

Brains in vats and model theory

TIM BUTTON

button@cantab.net

Hilary Putnam’s anti-sceptical BIV argument first occurred to him when ‘thinking about a theorem in modern logic, the “Skolem–Löwenheim Theorem”’¹ One of my aims in this paper – following Putnam, Thomas Tymoczko, and Adrian Moore – is to explore the connection between the argument and the Theorem. But I also want to draw some further connections. In particular, I think that Putnam’s BIV argument provides us with an impressively versatile template for dealing with sceptical challenges, and I want to probe that template.

The template I have in mind was most clearly set out in a remarkable – and remarkably brief – passage:

Suppose we (and all other sentient beings) are and always were ‘brains in a vat’. Then how does it come about that *our* word ‘vat’ refers to *noumenal* vats and not to vats in the image?²

Two points leap out from this passage. First, the intended refutation of scepticism turns on *semantic* considerations. Second, the refutation involves an accusation of *self-refutation*:³ if the sceptical scenario actually obtained, then the sceptic would be unable to formulate her sceptical scenario. These two points allow us to unify some of Putnam’s most enduring contributions to the realism/antirealism debate: his discussions of brains-in-vats, of Skolem’s Paradox, and of permutations.

After considering some general questions about our engagement with scepticism (§1), I shall present and defend my favourite version of Putnam’s BIV argument (§2). My aim is not to offer a complete defence of the argument, but to highlight its working parts, and also to explain why the argument is significant. The key point is that, having answered the BIV sceptic, we must jettison any philosophical picture which treats BIV scepticism as *unanswerable*.

I then apply the same considerations against two further varieties of scepticism. In particular, I first show how to answer skolemism (§3), and then show how to answer the kinds of semantic scepticism that arise during Putnam’s model-theoretic arguments (§4). In all three cases, the anti-sceptical argument does not merely defeat the sceptic; it also shows us that we must reject some *prima facie* plausible philosophical picture.

¹Putnam (1981: 7).

²Putnam (1977: 487); Putnam’s emphasis.

³The ‘self-refuting’ nature is emphasised by Putnam (1981: 7–8).

1 Why to engage with scepticism, and how

If a sceptic is any good, then she will present us with an *unanswerable* challenge. One might reasonably wonder, then, why we should even bother to engage with scepticism. Perhaps the right response is simply to dismiss unanswerable challenges as *uninteresting*,⁴ maybe even just *because* they are unanswerable.

With Putnam, I think that this is wrongheaded. There are at least three good reasons for engaging with scepticism, and understanding those reasons will also help us to understand how we ought to engage with it.

Reason 1. Even if we fail to answer the sceptic, we can learn from engaging with her. If the sceptic is any good, then she raises questions concerning concepts that are central to our way of thinking. Dialogues with sceptics therefore potentially provide us with a means for profound philosophical investigations.

This thought supplies a criterion for when it is worth engaging with scepticism. Suppose the sceptic starts by employing ‘assumptions which we ourselves hold’,⁵ and so raises a challenge which seems to arise from *within* our own worldview. Then, if she argues well, she will ‘confront us with an antinomy’ from within our own worldview, ‘and one always learns from an antinomy’⁶ (or, almost always). Call the sceptic who raises an antinomy from within our own worldview an *internal sceptic*. It is worth engaging seriously with such a sceptic, since we may learn about our own worldview, or even about the world itself.

Conversely, if the sceptic does not use our own tools against us – if she is not an internal sceptic – then her sceptical scenarios will be *genuinely* uninteresting. In that case, we are free simply to walk away.

Reason 2. To concede that the sceptical challenge is unanswerable is to capitulate too easily. As we shall see, some sceptical challenges *are* answerable.

Saying this, of course, raises the question of what it takes to *answer* the sceptic. After all, a really irritating sceptic might refuse to accept any premise in any argument we present against her; or she might refuse to accept any rules of inference; or whatever. Such a sceptic cannot be *answered*, in the sense that nothing we can say will stop her from pestering us. This, however, should not bother us. For if the sceptic refuses to accept *anything*, then she will have ceased to be an internal sceptic, in the sense outlined above. She will no longer be outlining an antinomy that arises from within our own worldview. And as such, we will be free to ignore her.

Consequently, when we engage with some sceptical challenge, our strategy for answering will always be to show that it is not an *internal* scepticism.⁷

⁴Something like this is suggested by Devitt (1984: 64, 75).

⁵Putnam (1994: 284).

⁶Putnam (1994: 284).

⁷This is emphasised by Putnam (1994: 284–5) and Button 2013, pp. 121–3.

Reason 3. When a sceptical challenge is answerable, those philosophers who have declared it to be unanswerable may be guilty of more than mere hastiness. The *unanswerability* of certain sceptical challenges is central to certain philosophical pictures. In such cases, if we discover that we can answer the sceptical challenge after all, then we also discover that we must jettison the associated philosophical picture.

I have presented three reasons for engaging with scepticism; reasons which also suggest how we should set the terms of that engagement. At the moment, though, these are just abstract reasons-in-principle. To make them concrete, we must start engaging with particular versions of scepticism.

2 BIV scepticism

I shall begin with BIV scepticism, and Putnam's refutation of it. Much has been written about Putnam's argument, and I have done my best to defend it fully elsewhere.⁸ Consequently, my main aim in this section is not to offer a complete defence of the BIV argument. It is simply to have a plausible argument against BIV scepticism, which can serve as a blueprint for further anti-sceptical arguments in subsequent sections.

Introducing the sceptic. Say that someone is a *BIV* just in case, for the entire duration of their lives – past, present and future – they and all of their worldmates are brains in vats. The *BIV scenario* is then the scenario in which everyone is a BIV. And anyone who thinks that the BIV scenario actually obtains is a *BIV sceptic*.⁹

BIV scepticism can be motivated by the following apparently natural line of thought:

Obviously, all the information that we obtain about the empirical world is obtained through our sense organs. Suppose, now, that those organs have been replaced with a computer, which is wired into our brain, but which feeds us exactly the same electronic signals as our sense organs typically do. We would obviously be unable to tell the difference. So: we cannot tell whether things are as they seem, or whether we are BIVs.

The challenge raised here is (at least initially) a variety of radical *Cartesian* scepticism. We are being threatened with the thought that appearances might be radically, nightmarishly deceptive, so that (almost) all of our beliefs are false. Indeed, the BIV scenario is not much more than a sci-fi reworking of the nightmare sceptical scenario discussed in

⁸Button (2013: chs. 12–14).

⁹It is probably more normal to use 'the BIV sceptic' for someone who maintains that we do not *know* whether the BIV scenario obtains. I avoid this for two reasons: first, it is best to avoid epistemological notions at the outset (see p.7 for more); second, though nothing ultimately turns on it, the discussion is easier to follow if the sceptic is actually *advancing* a position. Exactly similar comments apply to the other sceptical positions I discuss in this paper.

Descartes's First Meditation. However, precisely this sci-fi gloss allows the BIV sceptic to maintain that she is using own principles against us; that she is raising an antinomy from within our worldview. As such, the BIV sceptic purports (I stress, *purports*) to present radical Cartesian scepticism as an *internal* scepticism.

A Moorean response. Since BIV scepticism purports to be an internal scepticism, we can reasonably hope to learn something from engaging with it. However, before I consider Putnam's response to BIV scepticism, I wish to consider a *Moorean* response to BIV scepticism. To be clear: I do not want to argue for or against this Moorean response. Instead, I introduce it for two reasons. First, it will help me to flesh out the idea of *internal scepticism*, as discussed in §1. Second, this Moorean response usefully contrasts with Putnam's own BIV argument.

Here is the Moorean response to the BIV sceptic:¹⁰

- (iB) BIVs have no hands.
- (iiB) I have two hands.
- (iiiB) So: I am not a BIV.

Both of the premises in this Moorean argument sit well within my worldview: by definition, BIVs have no hands; whereas my best theory of the world tells me, among other things, that I have hands. Accordingly, the BIV sceptic must deny some claim that sits well within my worldview. And this might suggest that the Moorean response satisfactorily answers BIV scepticism, in the sense that it shows that BIV scepticism cannot be an *internal* scepticism.

That thought, however, is too quick. The BIV sceptic claims to motivate the BIV scenario by focusing on a central aspect of our empirical theorising. In particular, she focuses on our beliefs concerning how we receive information about the empirical world through our senses. She then uses this particular aspect of our theorising to throw shade on the remainder of our theorising. Her point is that, if I *am* a BIV, then it will falsely *seem* to me that I have hands; so that *even by my own lights* I cannot take for granted that I have hands. And it is not too difficult to get into the state of worrying that the Moorean anti-sceptical argument is no more suasive than the BIV sceptic's initial challenge:

- (iB) BIVs have no hands.
- (¬iiiB) I am a BIV.
- (¬iiB) So: I have no hands.

To repeat: I am not siding here with the sceptic against the Moorean. My point is just that an internal scepticism need not offer a challenge which is compatible with literally *all* of our beliefs. (Indeed, anything that was so compatible would scarcely constitute a *challenge*.) Rather: an internal scepticism need only invoke some of our most central

¹⁰G. E. Moore 1939.

principles of reasoning, in order to undercut our ability to appeal to certain other beliefs – often beliefs which we hold very dear – in our attempt to answer the sceptical challenge.

Now, since the internal sceptic uses only *some* of our principles against us, and since the centrality of such principles comes by degree, one might well expect that there is not always a sharp answer to the question of whether a particular version of scepticism counts as *internal*. Similarly, we might well expect that there is not always a sharp answer as to whether a particular anti-sceptical argument counts as *Moorean*. I am very sympathetic to these ideas, but I cannot explore them properly here.¹¹ For now, it suffices to note the following: there is at least a workable distinction (if not a sharp dichotomy) between internal and non-internal scepticism, and between Moorean and non-Moorean responses to scepticism.

The BIV argument. Putnam's BIV argument is not a Moorean response. Instead, it aims to show that BIV scepticism is *self-refuting* when considered as an *internal* scepticism. Here is the argument, as I understand it:¹²

- (1B) A BIV's word 'brain' does not refer to brains.
- (2B) My word 'brain' refers to brains.
- (3B) So: I am not a BIV.

The two premises establish that my language is not the language of a BIV, from which (3B) follows. For the remainder of this section, I shall sketch defences of both premises, and then explain why the conclusion is so significant.

The language of BIVs. Premise (1B) is to be justified by rejecting all *magical* theories of reference. In slightly more detail, here is a sketched defence of (1B):¹³

Consider Brian, a BIV. Because Brian is a BIV, no one in Brian's world has ever tokened the word 'brain' whilst interacting appropriately with a brain. So, on a naïve causal theory of reference, Brian's word 'brain' cannot refer to brains. And we will be led to the same conclusion, even if we offer a more subtle theory of reference, provided that the theory retains some plausible causal constraint on reference. Indeed, in order for Brian's word 'brain' to refer to brains, there would need to be some *magical* connection between words and objects, such that there was (for example) an intrinsic connection between the letters *b-r-a-i-n*, and brains.

¹¹They fit extremely well with what I say in Button (2013: chs. 15–16).

¹²I offer this version of the argument in Button (2013: 118, 125). It draws from Tymoczko (1989: 285), Wright (1992: 74), Brueckner (1992: 127) and Putnam (1992: 369).

¹³This sort of justification is given by: Putnam (1981, pp. 16–7; 1992, p. 369) and Wright (1992: 219). See also my Button (2013: 118–23).

To offer a full defence of (1B), more would need to be said. But my main aim here is not to defend the BIV argument itself in great detail. So I shall simply take it as established that (1B) holds if we reject all magical theories of reference.

One point, however, deserves comment. A magical theory of reference is not *formally* inconsistent. Consequently, the BIV sceptic could resist the BIV argument, without falling into formal inconsistency, by embracing magic. But, in so doing, she would have given up on the idea of employing our own tools against us. She would have given up on the idea of showing that, *by our own lights*, we cannot simply ignore the BIV scenario; for our own lights tell us not to countenance magic. In sum: if BIV scepticism must rely upon a magical theory of reference, then it is no longer an *internal* scepticism, and so, for reasons explained §1, we have earned the right to ignore it. Contraposing: since the BIV sceptic wants to be taken seriously, she must accept premise (1B).

My language. Premise (2B) is to be justified by defending disquotation in the mother-tongue. In slightly more detail, here is a sketched defence of (2B):

When I present a semantics *for* my own language, I must *use* my own language. Consequently, I have no option but to offer disquotational semantic clauses. And (2B) is just such a disquotational clause.

In any other context, no further justification for (2B) would be required. But we are in the strange situation of arguing with a sceptic; and the sceptic is likely to push back on this point.

In particular, the BIV sceptic will probably complain that the argument (1B)–(3B) is nothing more than a *Moorean* response to her sceptical challenge. Since she thinks that we are BIVs, she denies that we have hands; similarly, since she thinks that we are BIVs, and has just agreed with (1B), she denies that our word ‘brain’ refers to brains. Her sceptical challenge then renews itself, as follows:

- (1B) A BIV’s word ‘brain’ does not refer to brains.
- (¬3B) I am a BIV.
- (¬2B) So: My word ‘brain’ does not refer to brains.

However, affirming (¬2B) undercuts the BIV sceptic’s ability to formulate her own sceptical challenge, as I shall now explain.

The BIV sceptic attempts to present her challenge by talking us through the BIV scenario. That is, in saying ‘everyone is a BIV’, she hopes to present us with the worry *that* everyone is a BIV, i.e., with the worry *that* everyone ever is always an envatted *brain*. However, given (¬2B), saying ‘everyone is a BIV’ would simply fail to present us with any thought *about* brains. It is unclear what thought (if any) it would present us with; but that need not concern us. The point is just that she would simply have *failed* to confront us with the worry *that* we are brains in vats. Thus, premise (2B) is implicitly required

by the BIV sceptic herself in the very *formulation* of her sceptical challenge.¹⁴ For similar reasons, premise (2B) is required by the BIV sceptic even in order for her vocalisation of (\neg 2B) to present us with the thought *that* my word ‘brain’ does not refer to brains. In short: to deny (2B) is *self-refuting*.¹⁵

At the very least, then, the BIV sceptic needs to accept that the word ‘brain’, as it occurs in the formulation of the BIV scenario, somehow does refer to brains after all. In desperation, then, the BIV sceptic might suggest that the BIV argument is too crude; she might maintain that it ignores the fact that reference depends upon *context*. More specifically, she might abandon (1B) and (2B) in favour of the following:

- (aB) In *ordinary* contexts, my (i.e. a BIV’s) word ‘brain’ does not refer to brains.
- (bB) In *philosophical* contexts, my (i.e. a BIV’s) word ‘brain’ refers to brains.

This will allow her to formulate the BIV scenario, at least in philosophical contexts. Unfortunately, (bB) invokes a magical theory of reference, every bit as much as the denial of (1B). Granted, it is magic which is *local* to the philosophy classroom – or, rather, to the *appearance* of the philosophy classroom, since there are not *really* any philosophy classrooms in the BIV-world – but it is magic nonetheless. We have already rejected magic in general, and we should equally well reject local magic.¹⁶

The debate surrounding the BIV argument continues beyond this point, but I shall not pursue it any further.¹⁷ We have arrived at the following. The BIV sceptic must accept (1B), because the alternative involves invoking magic, and she wants to be an internal sceptic. The BIV sceptic must accept (2B), because she relies upon it in formulating her own sceptical scenario. We have therefore established the following: *if BIV is an internal scepticism, it is self-refuting*. The BIV sceptic is answered.

I shall now explain why this matters.

Epistemological versus semantic externalism. Sometimes the challenge posed by the BIV scenario is explicitly framed in terms of knowledge: *you cannot know that you are not a BIV*. In response to this challenge, we might try to unpack the concept of knowledge along externalist lines. We might then respond in something like the following fashion:

¹⁴Compare Tymoczko (1989: 284–6).

¹⁵Importantly, this is not the kind of self-refutation involved in the Moorean contradiction ‘it is raining and I believe it is not raining’. Whether the Moorean contradiction is true or false, I can *represent* the scenario it describes. By contrast, if the BIV scenario obtained, the sceptic would be unable even to *represent* that scenario. See Brueckner (1996: 277) and Button (2013: 126–7).

¹⁶Compare Ebbs (1996: 514–8) argument that content scepticism vacillates between a ‘subjective’ and an ‘objective’ point of view.

¹⁷A particularly interesting suggestion is that the BIV argument answers the initial sceptical challenge, but leaves us with a lingering, *ineffable*, anxiety. For more, see Wright (1992: 93) and Button (2013: 137ff.).

If we are not BIVs, then we have appropriately safe and sensitive true beliefs (which entail) that we are not BIVs. So, given our account of knowledge, we *know* that we are not BIVs (assuming that, in fact, we are not BIVs). Hence we have answered the challenge of how we *know* that we are not BIVs.¹⁸

Whatever the merits of this kind of argument, it cannot help to address the root anxiety induced by the BIV scenario. The sceptic might phrase her challenge in terms of knowledge, but she need not. She can instead say: *you cannot be certain that you are not a BIV*; or: *you cannot be justified-from-within that you are not a BIV*; or: *nothing guarantees that you are not a BIV*. Most simply, she can just ask: *might you be a BIV?*¹⁹

Putnam's BIV argument invokes *semantic* externalism, rather than *epistemological* externalism. And this makes his anti-sceptical argument more robust, in at least one clear sense. Semantic externalism allows us to establish that a BIV cannot entertain certain thoughts. In particular, Brian the BIV cannot even ask himself whether he is a BIV, let alone whether he might be one, or whether he is *that* he is not one. Consequently, Putnam's anti-sceptical argument guarantees that I am not a BIV, *because I can so much as ask whether I am one.*²⁰

Even then, of course, the BIV argument has its limitations. It tells me that I am not a BIV, but it does not straightforwardly allow me to conclude that my brain was not scooped from my skull last night and plunged into a vat. (In that scenario, my word 'brain' presumably *would* still refer to brains.)²¹ So the question still remains: Why should we *care* that we can answer the BIV sceptic?

Why this matters. At the start of this section, I said that we should engage with BIV scepticism because it purports to be an *internal* scepticism. It can just seem obvious that we cannot rule out the BIV scenario; obvious, that is, that we might actually be BIVs. So it can just seem to be a philosophically neutral starting point, that a theory which gives every appearance of being true might really be radically false. This last principle is characteristic of *external realism*; the realist position whose 'favourite point of view is a God's Eye point of view', as Putnam puts it.²² So: it might just seem obvious that external realism presents is the neutral starting point to begin one's philosophising.

In these terms, the BIV argument shows us that this philosophical picture is *not* a neutral starting point. In fact, so far from being neutral, it is *disastrous*. For if the external realist picture were correct, then BIV scepticism would be unanswerable. Since

¹⁸Zalabardo (2009: 80–2) presents a detailed development of this kind of strategy. (Although this indented paragraph is not a quote from Zalabardo.) See also Williamson 2000, ch. 8.

¹⁹So: Zalabardo (2009: 82) caveats his argument by saying 'Of course we might be wrong about this', i.e. about our assumption that we are *not* BIVs. The sceptic can walk away from issues concerning knowledge; this 'might' is all she needs to get a foot in the door.

²⁰Tymoczko (1989: 281, 284) puts the point this way.

²¹For extensive discussion of vat-variations and their significance, see Button (2013: chs. 15–16).

²²Putnam (1981: 49). In Button (2013: esp. ch. 8), I argue that this principle (which I call the Cartesianism Principle) is the *distinctive* feature of external realism.

BIV scepticism is answerable, the picture is fatally flawed. More specifically: the picture must either rely upon a tacit assumption of magic, as per the discussion of (1B); or it must be self-refuting, as per the discussion of premise (2B).

In short: the BIV argument teaches us that external realism is either incoherent or magical. That is why the BIV argument matters.

3 Skolemism

I now want to consider a sceptical challenge which arises in the context of mathematics. It falls out of reflections on Skolem's Paradox, in something like the way that BIV scepticism falls out of reflections on apparently mundane thoughts concerning how we find out about the world around us. It can be refuted via a similar argument, which I believe was first formulated explicitly by Thomas Tymoczko and Adrian Moore.

Skolem's Paradox. Consider the following result:²³

The Löwenheim–Skolem–Mostowski (LSM) Theorem. For every transitive model \mathcal{F} of ZF, there is a countable, transitive model \mathcal{H} which is elementary equivalent to \mathcal{F} .

Presumably, ZF has a transitive model. So, ZF has a countable, transitive model, \mathcal{H} , by the LSM Theorem. Now, ZF proves 'there is an uncountable set', so \mathcal{H} must make this sentence true. Formalising this slightly, we say:

$$\mathcal{H} \models \exists x \neg \exists y (y \text{ is an enumeration of } x)$$

Hence, for some b in \mathcal{H} 's domain:

$$\mathcal{H} \models \neg \exists y (y \text{ is an enumeration of } b)$$

But because \mathcal{H} is both countable and transitive, every member of \mathcal{H} 's domain is countable. So in particular:

²³To say that a model, \mathcal{M} , is *countable* is to say that its domain, M , is countable. In general, to say that x is *countable* is to say that there is an enumeration of x (i.e. an injection from x to the natural numbers), and x is *uncountable* iff x is not countable.

To say that a model, \mathcal{M} , is *transitive* is to say the following: If $a \in b \in M$, then $a \in b$ iff $\mathcal{M} \models a \in b$. The effect of this stipulation is that what \mathcal{M} takes to be membership *really is* membership.

To say that \mathcal{M} is *elementary equivalent* to \mathcal{N} is to say that they satisfy (i.e. make true) exactly the same (first-order) sentences.

The LSM Theorem is Theorem 3.8 of Mostowski 1969. Mostowski certainly deserves to be named in the statement of this result, since its proof involves Mostowski's Collapsing Lemma. It is well known that Skolem-style results (such as the LSM Theorem) can be avoided by moving from first-order logic to second-order logic with a *full* semantics. For a full survey of this point and its dialectical significance, see Button and Walsh (MS: §§5–6).

$\exists y$ (y is an enumeration of b)

This is (one version of) Skolem's Paradox. And, whilst it might initially be surprising, it is readily explained. There is an enumeration of b . However, this enumeration lies outside \mathcal{H} 's domain, and so beyond the range of the quantifier ' $\exists y$ ' as interpreted *within* \mathcal{H} . In a slogan: to be countable-in- \mathcal{H} is not to be countable *simpliciter*. Paradox dissolved.²⁴

Introducing the sceptic. Nonetheless, this version of Skolem's Paradox naturally gives rise to a sceptical line of thought.²⁵

Mathematical entities are abstract entities, in Platonic Heaven. Given their abstractness, it is obvious that we could only refer to them directly if we had some kind of supernatural, magical powers. We have no such powers, so the best hope of our picking them out is by employing formal theories. So: we pick out the sets, for example, as the entities satisfying some particular formal set theory, such as (some extension of) ZF.

But suppose, now, that we do not live in a universe with any uncountable sets, but instead live in a countable, transitive universe built along the lines of \mathcal{H} .²⁶ Then we would surely be unable to tell the difference. Every sentence of our set theory that 'ought' to be true would be true; every description in the language of set theory that 'ought' to be uniquely satisfied would be uniquely satisfied; and so forth. So: we cannot tell whether there really are uncountable sets, or merely seem to be.

In short, we have been led to the worry that every set is countable. This is the *skolemist scenario*. For brevity, I shall refer to the denizens of the skolemist scenario as *smallworlders*, and anyone who thinks that the skolemist scenario actually obtains is a *skolemist*.

Like the BIV scenario, the skolemist scenario seems to arise in a fairly mundane way. We begin with some fairly pedestrian reflections – not this time upon how our senses inform us about the world, but on incontestable results from model theory – and we are swiftly confronted with a foundational challenge. Like BIV scepticism, then, skolemism presents itself as an *internal* scepticism. This is why skolemism initially commands our attention.

This is the first of several similarities between skolemism and BIV scepticism. However, before I explore other similarities, it is worth highlighting an important disanalogy. As mentioned in §2, the BIV scenario is primarily a vehicle for Cartesian angst.

²⁴For more on the mathematics of Skolem's Paradox, see Bays (2009: §2).

²⁵See A. W. Moore (2001: 160–4).

²⁶This step is surprisingly contentious. The implicit idea seems to be that we can simply input 'the universe' into the LSM Theorem, generating \mathcal{H} . However, presumably 'the universe' is *class*-sized, and so does not form a model that can be plugged into the LSM Theorem. This point is pursued by Velleman (1998), Bays (2001, pp. 335–40; 2007, pp. 119–23) and Button (2011).

The initial worry that it generates is that appearances might be radically deceptive, so that (almost) all of our beliefs are false. By contrast, the skolemist scenario does not ask us to consider a situation in which (almost) all of our mathematical beliefs seem to be true but are really false. On the contrary, the skolemist deliberately places us in a scenario which is tailor-made to make our favourite mathematical theory true. Her sceptical challenge is that our theory is true for ‘unintended’ reasons; because Platonic Heaven is much smaller than usually thought. Skolemism is therefore not a version of *Cartesian* scepticism, but of something rather different, which I shall call *metaphysical* scepticism.²⁷

BIV scepticism and skolemism consequently begin life as different varieties of scepticism. However, they need not remain that way permanently. After some reflection, we might come to decide that Brian the BIV has mostly true beliefs, but is simply deceived about the nature of his world. In this case, the BIV scenario will equally have become a vehicle for *metaphysical* scepticism.²⁸ None of this, of course, affects the success of the BIV argument, since all the argument requires of Brian is that he fails to refer to brains. But it will be important, in what follows, to recall that the skolemist scenario is *never* a vehicle for Cartesian angst.

Moorean responses. I now want to return to the similarities between BIV scepticism and skolemism. First among them is that they are both presented as *internal* scepticisms. Now, when I considered BIV scepticism, I probed what this amounted to by considering a Moorean response. It is worth considering a Moorean response to skolemism:

- (is) If I am a smallworlder, then there are there are no uncountable sets.
- (iis) There are uncountable sets.
- (iiis) So: I am not a smallworlder.

Here, (is) is true by definition. And (iis) is surely as central to our mathematical practice as anything else. So: we have a Moorean response to the skolemist.

This kind of Moorean response seems to underpin arguments against skolemism which make appeal to ‘naïve’ or ‘informal’ mathematics; in such cases, ‘naïvety’ is introduced in order to defend (iis). The skolemist will claim, though, that (iis) simply begs the question against her. She will point out that, if I am a smallworlder, then I simply *cannot* take mathematical practice at face value, and so cannot infer much from the fact that all mathematicians are happy to claim that there are uncountable sets. Indeed, the skolemist will say that it is just as question-begging to take mathematical *utterances* at face value, as it is to respond to BIV scepticism by taking *appearances* at face-value.

As in the discussion of BIVs, I have not mentioned this Moorean response because I think it is good or bad. Instead, I have raised it as a contrast with the kind of response

²⁷Compare Moore’s (2011) discussion of Georg.

²⁸See A. W. Moore (2011: §3) and Button (2013: 141–8).

I want to offer, and also to emphasise the sense in which skolemism presents itself as an internal scepticism. Skolemism picks up on (what it takes to be) a central aspect of our philosophical picture: that our only access to mathematical objects is through the truth of certain mathematical sentences. And, on the basis of this principle, it purports (*purports*) to establish that *even by our own lights* we cannot take for granted that there are uncountable sets.

The anti-skolemist argument. The aim of the anti-skolemist argument, then, is to show that this is mistaken. In particular, it aims to show that if skolemism is an internal scepticism, then it is self-refuting. Consequently, we will be forced to reject the picture, according to which our only access to mathematical objects is through the truth of certain mathematical sentences.

To offer the anti-sceptical argument, it will help to recap how skolemism arose in the first place. The LSM Theorem convinced us of the existence of a countable, transitive model, \mathcal{H} , of ZF. No paradox ensued, since being uncountable-in- \mathcal{H} is not to be uncountable simpliciter. However, skolemism was driven by considering the sceptical thought that, metaphorically, we might be ‘trapped within’ \mathcal{H} . Stipulating that a set is *countable ^{\mathcal{H}}* just in case it is countable-in- \mathcal{H} , here is the anti-skolemist argument:

- (1S) A smallworlder’s word ‘countable’ applies only to countable ^{\mathcal{H}} sets.
- (2S) My word ‘countable’ does not apply only to countable ^{\mathcal{H}} sets.
- (3S) So: I am not a smallworlder

The argument is due to Tymoczko and Moore, both of whom are explicitly inspired by Putnam’s BIV argument.²⁹

The language of smallworlders. It is easy to see that the skolemist is committed to (1S). The motivating thought behind skolemism is the following: *If* our world is like \mathcal{H} in containing only countable sets, and *if* our words are interpreted along the lines suggested by \mathcal{H} , *then* all of our mathematical claims will come out true (just as true as before). But to interpret our words along the lines suggested by \mathcal{H} is precisely to say

²⁹Tymoczko (1989: 287–90) and A. W. Moore (2001: 165–7). The argument is also similar to Moore’s (2011: §3) discussion of *set-scepticism*, and to Einheuser’s (2010: 243–6) discussion of scepticism about absolutely general quantification. It is worth, though, highlighting some slight differences between all of these arguments.

First: my immediate aim is only to show that the *skolemist* is incoherently committed to both premises. By contrast, Moore and Tymoczko *themselves* embrace (analogues of) both premises. Second: Moore’s analogue of (2S) is something like: *My word ‘countable’ refers to (i.e. applies to all and only the) countable sets.* This is stronger than my (2S). To argue that the skolemist is committed to (2S) would raise issues concerning the skolemist’s commitment to *absolute generality*, and I would prefer to avoid such issues here. (This also explains why I have deviated slightly from the template laid down by the BIV argument.) I discuss a third difference in fn.39, below.

For a survey of some related anti-sceptical arguments in the philosophy of mathematics, see Button and Walsh (MS: §7).

that the word ‘countable’ should apply only to countable ^{\mathcal{H}} sets. The entire skolemist challenge, then, is built around acceptance of (1S).

To be clear, this defence of (1S) is explicitly *ad hominem*, and rather different from the defence of premise (1B) of the BIV argument. I make no claim to have defended (1S) from all possible challenges. On the contrary, (1S) is extremely vulnerable to attack. The skolemist line of thought began by reflecting on the difficulties of referring to mathematical objects, given their abstractness. But these difficulties might well push us, not towards skolemism, but towards the idea that it is just impossible that our mathematical vocabulary to apply to *anything*. Maybe, indeed, mathematics is just so much *symbol-shuffling*.³⁰

The anti-skolemist argument will not help us to refute this pessimistic thought. But it is not meant to: it is meant only to refute *skolemism*. And it can do that easily enough, since the idea that we cannot refer to mathematical objects *at all* is just as problematic for the skolemist as for the ordinary mathematical platonist. At the risk of repetition: the skolemist *herself* is committed to (1S), since it is presupposed in her attempt to motivate her sceptical scenario.

My language. I now turn to premise (2S) of the anti-skolemist argument. The defence of this premise closely mirrors the defence of premise (2B).

The skolemist is likely to complain that the argument (1S)–(3S) is nothing more than a *Moorean* response to her challenge. Indeed, since she thinks that we are smallworlders, (2S) is precisely what she contests. Her sceptical challenge then renews itself, as follows:

- (1S) A smallworlder’s word ‘countable’ applies only to countable ^{\mathcal{H}} sets.
- (¬3S) I am a smallworlder.
- (¬2S) So: my word ‘countable’ applies only to countable ^{\mathcal{H}} sets.

Affirming (¬2S), however, undercuts the skolemist’s ability to formulate her own sceptical challenge. The skolemist presents her challenge via the skolemist scenario. That is, in saying ‘every set is countable’, she hopes to present us with the worry *that* every set is countable. However, given (¬2S), what she has said simply does not present us with that thought. Instead, it presents us with the thought *that* every set is countable ^{\mathcal{H}} . And the skolemist denies this *herself*. She wants to say:

There are countable ^{\mathcal{H}} sets and uncountable ^{\mathcal{H}} sets; but even the uncountable ^{\mathcal{H}} sets are countable.

However, if this speech is to mean what the skolemist wants it to mean, then the last word in her speech cannot apply *only* to countable ^{\mathcal{H}} sets. In short, if the skolemist is going to be able to articulate, motivate, or defend her sceptical challenge, she needs to *accept* (2S).

³⁰Indeed Klenk (1976: 485–7) suggests that reflecting on skolemism should lead us to formalism.

At the very least, then, the skolemist needs to accept that the word ‘countable’, as it occurs in the formulation of the skolemist scenario, somehow does apply beyond the countable^H sets. In desperation, then, the skolemist might suggest that the anti-skolemist argument is too crude; she might maintain that it ignores the fact that application depends upon *context*. More specifically, she might suggest abandoning (1s) and (2s) in favour of the following:

- (as) In *mathematical* contexts, my (i.e. a smallworlder’s) word ‘countable’ applies only to countable^H sets.
- (bs) In *philosophical* contexts, my (i.e. a smallworlder’s) word ‘countable’ does not apply only to countable^H sets.

When I considered this kind of response in the case of BIV scepticism, I maintained that (b_B) invoked a magical theory of reference. A similar point holds here. In order for the skolemist to defend (as) and (bs), she needs to carve out a distinction between distinctively *mathematical* uses of the word ‘countable’ – as when she opens her set theory textbook and works through its contents – and distinctively *philosophical* uses of that word – as when she formulates the skolemist scenario. The difficulty she faces is that her scenario is formulated using entirely *mathematical* vocabulary (‘countable sets’) and is motivated entirely by the consideration of *mathematical* results (the LSM Theorem). What, then, can be meant by the idea that the statement of the skolemist scenario is *philosophical rather than mathematical*?³¹ Otherwise put: if the skolemist severs the connection between mathematical language and the formulation of her sceptical scenario, then skolemism ceases to be an *internal* scepticism.

As a final effort, the skolemist might maintain that context affects application-conditions in the following, more sophisticated, way.

I spend some time as an active mathematician. Then I stand back and reflect upon what I have done. I realise that I could be interpreted – without loss – as having failed to talk about any uncountable sets. Of course, even as I assert this, I must acknowledge that I am not now failing to talk about any uncountable sets; that is what the anti-skolemist argument teaches me. Skolemism is, indeed, refuted.

But now I stand back and reflect on what I just did in ‘refuting skolemism’. And I realise that everything I said could be interpreted – without loss – as having failed to talk about any uncountable sets, even when I reached the conclusion ‘I am not now failing to talk about any uncountable sets’. Of course, even as I assert this, I must acknowledge that it would be wrong for me to say (in the same breath) that every set is countable, for the anti-skolemist argument can be repeated at this level.

³¹This is very close to parts of Tymoczko (1989: 290); but see fn.39, below.

But now I stand back and reflect...

And so it goes. Whenever I run the anti-skolemist argument, a new sceptical hypothesis immediately pops up. I can answer every such challenge. But at no stage am I stably convinced that I was talking about any uncountable sets. So, whilst I cannot ever say 'every set is countable', on pain of the anti-skolemist argument, I worry that no set is absolutely uncountable, as it were.³²

This line of thought is ultimately, though, 'sophisticated' only in the sense that it is sophistry. It fails to express a serious sceptical challenge, for two reasons.

First: The skolemist keeps raising the same sceptical challenge, over and over again. However, she also acknowledges that the sceptical challenge can be answered whenever it arises. Moreover, the answer that we offer is *exactly* the same answer, every time. Despite this, the skolemist suggests that the sceptical challenge nevertheless remains unanswered, in some profound sense. But why should we not think that the challenge has been straightforwardly *answered*, over and over and over again, until we simply tired of repeating ourselves?³³

Second: The skolemist seems to be suggesting that, for every set, there is some perspective – the *absolute perspective*, as it were – from which that set appears countable. However, the skolemist also grants that, for any perspective that we *can* adopt, we can successfully run the anti-skolemist within that perspective. As such, the skolemist faces a straightforward dilemma. To the extent that I *can* make sense of the idea of the absolute perspective, I can also grasp that one could present the anti-skolemist argument from within that perspective. To the extent that I *cannot* make sense of the idea of the absolute perspective – to the extent that it is ineffable – I cannot make sense of the supposed sceptical challenge.

Why this matters. We could spend more time chasing down further skolemist-epicycles. Nonetheless, I think we have answered the skolemist. More precisely, and as in the case of BIV scepticism, we have shown the following: *if skolemism is an internal scepticism, it is self-refuting*. And, just as in the case of BIV scepticism, this discovery teaches us that a particular philosophical picture is either incoherent or magical.

The picture in question was characterised by Putnam as *moderate realism* about mathematical entities. It is *realist*, because it pictures mathematical entities as *bona fide* entities in Platonic Heaven. It is *moderate*, because it does not allow that humans have supernatural, magical powers which enable them to refer directly to these mathematical entities; indeed, the moderate realist admits that we can pick out mathematical entities only by employing formal theories (if at all). But, such moderate realism is all that is

³²Skolem uses the expression 'absolutely uncountable' in a number of places (1929, p. 272; 1941, p. 468; 1958, p. 635). And this use of the word 'absolutely' has *some* affinity with the purely technical notion of *absoluteness*.

³³Compare Lewis (1984: 225–6) on the Respondent/Challenger dialectic, and Putnam (1994: 284–5) on infinitely regressive scepticism.

required to take us from Skolem's Paradox to the skolemist scenario—just re-read the sceptical speech on p.10—and for skolemism to emerge as an *unanswerable* scepticism.

Our ability to answer the skolemist, then, shows us that moderate realism is a fatally flawed picture. Far from being a plausible starting point for mathematical philosophising, it is either incoherent or reliant upon magic. We learn, in other words, that we must abandon the picture whereby mathematical entities are 'lost noumenal waifs looking for someone to name them.'³⁴

This last line is the conclusion which Putnam himself draws, after arguing at length against moderate realism. However, Putnam criticises moderate realism by deploying his *model-theoretic arguments*. By contrast, I have tried – following Tymoczko and Moore – to translate Putnam's attack on external realism, via the *BIV argument*, into an attack on moderate realism. This suggests a pleasing convergence between Putnam's model-theoretic arguments and the BIV-style of argumentation. I shall explore that convergence in the next section.

A brief digression on indeterminate language. Before that, for the sake of completeness I wish to compare the BIV-style argument against skolemism with a related, but flawed, anti-sceptical argument.

Let us introduce a new kind of sceptic, the *set-indeterminist*. Noting that the 'intended' model and the countable model \mathcal{H} make exactly the same sentences true, the set-indeterminist maintains that they present equally good reference candidates for my set-theoretic vocabulary. In particular, she maintains that there are two equally good reference candidates for my word 'countable': the countable sets on the one hand, and the countable ^{\mathcal{H}} sets on the other. She therefore maintains that set-theoretic vocabulary has indeterminate conditions of application.

Someone might try to offer a purely disquotational argument against the set-indeterminist, as follows:

The word 'countable' applies to all and only the countable sets.

So: the word 'countable' has determinate conditions of application

Harry Field describes this argument as 'a clearer version' of the argument-strategy that Putnam pursues in his 1980. However, Field immediately notes that the word 'bald' applies to all and only the bald things, even though the word 'bald' has *vague* (and hence indeterminate) conditions of application.³⁵ This simple observation elegantly undermines the purely disquotational argument.

Fortunately, Field's observation does not affect my BIV-style argument (which I regard as the 'clearer' reconstruction of Putnam's aims in his 1980). Indeed, set-indeterminism is just as vulnerable to a BIV-style argument as skolemism. I shall briefly explain why.

³⁴Putnam (1980: 482).

³⁵Field (1994: 398–9).

I can coherently say that the people with fewer than n hairs and the people with fewer than $n+1$ hairs are equally good reference candidates for my word ‘bald’. This makes it easy to see why ‘bald’ has vague conditions of application. But the key feature of \mathcal{H} is that some uncountable ^{\mathcal{H}} sets are countable. So I cannot coherently say, as the set-indeterminist wants to, that the countable sets and the countable ^{\mathcal{H}} sets are equally good reference candidates for my word ‘countable’. That would be like saying that the bald people and the non-bald people are equally good reference candidates for my word ‘bald’. It would contradict the disquotational platitude that the word ‘bald’ (respectively, ‘countable’) applies to all and only the bald (respectively, countable) things, which is needed to formulate the sceptical challenge. Otherwise put: the sceptical challenge would be unformulable if it were true, which is just the point of the BIV-style argument.

4 Semantic scepticism and the model-theoretic arguments

Semantic scepticism. Skolemism was motivated by a result from model theory. We can motivate an alternative kind of scepticism via a much simpler result:³⁶

The Permutation Theorem. For any non-trivial structure, there is a distinct, isomorphic structure.

We can understand the main idea behind this result without invoking any fancy model theory. It simply requires a two-step imaginary exercise. First: imagine laying out all the objects in the world, together with various labels (names) for them, and with other labels (predicates) for collections of them (perhaps under a certain ordering). Second: imagine shuffling the objects around, willy nilly, without disturbing the labels (so the labels do not ‘follow’ the objects during the shuffling, but end up applying to new objects). After a moment’s reflection, it should be clear that this process will not affect the truth values of any sentence. (And after another moment’s reflection, it should be clear that the same process would have worked, even if we had started with objects which were *intrinsically featureless*.)

There are plenty of ways to turn the Permutation Theorem into a sceptical challenge. Here is a simple one. Suppose we start by believing that our word ‘cat’ refers to cats. We then apply the Permutation Theorem, shuffling the entities around. Say

³⁶For a brief technical discussion of the result, see Putnam (1981: 33–5, 217–18), Button (2013: 14–16, 225–31; ‘non-trivial’ is defined there) or Button and Walsh (MS: §§2–3). To say that \mathcal{M} and \mathcal{P} are *isomorphic* is to say that there is a structure-preserving one-to-one map between them; it follows that they are elementary equivalent (in the sense explained in fn.23, above).

that *cats** are the entities which we shuffle into the slots which were formerly occupied by cats. (Putnam makes this idea vivid, by pointing out that the *cats** might all just be *cherries*, provided that there are at least as many cherries as cats.) Then the *permutation scenario* is the sceptical scenario that our word ‘cat’ refers to *cats**, and the *permutation-sceptic* is the sceptic who thinks that the permutation scenario is actual.

Permutation-scepticism is not a version of Cartesian scepticism. Far from threatening us with the worry that (almost) all of our beliefs are false, the whole process of permutation is deliberately designed to preserve the truth values of our sentences. In that sense, permutation-scepticism is more like skolemism than like BIV scepticism. However, unlike skolemism, it need not even threaten us with the idea that we are ignorant as to the ‘nature’ of the world. Rather, the worry is *just* that our words do not refer as they are (intuitively) supposed to. Consequently, permutation-scepticism is a third kind of scepticism: a purely *semantic* scepticism.

I have not yet said anything to indicate *why* we might take permutation-scepticism seriously. I shall say a little more on that towards the end of this section. But of course, once we have opened the door to considering one ‘deviant’ permutation of our vocabulary, we should not stop there. Any permutation – and hence any interpretation – will be on the table as a way that our words might refer. And if there is nothing to tell between the rival interpretations, then we are likely to be pushed, from regarding permutation-scepticism as unanswerable, to treating as unanswerable the thought that each of our words refers equally to every thing. Call this last thought *complete semantic scepticism*.

Some philosophers have claimed to learn to live with the unanswerability of complete semantic scepticism. For example, in response to considerations like those above, J. J. C. Smart suggested:

I do not much mind the idea that reference slides freely over the surface of noumenal waters, except that I do not think of the noumenal waters as consisting of unknowable things in themselves: there are just the electrons, protons, stars, cats, cabbages and other perfectly knowable objects.³⁷

Unfortunately, this is literally unintelligible. Smart wants to tell us something about cabbages and cats. Well, which things are the cabbages, and which are the cats? If complete semantic scepticism is true, then Smart says of each thing ‘it is a cabbage’ and ‘it is not a cabbage’, to the same extent. Admittedly, nothing is both ‘a cabbage’ and ‘not a cabbage’ relative to the same interpretation at the same time; but so what? If complete semantic scepticism is right then, when Smart says that the world is made up of ‘electrons, protons, stars, cats and cabbages’, if he makes any assertion at all, he asserts only *that there are some objects*.

³⁷Smart (1995: 309); see also Lewis (1984: 231).

An anti-sceptical argument. Evidently we are once again in the territory of self-refutation. To make this even clearer, we can present an anti-sceptical argument – in the spirit of the BIV argument – against the permutation-sceptic. The argument has a familiar form:

- (1P) If the permutation scenario obtains, my word ‘cat’ does not refer to cats.
- (2P) My word ‘cat’ does refer to cats.
- (3P) So: I am not in the permutation scenario.

Premise (1P) is true by definition, so all of the action is with premise (2P).

One might try to defend (2P) on the grounds that it is as certain as anything. And, of course, (2P) *is* as certain as anything. But to stress *that* point would be to present the argument (1P)–(3P) as a *Moorean* response to the permutation-sceptic. And to do that would be to miss what is really interesting about Putnam’s model-theoretic arguments.

In keeping with the BIV argument and the anti-skolemist arguments, the point is that (2P) is required by the permutation-sceptic herself, in order to formulate her sceptical scenario. She attempts to raise a sceptical concern by saying ‘our word “cat” does not refer to cats’. This presents us with the worry *that* our word ‘cat’ does not refer to cats, only if the last word in the formulation of the worry refers to cats. That is, the very intelligibility of the sceptical challenge requires that our word ‘cat’ *does* refer to cats after all.

Similar moves and counter-moves are likely to be made here, as were made in the discussion of BIV scepticism and skolemism. However, they are no more convincing in this new setting. We thus conclude as follows: *If permutation-scepticism is internal, it is self-refuting.*

Why this matters. The question arises, though, of why we should ever have thought that permutation-scepticism – or, indeed, any kind of semantic scepticism – was worth taking seriously. The answer is that the two philosophical pictures that have already been attacked in this paper – (general) external realism and (mathematical) moderate realism – treat some form of semantic scepticism as *unanswerable*.

It is easy to see that moderate realists must find (some kind of) semantic scepticism unanswerable. Recall that the moderate realist grants no tools for picking out entities, beyond the use of formal theories. But the Permutation Theorem then teaches us that, if we can find any model of a formal theory, we can find many different models. So the moderate realist is left with no tools for picking out any *particular* entities with her mathematical words. Consequently, the moderate realist will find complete semantic scepticism (restricted to mathematical objects) wholly *unanswerable*.

It is rather less obvious that external realists will find some kind of semantic scepticism unanswerable. Indeed, to explain why this is so is a significant project in its own right.³⁸ Here, though, is a very brief pointer in the right direction. Towards the end

³⁸Interested parties should see Putnam (2000) and Button (2013: part A).

of §2, I explained that external realism is characterised by a commitment to the view that even an ideal theory might be radically, hopelessly false. Suppose that we have adopted this picture. Then, intuitively, hypotheses about what my words refer to should be just as amenable to sceptical concerns as hypotheses about what is around me. And this opens the door to various versions of semantic scepticism, including permutation-scepticism.

However, using a BIV-style argument, we have shown that *if permutation-scepticism is internal, it is self-refuting*. And our ability to answer the sceptic therefore shows us that certain pictures – those which treated the scepticism in question as unanswerable – are fatally flawed. More precisely: we have another demonstration that both external realism and moderate realism are either incoherent or magical. That, of course, is the moral which Putnam himself draws from his model-theoretic arguments.

To be clear: Putnam does not explicitly present the refutation of permutation-scepticism (and hence of external realism, for example) in the form of the BIV-style argument (1P)–(3P). However, Putnam evidently regards the external realist's inability to answer the (semantic) sceptical challenges raised by the permutation-sceptic as a *reductio* of external realism. I have argued that this *reductio* is successful. But more than that: I have argued that it is successful for the same reasons as the BIV argument and the anti-skolemist argument, for all three arguments can be accommodated within a single template.

5 Scepticism and philosophical pictures

In this paper, I have defended three related anti-sceptical arguments. I now want to comment on the general structure, not merely of the anti-sceptical arguments, but of the broad dialectic that I have employed.

At a fairly high-level of abstraction, Putnam gave us a step-by-step process for criticising certain philosophical pictures:

- Step 1.* Isolate a particular philosophical picture.
- Step 2.* Observe that some sceptical challenge is unanswerable, given this picture.
- Step 3.* Show that the scepticism in question is actually self-refuting (or reliant on magic).
- Step 4.* Conclude by rejecting the original picture as incoherent (or reliant on magic).

This process unifies some of Putnam's most famous contributions to the realism/antirealism debate from the late 1970s and early 1980s: his discussions of brains-in-vats, of Skolem's Paradox, and of cats and cherries. (And, in passing, it may help to correct a common misreading of Putnam. No one thinks that Putnam ever embraced

BIV scepticism. By exact parity of reasoning, no one should think that Putnam ever embraced skolemism, or radical semantic scepticism.)

In order for the anti-sceptical arguments to figure in Step 3 of this four-step process, it is crucial that the arguments themselves should not depend upon any very substantial philosophical picture. To make this point clear, consider again the way in which the external realist must treat BIV scepticism as unanswerable. Suppose, now, that we could show that BIV scepticism is self-refuting (or magical), but *only* given the additional assumption of some *constructivist* metaphysical picture (say). Then we would not have shown that external realism *itself* is incoherent (or magical). Rather, we would have shown that external realism cannot coherently be combined with constructivism. That would be rather unsurprising, and rather uninteresting.

Fortunately, the anti-sceptical arguments do *not* require that we should adopt any particular philosophical picture. Rather, our anti-sceptical arguments all operate on scepticism *from within*. The lynchpin of all of the anti-sceptical arguments is just that, if the sceptical scenario were actual, then we would be unable to articulate this (barring magic). No particular ontological approach is required to establish this point.³⁹

In using the anti-sceptical arguments during our four-step rebuttal of a certain philosophical picture, we do not, then, need to invoke some *alternative* philosophical picture. Moreover, having overthrown one philosophical picture, we should perhaps pause before putting any alternative picture in its place. Maybe the best way to prevent seemingly-unanswerable-but-actually-self-refuting sceptical challenges from arising is simply to learn how to philosophise in the *absence* of comforting pictures.⁴⁰

For now, though, we have seen a deep connection between Putnam's thoughts on BIVs, on Skolem's Paradox, and on permutations. A simple argumentative template shows us how to defeat certain versions of scepticism and, in so doing, how to overthrow certain philosophical pictures. That is why Putnam's BIV argument matters.⁴¹

³⁹In this respect, my treatment of the anti-sceptical arguments differs significantly from Tymoczko's. Whilst defending the second premises of the BIV argument and the anti-skolemist argument, Tymoczko (1989: 289–90, 296n3) makes several contentious philosophical claims which seem distinctly *constructivist*. He refers to this constructivism as 'internal realism', citing the closing image of Putnam (1980: 482); he might equally have cited the extremely constructivist imagery of Putnam (1981: 52, 54). Tymoczko then acknowledges that the official conclusion of his argument is: 'if internal realism is presupposed by Putnam's [BIV] argument, it is also presupposed by the standard resolution of Skolem's paradox' (1989: 296n3). For my purposes, Tymoczko's conditional conclusion is much too weak. However, this is not to deny something which I hope is obvious: I have benefitted hugely from Tymoczko's paper, both in my understanding of the BIV argument and of the anti-skolemist argument.

⁴⁰That is the lesson I try to draw in Button (2013).

⁴¹For correspondence and discussion, thanks to Tim Bays, Jane Heal, Adrian Moore, Michael Potter, Hilary Putnam, Tim Sundell, Trevor Teitel, Joshua Thorpe, and Rob Trueman.

References

- Bays, T. (2001). "On Putnam and His Models." *The Journal of Philosophy* 98.7, pp. 331–50.
- (2007). "More on Putnam's Models: A Reply to Bellotti." *Erkenntnis* 67.1, pp. 119–35.
- (2009). "Skolem's Paradox." In *Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. URL: <http://plato.stanford.edu/entries/paradox-skolem/>.
- Brueckner, A. (1992). "If I Am a Brain in a Vat, Then I Am Not a Brain in a Vat." *Mind* 101.401, pp. 123–8.
- (1996). "Modest Transcendental Arguments." *Philosophical Perspectives* 10, pp. 265–80.
- Button, T. (2011). "The Metamathematics of Putnam's Model-Theoretic arguments." *Erkenntnis* 74.3, pp. 321–49.
- (2013). *The Limits of Realism*. OUP.
- Button, T. and S. Walsh (MS). "Ideas and Results in Model Theory: Reference, Realism, Structure and Categoricity." URL: <http://arxiv.org/abs/1501.00472>.
- Devitt, M. (1984). *Realism and Truth*. Princeton: Princeton University Press.
- Ebbs, G. (1996). "Can We Take Our Words at Face Value?" *Philosophy and Phenomenological Research* 56.3, pp. 499–530.
- Einheuser, I. (2010). "The Model-Theoretic Argument against Quantifying over Everything." *Dialectica* 64.2, pp. 237–46.
- Field, H. (1994). "Are Our Logical and Mathematical Concepts Highly Indeterminate?" *Midwest Studies in Philosophy* 19, pp. 391–429.
- Klenk, V. (1976). "Intended Models and the Löwenheim–Skolem Theorem." *Journal of Philosophical Logic* 5.4, pp. 475–89.
- Lewis, D. (1984). "Putnam's Paradox." *Australasian Journal of Philosophy* 62.3, pp. 221–36.
- Moore, A. W. (2001). *The Infinite*. Second. London: Routledge.
- (2011). "Vats, Sets, and Tits." In *Transcendental Philosophy and Naturalism*. Ed. by P. Sullivan and J. Smith. Oxford: Oxford University Press, pp. 42–54.
- Moore, G. E. (1939). "Proof of an External World." *Proceedings of the British Academy* 25.
- Mostowski, A. (1969). *Constructible Sets with Applications*. Amsterdam: North-Holland.
- Putnam, H. (1977). "Realism and Reason." *Proceedings and Addresses of the American Philosophical Association* 50.6, pp. 483–98.
- (1980). "Models and Reality." *Journal of Symbolic Logic* 45.3, pp. 464–82.
- (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press.
- (1992). "Replies." *Philosophical Topics* 20.1, pp. 347–408.
- (1994). "Comments and Replies." In *Reading Putnam*. Ed. by P. Clark and B. Hale. Blackwell, pp. 242–95.
- (2000). "Das modelltheoretische Argument und die Suche nach dem Realismus des Common sense." In *Realismus*. Ed. by M. Willaschek. Trans. by B. Brinkmeier. Paderbon: Ferdinand Schöningh Verlag, pp. 125–42.

- Skolem, T. (1929). "Über einige Grundlagenfragen der Mathematik." In Skolem 1970. Ed. by E. J. Fenstad. Oslo: Universitetsforlaget, pp. 227–73.
- (1941). "Sur la Porté du Théorème Löwenheim-Skolem." In Skolem 1970. Ed. by E. J. Fenstad. Oslo: Universitetsforlaget, pp. 455–82.
- (1958). "Une Relativisation des Notions Mathématiques Fondamentales." In Skolem 1970, pp. 633–8.
- (1970). *Selected Works in Logic*. Ed. by E. J. Fenstad. Oslo: Universitetsforlaget.
- Smart, J. J. C. (1995). "A Form of Metaphysical Realism." *The Philosophical Quarterly* 45.180, pp. 301–15.
- Tymoczko, T. (1989). "In Defense of Putnam's Brains." *Philosophical Studies* 57.3, pp. 281–97.
- Velleman, D. J. (1998). "Review of Levin's 'Putnam on reference and constructible sets' (1997)." *Mathematical Reviews* 98c:03015, p. 1364.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wright, C. (1992). "On Putnam's Proof that We Are Not Brains in a Vat." *Proceedings of the Aristotelian Society* 92, pp. 67–94.
- Zalabardo, J. L. (2009). "How I Know I'm Not a Brain in a Vat." *Royal Institute of Philosophy Supplement* 64, pp. 65–88.