghossian's account. The main reason that he offers for over-riding our intuitive suspicion of rule-circular arguments is that we have no other explanation of how the validity of rules of inference can be known. If rule-circular arguments are not needed for this role, we have no compelling reason for ascribing to them an epistemic status that they intuitively lack.

9. Boghossian thinks that an adequate account of inferential knowledge would have to be more demanding than (SIE) but less demanding than (SII). He puts forward an account of inferential knowledge that occupies this middle ground. It can be formulated in the following terms:

(BI) S's deductive inference of p from $q_1,..., q_n$ will enable her to know p if and only if:

   (a) S knows $q_1,..., q_n$ (independently of knowing p),

   (b) $q_1,..., q_n$ entail p and

   (c) EITHER S knows that $q_1,..., q_n$ entail p OR the inference from $q_1,..., q_n$ to p instantiates a *meaning-constituting* rule of inference.

A rule of inference R is meaning constituting when inferring according to R is a condition for having one of the concepts ingredient in it (279). And Boghossian's examples suggest that the standard introduction and elimination rules for the connectives and other logical operators should be regarded as having this status. Notice that (BI) will confer on some rule-circular arguments the power to produce knowledge of their conclusions. If R is a meaning-

constituting rule, my belief that all instances of R are valid will receive the status of knowledge from a valid rule-circular inference whose premises I know. I want to close with a few remarks about Boghossian's account of inferential knowledge.

Boghossian declares that his account intends to capture "a broadly *internalist* notion of warrant" (259). He claims that his account of inferential knowledge "*does* satisfy the constraints appropriate to an internalist notion", once the contrast between internalism and externalism is properly conceived. He argues that the contrast turns on whether or not epistemic responsibility or blamelessness are treated as necessary conditions for epistemic justification. For the internalist, a belief can't be justified unless it exhibits these features, whereas the externalist rejects this constraint. And Boghossian thinks that internalism, in this sense, is correct—that "[…] being justified is, at least in part, a matter of being epistemically blameless" (270). And again: "It appears to be a condition on someone's being epistemically justified that they *not be epistemically irresponsible* in forming their belief." (259)

Boghossian takes issue with those who assume that the possession of a reflectively appreciable item of information that justifies the belief is the only way of satisfying this constraint. He argues that the constraint is also satisfied in the absence of reflectively appreciable warrant by a belief that has been formed with an inference underwritten by a meaning-constituting rule. The thought is that inferring according to, say, MPP, cannot be regardless as blameworthy or irresponsible if an inclination to infer in this way is a necessary condition for having the concept of the conditional.

I am not going to discuss whether internalism is best construed in these terms. I want to consider instead whether it is legitimate to treat responsibility and blamelessness as central ingredient of the notion of justification and whether inferences underwritten by meaning-constituting rules can be said to exhibit these features.

10. What do blamelessness and responsibility have to do with justification? The first point we need to notice is that, although Boghossian's discussion at this stage is framed mainly in terms of the notion of justification, knowledge should be his ultimate concern. (BI) purports to specify necessary and sufficient conditions for inferential knowledge. Hence the question, whether responsibility and blamelessness are connected with justification is only of interest in this context to the extent that justification is a central ingredient in knowledge. If we understand Boghossian's goals along these lines, we are entitled to rephrase the question by asking what role responsibility and blamelessness play, not in justification, but in knowledge.

I want to suggest that the connection is much weaker than Boghossian assumes. On the one hand, responsibility and blamelessness don't take a true belief anywhere near the status of knowledge. Consider, in this connection, the situation of an envatted brain, assuming, *pace* Putnam, that he is capable of forming beliefs about his physical environment. An envatted brain may well conduct his cognitive life in a perfectly responsible and blameless manner. And yet, if he occasionally forms a true belief (e.g., *there is a chair in this room*), we have no inclination whatsoever to ascribe to this responsible, blameless true belief the status of knowledge.

On the other hand, it can be argued that blamelessness and responsibility are not necessary conditions for knowledge—that we can acquire knowledge in ways that can be described as irresponsible or blameworthy. Suppose, for example, that I have compelling evidence, provided by my parents, doctors, etc. that my childhood memories of a trip to the zoo are fabricated. I seem to remember seeing all the animals, but all the evidence at my disposal indicates that these seeming memories are not to be trusted. Suppose, however, that this evidence is the result of a conspiracy to deceive me—that I really went to the zoo and my memories are perfectly real. Suppose now that I seem to remember eating an ice cream at the zoo that day, and that I actually ate an ice cream and I have a genuine memory of the event.

If, in spite of all the evidence, I can't refrain myself from believing that I ate an ice cream at the zoo that day, my belief can be naturally described as blameworthy and irresponsible. Even I would describe it in these terms, as I am convinced that I shouldn't be taking any notice of these fabricated memories. Nevertheless, there is a strong intuitive case for describing this irresponsible, blameworthy true belief as knowledge. If we are prepared to contemplate, as I think we should, cases in which knowledge is acquired in this kind of situation, we will have to say that responsibility and blamelessness are not necessary conditions for knowledge.[9]

I conclude that the question, whether a belief was formed responsibly and blamelessly, is largely tangential to the question, whether the belief has the status of knowledge.

11. I want to turn now to the issue, whether being underwritten by a meaning-constituting rule of inference can be said to secure responsibility and blamelessness. It is important to appreciate in this connection that, at least for the purposes of Boghossian's argument, blamelessness and responsibility will need to be assessed from the thinker's own conception of his situation. If they were assessed, instead, from an objective point of view, then beliefs that satisfied an externalist criterion for knowledge would have to count as responsible and blameless, since forming those beliefs is what the subject would need to do in order to attain the objective goal of truth.[10]

---

[9] Alvin Plantinga has argued for this conclusion. See (Plantinga 1993: 45). One of the roles that Laurence BonJour's clairvoyance thought experiments are supposed to play is to lend support to the claim that epistemic rationality and responsibility are necessary conditions for knowledge, but I have argued elsewhere (Zalabardo 2006: 157-60) that BonJour's argument fails to achieve this goal.

[10] This point is accepted by Laurence BonJour, who is presented by Boghossian as the main source of his ideas in this area. BonJour tells us that the notion of epistemic rationality and responsibility that his anti-externalist argument needs to appeal to is the notion "[…] of such rationality as essentially dependent on the believer's own subjective conception of his epistemic situation". (BonJour 1985: 49-50)

However, once we accept that this is the point of view from which the responsibility and blamelessness of a belief would need to be assessed, it is hard to see why the meaning-constituting character of a rule of inference should make a difference to whether the beliefs that are formed with it are responsible or blameless. The crucial point here is that subjects don't typically have at their disposal information as to which of the rules of inference that they employ are meaning constituting. Assume that modus ponens is a meaning-constituting rule, but modus tollens isn't, and consider a subject who is inclined to infer according to both rules, but has no idea that one, but not the other, is a meaning-constituting rule. Now, it follows from Boghossian's views that when the subject infers according to modus ponens, her inferences will be responsible and blameless, even in the absence of 'reflectively appreciable warrant'. However, when she infers according to modus tollens in the absence of such warrant, her inferences will have to count as irresponsible and blameworthy.

This strikes me as an unacceptable outcome. From the subject conception of the situation there is no difference between the status of the two rules. Hence, if responsibility and blamelessness are assessed from this subjective point of view, inferences involving each of these rules should not attract different assessments. One might try to respond to this point by using the subject's beliefs as to which rules are meaning-constituting to determine when they can produce responsible and blameless inferences in the absence of reflectively appreciable warrant. However, the problem can't be solved in this way. Many inferrers have no beliefs at all as to which of the rules of inference that they use are meaning-constituting. Furthermore, the move would introduce into epistemic assessment the kind of subject-relativity that Boghossian is admirably anxious to avoid.

12. The second point that I want to make concerning meaning-constituting rules presents a more radical challenge to Boghossian's strategy. A meaning-constituting rule with respect to, say, the conditional is a rule with the following status: no one can count as having the concept

of the conditional unless they have the inclination to infer according to the rule. I want to suggest that there might not be any rules with this status. Clearly, not all valid rules of inference involving the conditional have this status. Someone who has no inclination to reason according to the pattern, say, 'from $(\phi \supset \psi) \supset \phi$ infer ϕ' might still count as meaning conditional by '$\supset$', provided that we have a plausible explanation of her failure to infer in this way and that her inclinations regarding '$\supset$'-involving inferences are otherwise sufficiently standard. Meaning-constituting rules are rules that don't behave like this. If a subject doesn't infer according to a rule of inference involving '$\supset$' that is meaning constituting for the conditional, then even if we had a very good explanation of why the subject doesn't infer in this way, and even if her inclinations regarding '$\supset$'-involving inferences were otherwise perfectly normal, it would be wrong to say that she means conditional by '$\supset$'.

I want to suggest that our meaning-ascription practices might never confer this status on a rule of inference. It might be, that is, that every rule is such that a subject's failure to infer according to it might be compensated for by a sufficiently good explanation of this failure and otherwise sufficiently standard inferential patterns.[11] Suppose, for example, that failure to infer consistently according to modus ponens were a well documented cognitive deficit, present in subjects whose inferential behaviour regarding the conditional is otherwise perfectly normal. I think it would be possible to fill in the details of the case in such a way that it seems entirely natural to say that these subjects have the concept of the conditional. If we were prepared to accept that there might be circumstances under which this would be the right verdict, we would have to conclude that modus ponens is not a meaning-constituting rule for the conditional. And if we accepted this possibility for every conditional-involving rule, we would have to conclude that none of them are meaning constituting. If this result

---

[11] Notice that this point is perfectly compatible with the view that the meanings that we ascribe to the connectives are determined exclusively by our inferential patterns involving them

could be extended to all rules of inference, (BI) would collapse into (SII), and Boghossian would have failed to articulate a viable middle ground between (SIE) and (SII).

13. I have identified several respects in which Boghossian's arguments concerning inferential knowledge would need to be clarified or supplemented. As they stand, they don't offer compelling reasons for rejecting any of the following views:

- Simple Inferential Internalism doesn't render inferential knowledge impossible.

- Simple Inferential Internalism does not entail that rule-circular inferences can't produce knowledge.

- We can explain how we know that rules of inference are truth preserving without condoning rule-circular inferences.

- Responsibility and blamelessness do not play a significant role in the concept of knowledge.

- Whether an inference is responsible or blameless is not affected by whether it is underwritten by a meaning-constituting rule of inference.

- There are no meaning-constituting rules of inference.

If these views can be upheld, a certain account of inferential knowledge starts to look much more appealing than Boghossian makes it sound. On this account, deductive inferential knowledge does require that the subject knows that the inference is valid, but knowledge of the validity of specific inferences and of inference rules does not need to be explained inferentially. Boghossian has not shown that there's anything wrong with a position along these lines.

REFERENCES

Alston, William P. 1986. Epistemic Circularity. *Philosophical Studies* 47:1-28.

Boghossian, Paul. 2008. *Content and Justification. Philosophical Papers*. Oxford: Clarendon.

BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, Massachusetts: Harvard University Press.

Carroll, Lewis. 1895. What the Tortoise Said to Achilles. *Mind* 4:278-80.

Plantinga, Alvin. 1993. *Warrant: The Current Debate*. New York and Oxford: Oxford University Press.

Wright, Crispin. 2001. On Basic Logical Knowledge. Reflections on Paul Boghossian's "How Are Objective Epistemic Reasons Possible?". *Philosophical Studies* 106:41-85.

Zalabardo, José L. 2005. Externalism, Skepticism and the Problem of Easy Knowledge. *Philosophical Review* 114:33-61.

———. 2006. BonJour, Externalism and the Regress Problem. *Synthese* 148:135-68.