

The Skeptic and the Veridicalist

1

Although the whole of this life were said to be nothing but a dream, and the visible world nothing but a phantasm, I should call this dream or phantasm real enough, if, using reason well, we were never deceived by it.

Leibniz (1923)

1 Introduction

Skepticism about the external world has been formulated in many ways through the centuries. It has been taken to concern not only knowledge, but justified belief, rational degrees of confidence, doubt, and certainty; it has taken the form of a theory, a challenge, a paradox, a way of life, and an invitation to doubt. These days, skepticism is usually associated with the question,

Can you know that external, ordinary objects such as tables exist?

This concerns what there is, or what things there are. Sitting at your table, you look down and see its surface, and feel its pressure under your hand. Is the table really there, or is it part of a vast computer simulation, or perhaps an idea in the mind of Descartes' demon? If you conclude that you don't know whether some such skeptical scenario is playing out, then you are supposed to conclude that, therefore, you don't know that there is a table there. This is nothing special about tables, so the conclusion generalizes: you don't know whether any external objects exist. This standard formulation ignores the question of what things are, or the ultimate nature of the things whose existence is in question. The nature of things such as tables is reasonably taken to be a question of metaphysics or science. In fact, skepticism calls into question whether the metaphysics and science of ordinary objects even have a known subject matter! The possibility that you are in a simulation or a victim of Descartes' demon is an *epistemological* problem about what there is, and metaphysical questions which presuppose the existence of objects such as tables depend on a solution to that problem.

But notice that this standard formulation of skepticism is not entirely metaphysically innocent. There is a metaphysical conception of external world objects already at play, and required for the skeptical reasoning to get off the ground, to the effect that tables are not the kinds of things that could exist in a simulation, or in the mind of a demon. And accordingly, one might try to stop skepticism in its metaphysical tracks by rejecting this conception of objects. Berkeley famously proposed that the external world is ultimately ideal, or made up of objects that are themselves ideas about which there can be no serious or even coherent skeptical worry. This would solve the problem of skepticism, perhaps. But this sort of idealism seems too extreme to most of us today, perhaps even more extreme than skepticism itself. Instead, in this Element we will explore a different metaphysical, antiskeptical strategy, one that rejects some standard metaphysical assumptions about the external world, but stops short of idealism.

Setting aside Berkeley's idealism until Section 2, consider this alternative metaphysical view: If this is all a vast simulation, then tables still exist. In that case, tables turn out to be simulated or virtual objects. If this is all one big dream, then tables still exist. In that case, this is a table "in the dream," and tables turn out to be ideas rather than mind-independent matter. What tables are depends on which scenario you're in. So, goes the alternative approach, there's no problem. Look down again at your table. When you realize that you don't know whether or not you are in a simulation or a victim of Descartes' demon, you should not conclude from this that therefore you don't know whether there is a table. If this is all a simulation or a demon scenario, then the table turns out to be a virtual object or an idea in the mind of Descartes' demon. *That's just what things are* in those scenarios. If you don't know whether you are in some such skeptical scenario, then although you know what there is – there is a table – you don't know what things are – what is this table, a computation in a simulation, an idea in the demon's mind? But that is no skeptical catastrophe. You don't know whether things are ultimately simulations or ideas in a demon's mind, but this is no different from perfectly acceptable metaphysical humility about the ultimate nature of things. We should already accept that the best metaphysics or science will not necessarily tell us everything about what things are, and that's all the skeptical possibilities show.

This alternative way of thinking swaps skeptical ignorance of what there is for skeptical ignorance about what things are. It is at least as old as Leibniz, but following its most recent proponent, Chalmers (2018), I call it *veridicalism*. The key point of this Element is that veridicalist skepticism about what things are is just as epistemically disastrous as standard skepticism about what there is. It isn't mere metaphysical humility. As we will see, veridicalism might address specific *formulations* of skepticism, but it does not solve the problem posed by the skeptical argument. In other words, the metaphysics of ordinary objects affects the *formulation*, but not the *epistemic significance*, of the skeptical conclusion. And this will reveal something important about the problem of skepticism. A satisfactory solution to the skeptical problem requires dealing with the part of the formulation that *doesn't* depend on one's metaphysical outlook. It takes getting clear about what part of the skeptical problem *does* depend on our metaphysics to see this.

An entirely reasonable reaction at this point is: But how could this be? The alternative, veridicalist metaphysics vindicates the claims that you know that there are tables, that you can put cups on tables, that tables are made of wood, and so on. (You merely lack knowledge of whether all of those things are ultimately simulations, or ideas.) How could this fail to solve the problem of skepticism about the external world? It shows that you know a lot! Consider that

a small child, or our ancestor in the ancient world perhaps, knows a lot about what there is. This knowledge is significant to their lives, and they know how to act in light of this knowledge. And yet, we can imagine, they either have no beliefs about what things are beyond how they appear (in the case of the child), or else they have a bunch of false beliefs about this (in the case of our ancestor, who perhaps thought that all things were ultimately made of water, or fire). Whatever it is that one worries about when one worries about the skeptical problem – that we know nothing about the world – surely these worries don't apply to the cases we are imagining, as described! Yet, *all* these people lack, it seems, is knowledge about what things are. So, how could veridicalism, which posits that we are roughly in the same position, since we know what there is but not what things are, possibly fail to satisfy us? It shows that we know at least as much as the child or the ancestor, and their situation does not seem worthy of philosophical panic.

The bulk of the arguments in this Element aim to show that this attitude, though initially appealing, is mistaken. The child's and the ancestor's ignorance about what things are is not the same as skeptical ignorance about what things are. As we will see, the former is a matter of immaturity (whether of the individual or of the science of their time). The latter is a matter of necessary, complete hopelessness. And, as we will see, it is equivalent in important ways to ignorance about what there is. It is true that not *all* kinds of ignorance about what things are constitutes an epistemic disaster. But a particular, skeptical sort of ignorance about what things are is different. And so, one can know a lot about what there is while knowing very little about the external world.

Whether skeptical hypotheses raise doubts about what there is or what things are is, as veridicalism helpfully shows, a matter of metaphysics. But if this table could be anything from a simulation to a demonic idea, then the *mere* conclusion that you know that there is a table is still compatible with almost total ignorance about the external world. As I will suggest, this also serves to clarify what a skeptical scenario is, and its relation to other metaphysical hypotheses about the world. Skepticism isn't *really* only about knowledge of what there is, rather it is standardly formulated that way because of a standard metaphysical picture.

Here is an overview of what is to come. In the rest of this section, after I introduce the standard, contemporary skeptical argument, I describe and set aside some more familiar, nonveridicalist antiskeptical strategies.

In Section 2, I introduce the veridicalist strategy and some of the different ways it has been pursued. This strategy avoids skepticism about what there is only by accepting skepticism about what things are. I then distinguish veridicalism from some other, externalist and idealist theories, and describe skepticism about what things are and how veridicalism is committed to it.

Sections 3 and 4 are where most of the action (and novel argumentation) occurs.

Section 3 presents arguments for thinking that skepticism about what things are leaves us with as little knowledge about the world as skepticism about what there is. One important consequence of this is that the veridicalist antiskeptical strategy fails – at least, *if* you were ever worried about standard skepticism about what there is, then your worry should not be assuaged by veridicalist skepticism about what things are.

In Section 4 I consider some objections to my arguments: veridicalism at least grants us knowledge that the world we see around us is reality; skepticism about what things are is not a radical skepticism at all, but rather mere scientific or metaphysical humility about the ultimate nature of things; it doesn't matter what things ultimately are pragmatically or for other values; the problem of skepticism is only about the existence of things, and other concerns about the nature of things are a change of topic. In reply to this last objection, I discuss some of the history of skepticism, from Sextus to Hume and Moore, to show that the exclusive focus on the *existence* of ordinary objects is a relatively recent phenomenon, and that skepticism about what things are has often been taken to be a concern.

Finally, in Section 5 I describe some more general lessons that can be learned from the way in which veridicalism fails to solve the skeptical problem. One can know what there is without knowing much about the world. Some other anti-skeptical strategies besides veridicalism fail for this reason, and furthermore the standard argument is inadequate as an expression of the skeptical problem. I also make some observations about what a successful antiskeptical strategy would have to look like.

1.1 A Standard Skeptical Argument

To begin, consider one currently standard formulation of skepticism and the skeptical argument. Skepticism, as Barry Stroud's (1984) seminal work formulates it, is the view "that no one knows anything about the world" (p. 1). If true, that would be epistemically devastating. But what exactly is knowledge "about" the world? Philosophers these days interpret this to be a matter of what there is: you don't know that there are tables. How is this skeptical conclusion reached?

This standard skeptical argument revolves around skeptical scenarios which purport to cast doubt on vast swaths of our beliefs. In his *Meditations on First Philosophy*, Descartes appeals to such scenarios, for example one in which you are being deceived by an all-powerful demon. So this standard skeptical argument is often referred to as the "Cartesian" skeptical argument. In Descartes' demon scenario, the demon makes you have the experience you are currently having, so that it seems like there is an external world full of ordinary things such as tables

and people. But, in fact, there are no such things – at least according to the standard description of the scenario. Let us focus on such *global* skeptical hypotheses, which on a standard metaphysics call all of our knowledge (or justified beliefs) about the external world into doubt.¹ For examples, the hypotheses that your entire life has been like one long dream, that you are a brain in a vat being stimulated by a computer in an otherwise empty world, and that you are in the scenario depicted in the film *The Matrix* since birth are all global skeptical hypotheses. Let “sk” stand for your favorite such hypothesis:

- (1) I do not know that ~sk.
- (2) If I do not know that ~sk, then I do not know that there are tables.
- (3) I do not know that there are tables (from (1) and (2))
- (4) Generalizing from tables, I do not know that there are any ordinary things (hands, tables, sandwiches, etc.).²

(I will remind the reader what each of (1)–(4) is each time it is used in a new section.) This is one standard formulation of the argument, but a few qualifications are helpful here. First, arguably the dream version of sk raises a different set of issues, which I address in Section 2. Second, one might reasonably construe the “standard” skeptical argument as the argument for (1), leaving the existence of ordinary objects as a further, different matter. Set those arguments aside here, since our concern is the relation between global skeptical hypotheses and the existence of ordinary objects. Third, there has been some discussion in the literature about so-called underdetermination principles which generate (3) and (4) in an alternative way. Veridicalism can be understood to apply to those versions of the argument as well, but at least the latest versions, due to Chalmers and Valberg, seem to focus on (1) and (2), so I will follow them.³ Ultimately, I conclude that identifying the skeptical problem with (4) is not quite right, since some views on which

¹ I briefly address one veridicalist’s (Chalmers’) view on nonglobal skeptical scenarios in my review of his book, Avnur (2023).

² This formulation closely follows Chalmers’ (2018, 1).

³ One recent and helpful discussion of the relation between underdetermination principles and the sort of closure principles that are typically taken to motivate (1)–(4) is Pritchard (2016, especially ch. 2). The underdetermination principle states, roughly, that if two propositions are (known to be) incompatible, then in order to know either one of them one’s evidence must favor that one over the other. Briefly, the underdetermination version of the skeptical argument is:

- (1*) My evidence does not favor the proposition that there are tables over sk.
- (2*) In order to know that there are tables, my evidence must favor the proposition that there is a table over sk. (by the underdetermination principle)
- (3) I do not know that there are tables. (from (1*) and (2*))
So,
- (4). Generalizing from tables, I do not know that there are any ordinary things.

(4) is false still fail to avoid the sort of epistemic disaster that makes (4) problematic. It is not necessarily enough to show that we know what things there are.⁴

Why is (4) such an epistemic disaster? The ignorance (4) 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 don't even know that there are tables, and the most you can know is that there are tableish sensory experiences, which is a nightmarish relation to have to reality. In other words, (4) makes explicit just how little you know about the world.

Although I will mostly be setting this aside, external world skepticism is not all about tables, it's also about people. As Stroud put it when contemplating the skeptical conclusion:

Other people, as I understand them, are not simply sensory experiences of mine; they too, if they exist, will therefore inhabit the unreachable world beyond my sensory experiences, along with the tables and chairs and other things about which I can know nothing. So at least with respect to what I can know I could not console myself with thoughts of a like-minded community of perceivers all working together and cheerfully making do with what a communal veil of perception provides. (Stroud 1984, 38)

You know nothing about things, or people, beyond your sensory experiences. Though the existence of other people is a central issue for those who worry about classic skeptical scenarios, as we will see in what follows, there is some question whether veridicalists mean to be addressing such scenarios at all. If they intend to address only those scenarios in which other people still exist (as in *The Matrix*), then they don't address a significant part of what has been worrying philosophers about (1)–(4). I will leave this point aside, though, to see how veridicalism fares even with respect to a narrower set of skeptical hypotheses, and even setting the existence of other people aside.⁵

The veridicalist strategy here would presumably be to reject the claim that the proposition that there is a table is incompatible with sk, and so to maintain that one can know it without possessing the requisite evidence that *favors* it over sk. Accordingly, (2*) is not supported by the underdetermination principles after all, and can be rejected. Other strategies may go the route of rejecting the principle, or rejecting (1*). But the veridicalist, it seems, can accept these and still hold that (2*) is false. Much of what I say about veridicalist objections to (1)–(4) will apply to this version as well, and I note that at least two recent proponents of the veridicalist strategy, Valberg and Chalmers, explicitly target (2) rather than (2*).

⁴ For the record, Stroud, in a different essay, seemed to agree, at least about the point that what things there are does not settle what sort of world we are in. See Stroud (1986, 263–264).

⁵ See Helton (forthcoming).

To be sure, some people, and even some philosophers, don't find (4) so worrying at all. What I will argue in the next section is consistent with that: *if* one finds standard skepticism about what there is disturbing, then one's worries should not be assuaged by veridicalism.

1.2 Topics to Set Aside

Since the focus in this Element is the metaphysical underpinnings of the standard skeptical argument, and since there is a large variety of antiskeptical strategies and some overlap between them, it will be helpful to explicitly set aside strategies and views that, though independently interesting, can distract from our target. Recall the standard argument:

- (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.
- (3) I do not know that there are tables.
- (4) I do not know that there are any ordinary things.

Philosophers try to avoid (3) and (4) by rejecting (1) or (2). The most direct strategy is to reject (1): show that your information is sufficient to rule out sk. Denying (2) is indirect in the sense that it targets what seems to follow from our ignorance about sk, rather than sk itself. Though I will be focusing on a particular kind of indirect approach, and setting aside direct approaches, the direct approaches will come up again in a couple of contexts.

First, there are well-known strategies, discussed in the next section and associated with Hilary Putnam's "semantic externalism" and George Berkeley's idealism, that superficially resemble veridicalism but are crucially different in various respects. As we will see in Section 2, Putnam-style externalism aims to show that we can rule out sk a priori, rather than showing that even if we don't know whether sk, we could still know that there are tables. We will also see that Berkeley's idealism is incompatible with various skeptical hypotheses (and posits a metaphysics of objects that is incompatible with veridicalism), and so also rules them out a priori. So these views reject (1), rather than (2). We will mostly set them aside.

Second, such rejections of (1) will come up again in Section 5. As we will see, the only way to satisfactorily solve the skeptical problem about the external world, insofar as this problem arises from consideration of skeptical hypotheses such as sk, is to pursue the direct approach of rejecting (1). In other words, much of what follows can ultimately be taken as an argument *for* the direct strategy as superior to the indirect one.

So, our focus is on the indirect strategy, which accepts (1) but rejects (2). But even here there is plenty of variety, and plenty to set aside. One familiar indirect strategy is to deny (2) for *epistemological* reasons. Implicit in the justification for (2) is the idea that, since the existence of tables implies the falsity of *sk* (recall that *sk* is a global skeptical scenario), knowing that there are tables implies knowing that not-*sk*. But if that is the justification for (2), then one way to avoid commitment to (2) is to avoid commitment to such conditions for knowledge. Maybe it isn't *always* the case that knowing something, *p*, and knowing that it entails something else, *q*, puts you in a position to deduce or otherwise know *q*. The considerations involved in this kind of strategy, often called "denying closure," are all epistemological, since they involve what is required to know something. That is, this strategy does not challenge the idea that if there are tables, then *sk* is false. It does not challenge the underlying *metaphysics* of objects. So, for the most part, we will set these strategies aside. However, as with direct strategies, these epistemological indirect strategies will come up again at the end, though this time not as beneficiaries. As we will see, much of the trouble for veridicalism is also trouble for some epistemological indirect strategies as well.

2 Veridicalism

In this section I describe a few versions of the veridicalist antiskeptical strategy and explain its commitment to skepticism about what things are. They each deny (2):

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

on the basis that there can be tables even if *sk* is true, because what appear to be tables in any global skeptical scenario *are* tables. So, accordingly, we don't need to know whether or not we are in a skeptical scenario in order to know that there are tables.⁶ Crucially, veridicalism is compatible with (1):

(1) I do not know that $\sim sk$.

Part of the appeal of the view is that it need not bear the burden of explaining how we know $\sim sk$.

Veridicalism can seem far-fetched. How could something in the mind of a demon, or in a simulation, be a table? Tables are real things, we want to say, and while simulations themselves can be real, what they are simulations *of* are

⁶ Notice that this does not require rejecting closure for knowledge. For a different take on the relation of (2) to skeptical hypotheses, see Roush (2010) and Avnur, Brueckner, and Buford (2011).

not, by virtue of being simulated, real. Let us now see how a few different philosophers motivate an alternative, veridicalist metaphysics.

2.1 Three Paths to Veridicalism

Veridicalism has been around at least since Leibniz (1923) and, in the last century, Bertrand Russell (1927) and O. K. Bouwsma (1949).⁷ More recently, veridicalism has been defended in different ways by David Chalmers (2005; 2012, 431–440; 2018; 2022), J. J. Valberg (2007), and, on one prominent interpretation, Donald Davidson (1986). Here I will discuss the latter three.

2.1.1 Chalmers' Structuralist Veridicalism

Chalmers has defended veridicalism in various works over many years (Chalmers 2005; 2012, 431–440; 2018; and 2022). He arrives at veridicalism through structuralism about physical properties.⁸ Roughly, structuralism about some property is the view that that property is individuated (or the term expressing the property defined) by its causal role (or by its relation to other terms in a theory) (2018, section 3). For instance, having *mass* is having a, or being the, property that plays the mass role (2018, 11). This also applies to other physical properties which things such as tables have, such as spatial properties. Details aside, the result is that being ignorant about *sk* does not itself imply that we are ignorant about whether there are tables. For, even if *sk* is true, there is something playing the role(s) associated with being a table, and that is a table.⁹

Presumably, on this view, if the table is a demonic idea (for example), then the fact that it has the causal role of a table amounts to the fact that it causes, in some sense of “cause” that we will grant the structuralist, other demonic ideas to occur as well as the relevant experiences in the mind of the victim. Oddly, when you knock on the wood of the table, on this view, you are knocking the demon's mind, and when you taste the coffee that is in the cup on the table, you are tasting the demon's mind. And if you are in the Matrix, you are tasting a computer chip in a certain state!

Whether one is in a simulation or a demon victim will make some difference to the causal profile of the table. If the table is a demonic idea, then a demon

⁷ Bouwsma is discussed in Chalmers (2018, 627–629).

⁸ Sklar (1985, 60–161) and Vogel (1990, 661) both consider, and reject, a sort of structuralist veridicalism view as well, as Chalmers points out.

⁹ See Chalmers (2018, 24–27) for some important details, including whether mass is identified with the property that actually realizes the requisite causal role, or whether it is identified as the property of having a property that plays that role. Though this is an important distinction for Chalmers, it won't play an important role here.

could, if it wanted, make the table disappear. Some (perhaps those more comfortable with theistic views according to which God constantly sustains all objects) might find this to be perfectly compatible with what they usually take the causal profile of tables to be, while others might find this to be inconsistent with how they assume tables behave. Presumably, though, all have tables in mind. So, in some respects, the causal profile of tables permits of some flexibility. Perhaps there is some range of different causal profiles that, if satisfied, is sufficient for tablehood. So, different “kinds” of tables will have different total causal profiles. But structuralist veridicalism holds that your ignorance about whether *sk* is true does not imply that you do not know that there are tables, because you know that something within that range exists. You still know that they have a certain mass, that they have four legs, that you can put cups on them, and so on; so (2) is false. On this view, *sk* is a metaphysical hypothesis about the ultimate nature of things like tables, and perhaps also about the specific causal profile of the table. If you are a demon victim, there is still a table here, but the table is ultimately a demonic idea that would vanish if the demon willed it. If you are in the Matrix or a brain in a vat (BIV), then the table is ultimately a computational state and could vanish according to the program, or by a glitch. And so on for other *sk*’s.

There are metaphysical objections to structuralism, independent of its implications for skepticism, that I will set aside.¹⁰ What I will argue in what follows is that structuralism (and other versions of veridicalism, and some other denials of (2)), even if true, would not suffice to solve the global skeptical problem.

2.1.2 Valberg’s Phenomenological Veridicalism

Valberg’s (2007) view is not intended to solve *all* relevant skeptical problems, but only the standard sort. One cannot do his nuanced phenomenological observations justice in summary form. But, roughly, he observes that the objects of experience are arrayed around you in a spatial expanse with you (the subject) at the center. He calls this array, or that within which things appear to you, your “personal horizon” (Johnston 2010, ch. 2, following Valberg in most details, posits a similar phenomenon, calling it your “arena of presence.”) According to

¹⁰ Vogel (2019) presents some, focusing on whether concepts such as distance and causation can be construed structurally. See Chalmers (2018, 18–19) for replies. But Vogel agrees with Chalmers that structuralism, if true, would solve the global skeptical problem (pp. 3–4), so on this point I differ from them both. Another sort of problem for structuralist veridicalism might be developed from Korman’s (2015, 39) argument that our ordinary object concepts are not so undemanding as to apply to *whatever* is out there in our environment.