PHILOSOPHICAL ISSUES, 9
Concepts, 1998

Replies to Commentators

Paul A. Boghossian

1 Reply to Loar

In "What the Externalist Can Know *A Priori*" (hereafter, EAP), I argued that combining an externalism about concepts with a traditional doctrine of privileged access to the contents of one's own thoughts leads to absurd results. For, I argued, such a compatibilist combination would allow us to argue purely *a priori* as follows:

(1) If I have the concept water, then water exists.

(2) I have the concept water.

Therefore,

(3) Water exists.

Strictly speaking, Brian Loar does not dispute this central claim. He says:

> Let "twin–Earth" externalism be a strong externalist theory that entails that "water" expresses different concepts for me and my twin. It seems to me that Boghossian's argument is correct about twin–Earth externalism about concepts. He is quite persuasive that compatibilism is difficult to defend for such an externalist. (p. 215)

This content downloaded from
128.122.149.96 on Sat, 17 Apr 2021 01:23:18 UTC
All use subject to https://about.jstor.org/terms

What Loar wishes to argue, however, is that a compatibilist position may not be so implausible for a weaker form of externalism, one that applies only to certain very basic concepts, such as spatial concepts, concepts of physical objects and their parts and the like.

I am not averse to the idea that there might be *a priori* routes to the truth of certain very basic propositions —for example, that there is an external world, that it contains things that occupy space. To that extent, I have nothing to say against a view, or cluster of views, that issues in *a priori* knowledge of such very basic propositions. The sort of anti-compatibilist argument I offered in EAP works only to the extent that the knowledge yielded by the compatibilism at issue is clearly and manifestly not knowable *a priori* —for example, that water exists. If I am right, then this style of argument works against a compatibilism that seeks to combine self-knowledge with a *generalized* externalism about concepts, the sort of externalism that would be yielded by twin-Earth experiments. If the compatibilism at issue is restricted to cases where the results are not embarrassingly strong —for example, to the sorts of results that have always been envisioned by proponents of transcendental arguments for the existence of an external world— then, to that extent, the argument of EAP has nothing against it.

2 Reply to Yablo

Steve Yablo has a lot of interesting things to say in his comment, but his main counterclaim is that it is none too clear that sentences involving the word 'water' have no determinate truth conditions on Dry Earth. He says:

> One thing is clear: if this be Dry Earth, then the great majority of worlds are not classifiable either as containing water or lacking it. It's the next step that bothers me. Can we really conclude that "*there is* water" is lacking in truth conditions?
>
> That depends. It probably does follow, if truth conditions are seen as singular propositions made up *inter alia* of full-blooded properties; in the absence of water there can be no full-blooded property of *being* water, which is curtains for the proposition.
>
> But the singularist approach is surprising in a context where truth-conditionality is being treated as a condition of thought. After all, the capacity for water-thought is intuitively quite independent of ontological disputes about what sorts properties there may be. (p. 223)

This content downloaded from
128.122.149.96 on Sat, 17 Apr 2021 01:23:18 UTC
All use subject to https://about.jstor.org/terms

If we are not to think of satisfaction conditions and truth conditions in terms of full-blooded properties, how should we think of them?

> The answer is that any object will do that encodes the possibility of interrogating a world on such matters as whether it contains water; that's what it takes to have water-thought, hence that's all that can be asked of truth conditions as a requirement of such thought. It seems to me that the encoding role is most directly and efficiently played by (truth conditions conceived as) rules or recipes for classifying worlds. Singular propositions can serve in some cases as handy repositories of classificatory information... [but] it's the rule that matters... (p. 223)
>
> If we are unlucky enough to be living on Dry Earth, then as Paul says, there is no (full-blooded) property of being water. That doesn't in and of itself make nonsense of the question "which actual and counterfactual stuffs deserve to be described as water?" It doesn't rule out that "there is water" puts a condition on worlds, albeit a condition with less taxonomical power than might ideally be wanted. (And it does have *some* taxonomical power, for it rules negatively on Dry Earth.) (Pp. 224-5)

Now, it is true that in EAP, I did identify the truth condition of a sentence with the singular proposition, if any, that it expresses. But I don't really see that that is essential to my argument. I could equally well put it in terms of *kinds*, for example; and that, I take it, no one could object to: talk about kinds is presumably not controversial in the context of a discussion of the semantics of natural kind words.

So my idea was this. If there is no kind denoted by a given natural kind term —say, 'water'— then no satisfaction conditions for that term will have been pinned down. The actual kind of stuff at the end of the relevant causal chain is supposed to fix the kind of stuff that is denoted by 'water'. If there's nothing there, then, it would seem, there is no fact of the matter what kind of stuff there would have to be for the extension of the term not to be empty.

Steve wants to deny this. But I wish he had explained at somewhat greater length why he thinks he is entitled to deny it. What his remarks suggest is this: There are conditions laid up in the meaning of 'water' that are sufficient to rule out certain things' counting as 'water', even if there isn't enough information there to rule anything in. In particular, the concept rules that there is no water on Dry Earth. So we can say that there is something thinkable there, even if it has "less taxonomical power than might ideally be wanted".

If this is what Steve has in mind, then I wish he had filled it out a little more, because in the way that I understand his proposal, it flies in the face of what most externalists believe about natural kind terms. Let me explain.

This content downloaded from
128.122.149.96 on Sat, 17 Apr 2021 01:23:18 UTC
All use subject to https://about.jstor.org/terms

What considerations drive someone to be a concept externalist in response to the Twin Earth experiments? Why not be a narrow content theorist?

One of the principal considerations has to do with how *thin* the narrow concept associated with a natural kind term would have to be, if we removed the information associated with its extension. For let us ask this: What can we plausibly insist upon up front is packed into the concept of a natural kind, say, to use Putnam's famous example, the concept cat? Furriness? Having four legs? Meowing? Not even the concept animal, Putnam argued, could be thought of as built into the concept cat (cats might turn out to be robots). Similarly for gold (not required that it be yellow or a metal or precious) and water (not required that it be a liquid or odorless or tasteless) and so on. If we insisted that narrow contents *are* thinkable contents, then we would have to conclude that most of our thoughts about natural kinds are virtually blank and so don't determine, in and of themselves, any sort of determinate truth conditions for themselves.

To avoid this unacceptable construal of content, the externalist reasons, we should think of the truth conditions as determined by the causal facts and then let them individuate the thinkable contents, the concepts. That way we have a notion of content that is rich enough to distinguish between cat and cat–robot, water and twin water, elm and beech, and which determines the right satisfaction conditions for them. The price is letting the externally determined reference fix the thinkable concept.

If I understand Steve correctly, however, he is basically saying that if we happen to be living on Dry Earth, then we should think of the thinkable contents associated with 'water' as given by whatever *narrow contents* are associated with 'water'. But this proposal is problematic for two reasons. One is that there may be nothing in the narrow concept except the bare idea of some kind of stuff or other, which by virtue of its blankness would rule nothing out. And the other is that, if this is an acceptable reaction to Dry Earth, why is it not an acceptable reaction to Twin Earth? If there is nothing inherently wrong with thinking of virtually blank narrow contents as thinkable contents, how do we get an *argument* for concept externalism on the basis of the Twin Earth experiments?

But even if we were to set aside all of these objections and accept Steve's proposal, how does it help vindicate the compatibility of privileged self–knowledge and concept externalism? If the only sense in which I know my thoughts *a priori* doesn't enable me to tell whether I have a virtually blank concept of 'water' that doesn't have

This content downloaded from
128.122.149.96 on Sat, 17 Apr 2021 01:23:18 UTC
All use subject to https://about.jstor.org/terms

anything to say about whether most actual and possible substances are or are not water, or a concept that's so rich that it distinguishes between water and twin water, then I, for one, will have lost my grip on what this vindicated knowledge is supposed to amount to.

3 Reply to Corbí

Corbí essentially makes two claims. First, he argues that a reasonable account of privileged access would not yield the *a priori* knowability of premise (2), so compatibilism as I conceive it isn't what needs arguing about. Second, he contends that the externalist is not committed to the *a priori* knowability of (1) in any event, so my argument against the misconceived target misfires anyway. I shall take these points in turn.

To claim that I can know that I have the concept water *a priori* is to claim that I can know without empirical investigation that I have thoughts that involve the concept water. Since for me, as for Frege and lots of others, concepts just are what thought contents are composed of, this is just another way of saying that I can know without empirical investigation what my thought contents are.

This claim I take to be, along with the principal compatibilists involved in this debate —for example, Tyler Burge and Donald Davidson— the minimal core of a doctrine of privileged access. It is this claim, as I understand them, that they insist can be reconciled with externalism. My argument is designed to show that this is not true. Claiming, as Corbí does, that no one should construe the doctrine of privileged self-knowledge in this way, does not engage the point of my argument, which is that *this* doctrine, as all parties to the debate agree in construing it, is not compatible with externalism.

If (2) isn't the way we should construe privileged access, then how, according to Corbí, should we construe it? Corbí's view appears to be that we know what our thought contents are in the same way that we know that we have certain abilities, for example, that we know that we know how to ride a bike. He says:

> How do I know about this ability of mine? There is a trivial sense in which I needn't carry out any empirical investigation to have this valuable piece of knowledge. But, indeed, no one would deny that my knowing that I possess that ability depends on my knowledge of certain facts of the world that involve my previous performances on a bike. (p. 233)

But this is like saying that there is "a trivial sense" in which I needn't

This content downloaded from
128.122.149.96 on Sat, 17 Apr 2021 01:23:18 UTC
All use subject to https://about.jstor.org/terms

conduct an empirical investigation to know that the next time I sit on this chair it will support me, although no one would deny that my knowing this depends on my previous experiences with this and other chairs. The point is that the knowledge in these cases is through and through empirical: it's based on induction on past experience. There is no useful sense in which my knowledge of the content of my occurrent thoughts is based on induction. Indeed, the suggestion is probably not just wrong but incoherent. In any event, when Tyler Burge seeks to vindicate compatibilism with his deployment of the idea of a logically self–verifying second–order thought, he is evidently not trying to reconcile inductively based knowledge of thought content with externalism.

Corbí's second claim is that an externalist is committed to (1) only if water, in addition to aiming to name a natural kind, actually does name a kind. Since, however, it is not *a priori* knowable whether a concept does name an actual kind, (1) is not *a priori.*

I confess to not fully understanding Corbí's remarks here. To support his claim he spends some time rebutting a passage of mine, a passage which is used merely to pave the way for the Dry Earth argument. That's the crucial argument that needs rebutting. Corbí does eventually come around to discussing this argument. His idea, if I understand him, is that on Dry Earth 'water' will express a complex concept that turns out not to name any kind. He rightly notes that I would object that this has 'water' expressing a compound concept in one person and an atomic concept in another functionally identical person, with the two differing only in respect of the external circumstances that they happen to be in. But he expresses puzzlement as to what would be problematic about this.

He thinks that on my view a concept is compound if it's individuated solely in terms of its functional role (p. 241). But this is a bad misunderstanding. I understand what it is for a concept to be compound in that way that I take it it is commonly understood —i.e., as saying that a concept is literally composed, typically by Boolean means, out of other concepts. Thus, bachelor is compound because it *is* the concept unmarried and male. Clearly, this notion of compoundness has nothing to do with being individuated solely in terms of functional role. If the constituent concepts of a compound concept are themselves externalistically individuated, then, naturally, the compound concept won't be individuated solely in terms of its functional role.

Now, having clarified what we mean by the compoundness or compositionality of a concept, it is easy to see why it couldn't supervene on external circumstances. The point to bear in mind is that having

This content downloaded from
128.122.149.96 on Sat, 17 Apr 2021 01:23:18 UTC
All use subject to https://about.jstor.org/terms

a compound concept requires one to have all the ingredient concepts that compose that concept, including the concepts of conjunction or disjunction. Suppose 'water' expresses an atomic concept A on Earth, and that on Twin Earth it expresses the compound concept B and C and D and E. Then we would have to say that external circumstances determine not only *which* concept a given word expresses, but also *how many* concepts a person has. Without any change in one's cognitive or functional capacities, and just by being in one environment as opposed to another, one might have either one concept or five. I don't see how this could happen and that is my basis for saying that the compositionality of a concept does not supervene on external circumstances.

4 Reply to Moya

Moya echoes Corbí's misgivings about whether the apriority of (2) correctly expresses the doctrine of privileged self–knowledge. He says:

> Though Boghossian takes this premise to be obviously acceptable, I find the way it is stated rather controversial. Normally, self–knowledge is held to extend to thought contents, rather than to the concepts one possesses in thinking those contents. In fact, it is hard to accept that one enjoys privileged access to the concepts one has, at least if this means the concepts one *masters* in having a certain thought. Normal adults, I would say, know directly, non–inferentially, what they are currently thinking in that they can give true reports of these thoughts, but I am pretty sure that they are often quite wrong about their mastery of the concepts that make up those thoughts. (p. 245)

This confuses knowing what concept one has with knowing how to explicate or analyze that concept. On behalf of the doctrine of privileged access, I claimed the former but not the latter. I never assume, merely on the basis of the doctrine, that we know that arthritis can't be a disease of the thigh, but only that if we know that we have arthritis thoughts then we know that we have the concept arthritis. Trivially, the one item of knowledge entails the other. If I know that I am thinking that arthritis is wet (or whatever) I know that my thought involves the concept arthritis.

More importantly, Moya claims that, *pace* what I say, an externalist could accept that a term with an empty extension expresses an atomic concept. He gives the truth conditions of 'Water is wet' on Dry Earth as follows:

This content downloaded from
128.122.149.96 on Sat, 17 Apr 2021 01:23:18 UTC
All use subject to https://about.jstor.org/terms

'Water is wet' on Dry Earth iff Dry Earth's water is wet.

He claims that we could say that because "Dry Earth's water" lacks reference, the right hand side of the biconditional is false and hence that its left hand side is false as well. Since something can't be false without being meaningful, that vindicates the existence of a 'water' concept on Dry Earth. But then in what conditions would 'Water is wet' as used on Dry Earth be true? He answers as follows:

> Well, it would be true if there were in fact a substance, a natural kind, in the extension of 'water' and this substance were wet... [An externalist could then say] what is the concept expressed by tokens of 'water' on Dry Earth: it is the concept Dry Earthians would have if there were in fact a natural kind with the right appearance denoted by tokens of 'water'. (p. 250)

I take it that neither 'mimsy' nor 'borogrove' has meaning, so that someone uttering 'All mimsies are borogroves' doesn't say anything and his statement has no truth conditions. Mimicking Moya's suggestion, however, we seem to be able to say that it is meaningful after all. For consider:

Mimsies are borogroves' is true iff mimsies are borogroves.

Since neither 'mimsy' nor 'borogrove' refers to anything, we can take the view that 'Mimsies are borogroves' is false. What are its truth conditions? It would be true if there were some things in the extension of 'mimsy' and some things in the extension of 'borogrove' and if all the former were instances of the latter. And the concept we do express by 'mimsy' is the concept that we would express if there were in fact things in the extension of 'mimsy'.

This content downloaded from
128.122.149.96 on Sat, 17 Apr 2021 01:23:18 UTC
All use subject to https://about.jstor.org/terms