

## VERIDICALISM AND SCEPTICISM

BY YUVAL AVNUR

*According to veridicalism, your beliefs about the existence of ordinary objects are typically true, and can constitute knowledge, even if you are in some global sceptical scenario. Even if you are a victim of Descartes' demon, you can still know that there are tables, for example. Accordingly, even if you don't know whether you are in some such scenario, you still know that there are tables. This refutes the standard sceptical argument. But does it solve the sceptical problem posed by that argument? I argue that it does not, because we do not know substantively more about the external world according to veridicalism than we do according to the sceptical argument. Rather, veridicalism merely reformulates what little knowledge we have. I then draw some general conclusions about the nature of the sceptical problem, the formulation of the standard argument, and the significance of this for some other, non-veridicalist strategies.*

**Keywords:** epistemology, scepticism, structuralism, metaphysics, ordinary objects.

### I. INTRODUCTION

Veridicalism is the view that even if you are in a global sceptical scenario, many of your ordinary beliefs about the external world are still true and can constitute knowledge. It is often taken by its proponents as a decisive solution to the sceptical problem. For years, and in various works, David Chalmers has championed this view based on structuralism about physical concepts (and properties).<sup>1</sup> However, there are other versions of veridicalism: Davidson (1986), Russell (1927), Valberg (2007), and arguably Leibniz (1923: 6.4: 1502–3).<sup>2</sup> This paper critically discusses the veridicalist strategy in order to draw some more general lessons about the problem of scepticism and its formulation. I grant veridicalism (and structuralism) for the sake of argument, and show that, though it might refute the letter of the now standard formulation of

<sup>1</sup> Chalmers (2005, 2012: 431–40, 2018, 2022).

<sup>2</sup> Bouwsma (1949) was also a proponent of non-structuralist veridicalism, discussed in Chalmers (2018: 627–9).

the sceptical argument, it does not solve the sceptical problem posed by that argument. Though my objections apply to all, even non-structuralist forms of veridicalism (as well as some denials of closure), I will often use Chalmers (2018) as an exemplar since it is the most current version in the literature. The positive lessons drawn about the nature of the sceptical problem shed light on other anti-sceptical strategies as well.

Focus on ‘global’ sceptical hypotheses, which purport to cast doubt on all empirical beliefs, rather than, say, only those about our current environment, or only about the past. For example, that you are and have been from birth a brain in a vat (BIV) and that you are and have always been a victim of Descartes’ demon are global sceptical hypotheses. Call your favourite such hypothesis ‘sk’. Here is a sceptical argument that has been a mainstay of the epistemology literature at least since Barry Stroud (1984):

- (1) I do not know that  $\sim$ sk.
- (2) If I do not know that  $\sim$ sk, then I do not know that there are tables.

Therefore,

- (3) I do not know that there are tables.

Generalising from tables,

- (4) I do not know that there are any ordinary physical objects.

Call this ‘the standard argument’. The conclusion, (4), has seemed unacceptable to most philosophers. Accordingly, much of the literature on the topic has focused on explaining which of (1) or (2) is false. *Veridicalism* targets (2) while accepting—or at least not disputing—(1).<sup>3</sup>

Chalmers’s route to veridicalism is via *structuralism* about physical claims: Our beliefs about the external world are equivalent to structural claims, and sk has the right structure to make those beliefs true. The details of his brand of structuralism will not matter here. Even if sk is true, there is something causing your table experience, and that is a table. If you are a BIV, then the table is ultimately digital, a computer state in the machine running the vat. If you are a victim of Descartes’ demon, the table is ultimately a mind-dependent, demonic idea. But either way, there is a table here. So, sk is a hypothesis about the ultimate nature of the table, or what ultimately plays the causal role of the table. Accordingly, even granting (1), (2) is false because there are (knowable) things such as tables regardless of whether sk is true.

<sup>3</sup> In contrast, semantic externalists like Putnam (1981) target (1) instead of (2). Another departure from Putnam’s strategy is that veridicalists like Chalmers accept what he calls ‘non-twin-earthable’ terms in (2018) (in 2005: 12 and 18, he called them ‘semantically neutral terms’) with which to conceive of and consider sk as a real possibility. One advantage of veridicalism over Putnam’s strategy is that veridicalism does not bear the burden of proving a priori that you are not a BIV.

There are other ways to arrive at veridicalism besides structuralism. Davidson's (1986) idea is that a correct interpretation of our belief that there is a table, even if *sk* is true, would make that belief true and presumptively constitute knowledge (236). However, the same cannot be said for the belief that  $\sim sk$ , which is not the sort of simple belief that has 'special' interpretive weight (237). So (2) is false, although (1) may be true. And Valberg (2007) holds that things that appear within, say, a dream are simply objects in a dream, so that if you are dreaming, your belief that there is a table is true (and presumably known), since there is a table in the dream (30, 85, 93). So, Valberg, too, rejects (2) even on the assumption that (1) (which he later disputes for other reasons).

In what follows, I grant veridicalism for the sake of argument, and assume that Chalmers and the others succeed in refuting (2), so that the standard argument is unsound. This still does not solve the sceptical problem posed by the standard argument, because it fails to show that we know substantially more about the world than (4) implies.

## II. OBJECTIONS TO VERIDICALISM AS A SOLUTION

A valid argument with plausible premises poses a philosophical problem if its conclusion implies something worrying or otherwise unacceptable from a philosophical perspective. To try to solve the problem, one might refute the argument. However, if such a refutation implies that same, worrying thing, then one has not thereby solved the problem. Rather, one has arrived at the same problem via a different commitment. In this case, if one refutes the standard argument ((1)–(4)) and thereby avoids (4), but only by committing to something as worrying as (4) was originally, one has not thereby solved the problem raised by the standard argument. Granted, what makes something worrying is a slippery matter, and some philosophers might think that the standard argument never presented a problem in the first place. However, in this section, I argue that *if (4) was worrying in the first place*, then the veridicalist rejection of (2) does not solve the problem raised by the standard argument, because it raises the same worry.

What is worrying about the conclusion of the standard argument, (4)? The ignorance it posits implies that your cognitive life, in relation to the world, is a joke: for all you know, you are thinking about the world and your place in it completely wrong. You do not even know that there are tables, which is a nightmarish relation to have to reality.<sup>4</sup> (4) makes explicit just how little

<sup>4</sup> As Stroud (1984: 38) points out, the worry also involves your lack of knowledge of the existence of other people. Though the existence of other people is a central issue for those who worry about classic sceptical scenarios, Chalmers (2018: 625–6) leaves scenarios in which there are no other minds aside.

information you have about the world by extracting, from the premise that your information is insufficient to rule out *sk*, the conclusion that you do not even know the simplest thing about your surroundings. As Stroud (1984: 1) formulated it, scepticism is the view ‘that no one knows anything about the world’, and that would be deeply disturbing. If, as (4) states, you do not even know that there are any ordinary objects around you, you do not know much ‘about’ the world.

Following are some reasons to think that veridicalism, even if it avoids (4), commits to a state of ignorance about the world as worrying as (4), leaving you with just as little knowledge ‘about’ the world.

### *II.1 Intuitive equivalence*

Here is an initial intuition. Insofar as the standard argument presents a problem it is by supporting this, ‘standard conclusion’:

I do not know whether there is a table here, because I might be a BIV, a demon victim, etc.

Veridicalism avoids the standard conclusion, but only by committing to this, ‘altered conclusion’:

I do not know whether this table is a BIV-table, a demon-table, etc.

That is, veridicalism, since it does not reject (1), commits to the claim you might be, say, a BIV. And if you are a BIV, then this is a simulated table, something that is ultimately a computer state in a BIV machine. Intuitively, it would be strange to find the standard conclusion worrying while not finding the altered conclusion worrying. For, they seem to express the same basic ignorance about the world. The altered conclusion seems like a mere re-formulation of, rather than an improvement over, the standard conclusion.<sup>5</sup> For those who do not share the intuition, consider that since they both accept (1), the standard sceptic (i.e., one who accepts (1)–(4)) and the veridicalist accept this, ‘basic conclusion’:

I do not know whether the thing causing my table experience is a BIV machine, a demon, etc.

Once we notice that veridicalism does not avoid the basic conclusion, it becomes hard to feel relieved by swapping the standard for the altered conclusion. For, the basic conclusion is the root of what is worrying about the standard conclusion: You do not know about the world behind its sensory effects on you, or the external world. It was only due to your assumption that

<sup>5</sup> Craig (1990: 213–4) gives a similar objection to Davidson’s veridicalism.

simulated tables (as they appear to a BIV) are not tables that you expressed this worry as the standard conclusion about the existence of tables.

One might object that at least according to veridicalism, we have saved your ordinary beliefs from the sceptical problem, since beliefs like that there is a table can still be knowledge.<sup>6</sup> However, there are two problems with this idea. First, it is only relatively recently that scepticism has been formulated so as to focus primarily on the knowledge status of our ordinary beliefs. Much of the history of scepticism has been more clearly focused on what I call the ‘basic’ sceptical conclusion above. Certainly, this is what David Hume<sup>7</sup> and George Berkley<sup>8</sup> regarded as the sceptical problem. This suggests that the standard conclusion is only one, recent formulation of scepticism, rather than the core of the problem. Second, the reason that the failure of ordinary beliefs to be knowledge is disturbing is that it suggests that we know very little about the world. If it turns out that veridicalism fails to establish that we know substantively more about the world than we do according to (4), it will have failed as a solution to that problem. In that case, it offers only a cosmetic or superficial solution by vindicating the status of ordinary beliefs without vindicating what we know *about* the world. This is what I will argue is the case in the next subsections.

## II.2 *Same possibilities*

For convenience, focus for now on the Chalmers-style structuralist, who holds that a demonic table is still a table because it plays (roughly) the same causal role as a normal table. This kind of veridicalist and the standard sceptic leave you ignorant about the same set of possibilities for what might be going on around you. Both agree that you know that there is either a demon-caused table appearance, or a BIV-caused table appearance, or a ‘non-sk-caused’ table appearance, and so on. The standard sceptic merely fails to count knowledge of that existentially quantified disjunction as knowledge of the existence of a *thing*, a table, whose existence conditions are so vastly disjunctive. Veridicalism teaches us to apply the concept ‘table’ to whatever satisfies the disjunction in

<sup>6</sup> David Chalmers, personal correspondence.

<sup>7</sup> ‘By what argument can it be proved, that the perceptions of the mind must be caused by external objects, entirely different from them, though resembling them (if that be possible) and could not arise either from the energy of the mind itself, or from the suggestion of some invisible and unknown spirit, or from some other cause still more unknown to us? It is acknowledged, that, in fact, many of these perceptions arise not from anything external, as in dreams, madness, and other diseases....It is a question of fact, whether the perceptions of the senses be produced by external objects, resembling them: how shall this question be determined?’ (*Enquiry* Part 12, Section I).

<sup>8</sup> ‘Hence arise *skepticism* and *paradoxes*. It is not enough that we see and feel, that we taste and smell a thing. Its true nature, its absolute external entity, is still concealed’ (Berkeley and Mathias 2007, *preface*).

our actual situation, but it does not put us in a position to rule out any of those disjuncts. This suggests that the veridicalist merely reformulates what little information we have about the world according to the standard sceptic. If what worried you about (4) was that it implied that you have very little information about the external world, how could veridicalism assuage that worry if you are not in a position to rule out any more of those disjuncts, or any more of the possibilities expressed by those disjuncts, than the standard sceptic claimed?

One might object that structuralist veridicalists can rule out the possibility that there are no tables. However, the standard sceptic can rule out the same set of possibilities, even if those possibilities are described differently. The possibilities in which there are no tables according to structuralism are those in which there is nothing playing the causal role that we take tables to play; the standard sceptic also knows that we are not in any such scenario, because she knows that something is playing a table-like role. The difference is that the sceptic does not describe those other possibilities, which you cannot rule out, as ones in which there are tables. However, what matters, when it comes to what is going on in the world, is which possibilities might be actual, not how those possibilities are described. Whether to call whatever is causing our experiences ‘tables’ is a paradigmatically *empty* question in this respect: not a question about which scenario you are in, but how to describe that scenario.<sup>9</sup> As I argue next, this concerns how to formulate your information; it is not a matter of additional information about that which causes our experience or about the external world.

### II.3 *Clever neologisms*

The altered conclusion is no better than the standard conclusion, with respect to how good your knowledge about the world is, because by merely inventing some clever terms, the standard sceptic can know as much about the world as the veridicalist claims to know. The standard sceptic can still know that *something* or other is causing her table experience, and that whatever it is has a causal profile similar enough to what we usually think tables have.<sup>10</sup> The only difference is that the veridicalist claims that whatever the thing that causes the experience is, it falls within the extension of ‘table’. However, the standard sceptic could come up with a clever term, ‘tabby’, which she stipulates to apply to the thing that causes her table experiences, whatever it is, as long as it has the causal role of a table. She has thereby turned knowledge that something caused her table experience into knowledge that there are tabbies. So, she

<sup>9</sup> The notion of an empty question is due to Parfit (1984: 235).

<sup>10</sup> Other sceptics might deny that there is even any causation, but veridicalists are considering a standard sceptic who accepts (1)–(4) but accepts that *something* causes her experiences.

accepts (4), and yet now she knows as much about tabbies as the veridicalist claims to know about tables.

Knowledge that there is a table and knowledge that there is a tabby amount to the same information about the world: They exist in all the same possible (global) scenarios, and regardless of whether *sk* is true. Veridicalists hold that the clever standard sceptic has a false theory of tables, as we ordinarily think of them, because tables are tabbies, even though this sceptic says that *sk* is incompatible with tables and compatible with tabbies. However, that is a disagreement about what counts as a table; there is no substantial difference between their claims to knowledge about what is going on around them, or the external world, given that they both accept (1) and the basic conclusion. There is nothing of substance that the veridicalist claims to know that the clever, standard sceptic fails to know about the external world. They merely use different terms—‘table’ vs ‘tabby’—to express that knowledge.

One might object that veridicalism saves the beliefs we all always had about the world; we all believed in tables, not tabbies. However, first, on some views of content, the belief that there are tabbies is the same as the belief that there are tables, since they exist in all the same possibilities. And second, recall that whether to call that which causes our experience ‘tables’, rather than ‘tabbies’, is an empty question. There is not a *substantive* question about the world that veridicalism can answer that the clever standard sceptic, who merely invents new terms, cannot.

#### *II.4 The score does not matter*

Recall that veridicalists and standard sceptics both accept the basic conclusion: that you do not know whether the thing causing your table experience is a BIV machine, a demon, etc. However, according to veridicalism, a greater quantity of our ordinary beliefs is knowledge, even though nobody knows what ultimately causes our experiences. The amount of true, knowledge-constituting ordinary beliefs (not counting those gained by clever neologisms like ‘tabby’) is the *only* difference between what veridicalism and standard scepticism posit. Call the matter of how many more of our ordinary beliefs constitute knowledge according to veridicalism (over standard scepticism) the ‘score’. My fourth and final objection is that the score, first, is not so straightforward, and, second, is irrelevant to the worry (4) raises.

First, counting beliefs is notoriously a slippery thing, where one belief becomes a kaleidoscope of millions. You believe that there is a table, a table in front of you, and also a table closer than 10 feet from you, closer than 12 feet from you, etc. How many beliefs to count there is not obvious. This muddies the veridicalist’s objection, because there is a similarly indefinite number of things you also believe, but *do not* know concerning the table according to veridicalism: that the table is not a demonic idea, that machines which have

enslaved humanity did not cause the table to exist, that they did not cause it to exist while running low on batteries, while being powered by an even number of solar panels, and so on. You are at least *disposed* to believe such things about the table. Notice that the standard sceptic does not hold that there are tables, as ordinarily conceived, in any of those scenarios (she is not a veridicalist), so she is not disposed—and does not think that we are ordinarily disposed—to believe such things. So, she commits to fewer such beliefs that do not constitute knowledge on that count. Such beliefs do not often come up in ordinary life, and are seldom said out loud, but neither does ‘there is a table in front of me.’ So the score is unclear.

To be clear, the above was a comment on the comparison between beliefs according to veridicalism and beliefs according to the standard sceptic. I am assuming veridicalism is true for the sake of argument. The point is that, even if veridicalism is true, the *number* of knowledge-constituting beliefs is not clearly greater than according to the standard sceptic’s non-veridicalist view, so it isn’t clear that veridicalism gets the better ‘score’.

One might object that beliefs such as that there is a table play a role in guiding one’s behaviour and thinking, while the belief that the table is not, ultimately, a demonic idea does not. So, the score should be adjusted in the veridicalist’s favour. However, first, if you thought that you may well be in a vast simulation, this might affect how you think about your life; we are not *generally* indifferent to sk. Admittedly, this is not a matter of navigating around furniture, but why should navigation be the only benchmark of importance? The veridicalist anti-sceptical strategy itself does not provide an answer, since it is not a sort of pragmatism about the value of knowledge.<sup>11</sup> That would be an entirely different anti-sceptical manoeuvre. Moreover, recall that the standard sceptic can count on beliefs about tabbies to help navigate around furniture.

Secondly, and independently, the score does not matter. ‘What’ we know about the world is not the same as ‘how much’ we know about the world. ‘How much’ we know depends on how we are counting or how we formulate what we know.<sup>12</sup> True beliefs are easier to come by when their contents are more broadly construed or more easily satisfiable, but that is not due to a better informational situation or even shifting epistemic standards (as is the case with contextualist solutions to scepticism). Rather, it is merely a reinterpretation of what our beliefs are about. What was worrying about (4) is that it implies that you know very little *about* the world, and this is not equivalent to *how many*

<sup>11</sup> Though Chalmers (2022) does give various considerations in favour of the idea that most of what matters in your life is indifferent to whether you are living in a global simulation. I lack space to discuss the merits of this claim.

<sup>12</sup> Relatedly, Treanor (2014: 552–9) has pointed out that what looks grammatically like ‘one’ fact can often contain a conjunction of many.



*things* you know. Besides, as we have seen, stipulating clever new terms affords the sceptic the same quantity of empty items of knowledge.

### II.5 *The dialectical situation*

Since I have granted veridicalism for the sake of argument, it will be helpful to note that none of the above arguments depends on veridicalism being *false*. Rather, they indicate that veridicalism fails to solve the problem posed by (1)–(4) because it fails to posit better knowledge about the external world than standard scepticism, or (4), does. In short, this is because veridicalism is still committed to what I called the *basic sceptical conclusion*: I do not know whether the thing causing my table experience is a BIV machine, a demon, etc.

Consider the claim that there is a table. According to veridicalism, this is compatible with *sk*. According to the more standard metaphysics assumed by premise (2) of the standard sceptical argument, that there is a table entails  $\sim sk$ . For clarity, let us call tables as the veridicalist understands them ‘*thin* tables’, because the content of the claim that there are tables is relatively thin in not automatically implying  $\sim sk$ . And let us call the more standard understanding of tables, on which there being a table is incompatible with *sk*, ‘*thick* tables’. In Sections II.1–II.4, I did not question whether tables as we ordinarily think about them are thin. So, I have not questioned whether you know that there are tables despite not knowing whether *sk*. Rather, I argued that knowing that there are thin tables does not imply knowing more about the external world than (4) when (1)–(4) are interpreted to be about thick tables. This is because knowing that there are thin tables is compatible with the basic sceptical conclusion, or ignorance about what is causing the table appearance out there, in the world external to appearances (regardless of whether we call that a ‘table’).

One might here object: Why is it a problem that we don’t know more than the sceptic, given that we know that there are *thin* tables, and given that that’s all we ever thought we knew (granting, as we are, veridicalism about tables)? Why not conclude from this, instead, that there was never any real sceptical problem to begin with? It may have *seemed* that our knowledge was deficient because we were incorrectly thinking of tables as thick tables. If our ordinary claims are about thin tables, then the sceptical conclusion that we do not know that there are thick tables is no big deal. Doesn’t veridicalism accomplish at least this dissolution of the sceptical problem?<sup>13</sup>

There are two kinds of reasons to reject this thought. The first is dialectical: The purpose of the arguments offered above is to show that *if* (1)–(4) are problematic, veridicalism doesn’t help. This is compatible with the negation of that antecedent, that (1)–(4) is not problematic in the first place. According to that view, the basic sceptical conclusion does not constitute any serious

<sup>13</sup> Thanks to an anonymous referee for pushing this objection.

problem: It doesn't matter much whether you know that there are thick tables, or whether you know that what causes the table appearance is a demon, a computer simulation, and so on. This seems implausible. It is hard to imagine ever being concerned about (1)–(4) if one was never concerned about what was causing the table appearance, external to our experiences. And it seems disingenuous to go to the trouble of offering a solution to the sceptical problem when one's view is that there is no coherent problem to begin with. We can sum up this point as a dilemma. Either ignorance about the existence of thick tables is philosophically problematic or not. If it's not, then there's no problem for veridicalism to solve anyway—it is clear that a sceptic who accepts (4) thinks that *something* is causing the table appearance and has a table-like causal profile, so the sceptic accepts the existence of thin tables. However, if ignorance about the existence of thick tables *is* problematic, then surely positing knowledge of thin tables that veridicalism posits isn't going to help. Either way, there's no problem here that veridicalism can solve.

The second reason to reject the thought above is more substantial: The claim that we never took ourselves to know that there are thick tables is, at best, misleading, even if veridicalism about what we ordinarily call 'tables' is true. Suppose you used to worry because you thought you lacked knowledge of whether there were tables out there, beyond mere table appearances. However, now that you've learned about veridicalism, and that you know that there are thin tables, which is what you apparently meant by 'table' all along, you *still* have no idea what is out there beyond the table appearance (a demonic idea? A computer simulation?). What changed is merely how you express the ignorance from before. It is compatible with this story that what worried you in the first place is what is in the external world, and particularly whether the external world is a demon, a computer, or a natural environment as you'd always assumed. You worried because you didn't know whether there was a thick table, but mis-expressed the worry as a worry about whether there is a table, which veridicalism teaches is a thin table. Rather than showing that we never cared whether we know that there are thick tables, all veridicalism shows is that any such worry cannot be expressed as a worry about the existence of tables. Instead, it must be expressed as a worry *about* tables: Are they ideas, simulations, or something else?

So, it seems incorrect or at least unmotivated (by veridicalism) to say that *all* we ever worried about was the existence of thin tables. A veridicalist can coherently worry about what causes the table appearances; she just can't express this as a matter of the existence of tables. This is because we have beliefs, or dispositions to believe, that are triggered by sceptical musings, in things beyond the bare existence of ordinary objects. You are typically disposed to believe that the table is not a demonic idea, even if this is not correctly expressible as a belief about the existence of a table. If you don't know whether sk, then that belief is not knowledge. The question now is whether this is an

acceptable kind of ignorance, and the veridicalism I've granted for the sake of argument does not settle this. In the next section, I consider an attempt on behalf of the veridicalist to argue that ignorance of whether the table is a demonic idea is less significant than ignorance of whether there is a table.

Let us sum up the dialectical situation. I have argued that the knowledge veridicalism posits is not substantially better than the knowledge the standard sceptic posits. One might insist that the paucity of knowledge the standard sceptic posits, once correctly understood according to veridicalism, is not problematic in the first place. However, first, this is to deny there is a problem to solve, rather than to solve a problem. And, second, veridicalism itself, as a theory about the content of our ordinary beliefs about the existence of physical objects, does not show that the ignorance implied by (1) is insignificant, even given veridicalism. For, it leaves us ignorant of what causes our table experiences, and this can matter even if it is not expressed as an ordinary belief about the existence of a table.

### III. SCEPTICISM VS HUMILITY

In this section, I consider a reply on behalf of the veridicalist. I have argued that veridicalism affects only the form, and not the substance, of our knowledge of the external world. However, perhaps the form makes a difference after all. As Valberg puts it, ignorance of whether this is a dream-table amount to ignorance about that in which the table appears. Similarly, Chalmers emphasises that ignorance about sk amounts to ignorance about the ultimate nature of tables. This formulation of our ignorance might be thought to diminish the sceptical problem since ignorance about that in which tables appear, or the ultimate nature of tables, seems like an acceptable humility, rather than an epistemic disaster. So, doesn't veridicalism successfully address the sceptical problem posed by (1)–(4) after all?

For illustration, consider that the altered conclusion can be compared to scientific ignorance of what lies behind the manifest image. Some physicists speculate that physical objects may ultimately be digital, and others that the observable universe may exist within a black hole of a larger one (Chalmers 2005: 5–6, 10). Similarly, in the past, we did not know whether water is made of H<sub>2</sub>O, or whether it is instead fundamental (as Thales thought). These are items of ignorance about the ultimate nature of things, but they are not disturbing sceptical conclusions. Rather, they constitute acceptable humility about nature. Does the altered conclusion amount merely to more such humility? If so, then we should not worry about it any more than we do about mere scientific humility.

However, there is a significant difference between ignorance about sk and scientific humility. The considerations that show that one currently does not

know some scientific hypothesis about the nature of water concern the current state of one's empirical investigation, and this is compatible with one's knowing some *other* facts about the nature of things, at least at some other time and in other circumstances. In contrast, ignorance of whether *sk* is true implies that one can never, even in principle, know the nature of tables and water, other than that they are whatever explains various features of our experience. For, the empirical data required to know which hypothesis about the nature of the table is correct is presumably undermined by the failure to know whether, say, evil machines are generating all the empirical data in order to manipulate us. The Matrix robot overlords could easily mislead you to think that the table is non-computational by providing the requisite misleading evidence. Whatever scientific evidence you claim to have about the nature of things, presumably robot overlords, BIV machines, or the demon could have planted that evidence to mislead you. If you are ignorant of *sk*, then science, when it comes to the nature of things, is not just incomplete, it is hopeless.

It follows that ignorance about *sk* is not just as acceptable as scientific humility about the ultimate nature of things. For what makes scientific humility acceptable is that investigation is ongoing, that the scientific enterprise at any given time (at least before it becomes 'complete') cannot be expected to tell us *everything* about the fundamental nature of things. However, if we are ignorant of *sk*, then the immaturity of the scientific enterprise is irrelevant, since we can never know what things ultimately are. Science cannot be expected to tell us *anything* about the nature of things. This is worse than the truism that science cannot tell us *everything*.<sup>14</sup>

However, perhaps it is not the job of science to reveal the ultimate nature of things anyway. If the veridicalist thinks that all we can ever know about are 'thin' objects (see the previous section), then it is natural for her to also think that the realm of science is limited to thin objects and their interactions. If so, though, it cannot be that scientific humility is equivalent to ignorance about what lies behind appearances, or whether things like tables are demonic ideas, because science does not deal with what lies behind appearances or whether tables are demonic ideas. Perhaps, instead, that is the job of metaphysics. Is the altered conclusion as acceptable as metaphysical humility about the nature of things? Consider panpsychism. Most of us do not worry, as we might about (4), about the idea that we do not know whether panpsychism is true, so that all objects are mind-dependent (Chalmers 2018: 645–60). So why should we

<sup>14</sup> One might think that on structuralist veridicalism, you can still do some science, even given (1). You can study the hidden causes explaining photosynthesis or the stars, even though plants and stars may turn out to ultimately be simulations. This is questionable, because a clever BIV scientist or demon could presumably thwart those scientific efforts by producing misleading empirical data. Regardless, what I argue here is merely that you cannot pursue the scientific study of the ultimate nature of things, and this leaves much other science aside.

worry that we do not know whether the table is mind-dependent because it is a demonic idea?

As in the case of scientific hypotheses, our inability to know whether panpsychism is true does not imply our inability to know *anything* about the ultimate nature of things, or *anything* about the world beyond what the standard sceptic knows (as argued in the previous section).<sup>15</sup> So panpsychism cannot serve as an instance of sk, since the epistemic implications of our ignorance about it are more modest. Still, panpsychism does clash with some of our basic assumptions about the world (arguably), and so our (purported) ignorance about it does demonstrate *some* inability to know about the world. Perhaps there is no bright line, but rather a sliding scale measuring how much some ignorance undermines our knowledge about the world. Ignorance of panpsychism is compatible with much more knowledge about the world than (4) allows, so it is not near on the scale to any of the classic sk's.

We can grant, then, that not *all* ignorance about the nature of tables raises sceptical worries. This should not be surprising. Consider an analogy. Learning that there are other minds while remaining ignorant of whether minds are physical or non-physical would solve the problem of other minds. So ignorance of whether minds are physical does not itself raise the problem of other minds. In the same way, ignorance of whether panpsychism is true does not itself raise the sceptical problem. However, learning that there are other minds while remaining ignorant of whether minds are conscious would *not* solve the problem of other minds. This is because the root of the problem is that you do not seem to know whether there is consciousness like yours out there. This is expressed as a worry about *minds* only because it is assumed that minds are conscious. And the view that minds might not be conscious undermines this assumption. So, you can no longer express the worry as one about minds, but as long as you are ignorant of whether there is consciousness like yours out there, you have failed to address the problem, even if you've reformulated it. Ignorance of whether minds are conscious disqualifies the existence of minds as a solution to this problem. In just the same way, learning that there are tables while remaining ignorant of whether tables are demonic ideas or BIV simulations does not solve the sceptical problem. The root of that problem is that you do not know whether that which causes your experiences—the external world—is a demon or a BIV. This is expressed as a worry about tables only because it is assumed that things like demonic ideas are not tables. Ignorance of whether tables are demonic ideas disqualifies the existence of tables as a solution to the problem because it undermines the assumption operative in expressing the problem in terms of the existence of tables.

<sup>15</sup> Lewis (2009: 203–22) held that we can never know the realisers of causal roles in the world, and considered this to be acceptable humility. However, as Langton (2004: 134) points out, it is hard to see this as a non-sceptical view unless we combine it with something like Lewis's contextualist solution to scepticism.

Knowing there are tables but not whether panpsychism is true is (arguably) mere humility; knowing there are tables but not knowing whether they are demonic ideas is just a reformulation of scepticism.

#### IV. GENERAL UPSHOTS

There are three ways in which the considerations offered above go beyond criticism of the veridicalists' efforts.

First, the objections above apply to any anti-sceptical strategy that rejects (2) without addressing (1). This includes some denials of closure. Those denials of closure that do not require the subject to assume, believe, or otherwise trust that *sk* is false will be subject to all the same objections.<sup>16</sup> This is because they are committed to the altered conclusion that you do not know whether the table is virtual, demonic, etc.

Secondly, the standard argument has been shown to be inadequate, because refuting it does not necessarily solve the problem it poses. The sceptical problem is that, according to (1), we know very little about the world. Whether this is a problem concerning the *existence* of things rather than their ultimate nature depends on the conditions for things' existence, and that is a metaphysical question. So the nature of the problem depends on some metaphysics, and that is an important insight. However, we cannot solve the problem merely by settling what kind of problem it is. Either way, we know very little about the world if (1) is true.

Accordingly, and third, a satisfying solution to the problem posed by the standard argument must address (1), our apparent ignorance of whether we are in some global sceptical scenario.

#### ACKNOWLEDGEMENT

For discussions and comments, thanks to Dominic Bailey, David Chalmers, John Fischer, Adam Pautz, Karl Schafer, Eric Schwitzgebel, Jonathan Vogel, and Rivka Weinberg. An early version of this paper was presented in 2017 at the Claremont Colleges 'Work in Progress' seminar and the Skepticism Conference at UC Irvine under the title 'Unhinged Skepticism', in 2018 at a UC Riverside workshop, and more recently at Central and Pacific APAs. Thank you to all those commentators and audiences, and to the anonymous referees of this journal.

<sup>16</sup> For example, Dretske (1970) and Nozick (1981). In contrast, Avnur (2012) and Coliva (2015) incorporate rejection of *sk* on the part of the subject in their denial of closure, so the objections here do not straightforwardly apply to their view.

## REFERENCES

- Avnur, Y. (2012). 'Closure Reconsidered', *Philosophers' Imprint*, 12. <https://quod.lib.umich.edu/cgi/p/pod/dod-idx/closure-reconsidered.pdf?c=phimp;idno=3521354.0012.009;format=pdf>.
- Berkeley, G. B. and Mathias, M. B. (2007). *George Berkeley: Three Dialogues between Hylas and Philonous*. USA: Routledge.
- Bouwisma, O. K. (1949). 'Descartes' Evil Genius', *Philosophical Review*, 58: 552–9.
- Chalmers, D. J. (2005). 'The Matrix as Metaphysics', in C. Grau (ed.) *Philosophers Explore the Matrix*. Oxford: OUP.
- . (2012). *Constructing the World*. Oxford: OUP.
- . (2018). 'Structuralism as a Response to Skepticism', *Journal of Philosophy*, 115: 625–60.
- . (2022). *Reality+: Virtual Worlds and the Problems of Philosophy*. New York: W. W. Norton.
- Coliva, A. (2015). *Extended Rationality: A Hinge Epistemology*. London: Palgrave Macmillan.
- Craig, E. (1990). 'Davidson and the Sceptic: The Thumbnail Version', *Analysis*, 50: 213–4.
- Davidson, D. (1986). 'A Coherence Theory of Truth and Knowledge', in E. LePore (ed.) *Truth and Interpretation. Perspectives on the Philosophy of Donald Davidson*, 307–19. Oxford: Blackwell.
- Dretske, F. (1970). 'Epistemic Operators', *Journal of Philosophy*, 67: 1007–23.
- Langton, R. (2004). 'Elusive Knowledge of Things in Themselves', *Australasian Journal of Philosophy*, 82: 129.
- Leibniz, G.W. (1923). *Sämtliche Schriften und Briefe*. Berlin: Akademie-Verlag, 6.4: 1502–150.
- Lewis, D. K. (2009). 'Ramseyan Humility', in D. Braddon-Mitchell and R. Nola (eds.) *Conceptual Analysis and Philosophical Naturalism*, 203–22. Cambridge, MA: MIT Press.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: OUP.
- Putnam, H. (1981). *Reason, Truth and History*. Cambridge: CUP.
- Russell, B. (1927). *The Analysis of Matter*. London: Kegan Paul.
- Stroud, B. (1984). *The Significance of Philosophical Scepticism*. Oxford: OUP.
- Treanor, N. (2014). 'Trivial Truths and the Aim of Inquiry', *Philosophy and Phenomenological Research*, 89: 552–9.
- Valberg, J. J. (2007). *Dream, Death, and the Self*. Princeton, NJ: Princeton University Press.

*Scrpps College, USA*