

Structuralism as a Response to Skepticism*

David J. Chalmers

A Cartesian argument for skepticism about the external world runs as follows. First: we cannot know that the evil-demon hypothesis is false, where the evil-demon hypothesis says that all of our sensory experiences have been produced by an evil demon. Second: if we cannot know that the evil-demon hypothesis is false, we cannot know we have hands. Conclusion: We cannot know we have hands. Generalizing: as a version of the second premise holds for arbitrary empirical claims about the external world, we cannot know any such claims to be true.

Here the evil-demon hypothesis is put forward as a *global skeptical hypothesis*: a single hypothesis whose possible truth threatens knowledge of arbitrary empirical claims all at once. Other putative global skeptical hypotheses include the brain-in-vat hypothesis, holding that we are permanently envatted brains; the simulation hypothesis, holding that we have lived our lives in a computer simulation; and certain versions of the dream hypothesis, holding that we have always been dreaming.

I will call these putative global skeptical hypotheses *Cartesian hypotheses*, and the scenarios that they specify *Cartesian scenarios*. They do not all come from Descartes, and Descartes also put forward other sorts of skeptical scenario, but these scenarios have exerted an especially strong grip on the philosophical imagination. My focus here is on skepticism about the external world, so I will restrict my attention (and the label “Cartesian”) to scenarios in which a subject’s physical environment is varied globally, keeping the subject’s experiences the same. Scenarios in which only other minds are absent or in which an evil demon is interfering with one’s own mind also raise important skeptical issues, but they are beyond the scope of this paper.

Cartesian hypotheses contrast with *local skeptical hypotheses*, whose possible truth threatens empirical claims a few at a time. For example, the painted-mule hypothesis, holding that one is currently looking at a mule painted with black and white stripes, threatens knowledge that one

⁰Forthcoming in the *Journal of Philosophy*. Thanks to Jonathan Vogel and a reviewer for comments, and to audiences at ANU, Bled, NYU, Rio de Janeiro, USC, Sydney, and Villa de Leyva for discussion.

is looking at a zebra. Local skeptical hypotheses suffice for many skeptical purposes. Even if every empirical belief were threatened by a different local skeptical hypothesis, that would be enough to threaten all of our empirical knowledge. Nevertheless, global skeptical hypotheses yield a particularly strong form of skepticism, which we might call global skepticism. If global skepticism is correct, we cannot rule out the possibility that we are entirely out of touch with the external world. My focus here will be almost entirely on arguments for global skepticism, with local skeptical hypotheses set to one side.

Perhaps the best-known responses to global skeptical arguments deny the first premise, holding that we can know that the Cartesian hypotheses are false. These include theist replies (such as Descartes's) where we know this through reasoning about god; dogmatist replies (such as Moore's) on which we know this through perceptual evidence; abductivist replies (such as Russell's) on which we know this by inference to the best explanation; and incoherentist replies (such as Putnam's) on which Cartesian hypotheses are ruled out a priori as incoherent or contradictory.

I will focus on a response to global skepticism that denies the second premise. The general line of response I have in mind is *veridicalist*, holding that subjects in Cartesian scenarios have mainly true beliefs. On this line, even if the evil-demon hypothesis is true, I still have hands, and so on. Given this, my inability to rule out the evil-demon hypothesis does not rule out my knowing that I have hands. The specific veridicalist line of response I have in mind is *structuralist*: most of our claims about the external world (that I have hands, for example) are equivalent to structural claims, and these claims will be correct even for subjects in Cartesian scenarios, roughly because those scenarios have the right structure. I will discuss veridicalist lines in general in the next section, and structuralism specifically in what follows.

Veridicalist Responses

Veridicalist responses to skepticism can be motivated through a number of quite different philosophical positions, including idealism, phenomenalism, verificationism, coherentism, externalism, and structuralism.

The idealist and phenomenalist thesis that appearance constitutes reality yields a straightforward path to veridicalism. In a Cartesian scenario, it systematically appears that one has hands. Given that systematic appearance that one has hands constitutes one's having hands, it follows that one has hands. The same goes for other aspects of external reality. Surprisingly, Berkeley never discusses Descartes's skeptical hypotheses explicitly. Perhaps he would hold the evil-demon hy-

pothesis to be incoherent, since he thinks that only a god could produce sensory appearances with the coherence and complexity of ours. But Berkeley's thesis that god produces our experiences is at least a cousin of the evil-demon hypothesis, and Berkeley certainly thinks that given his thesis about god, our beliefs and experiences are accurate.

The locus classicus of a veridicalist response to global skepticism is O.K. Bouwsma's 1949 article "Descartes' Evil Genius". Here Bouwsma argues that although the evil demon tries to deceive us, it does not succeed. As he sums things up:

I have tried in this essay to understand the boast of the evil genius. His boast was that he could deceive, deceive about "the heavens, the earth, the colors, figures, sound, and all other external things." In order to do this I have tried to bring clearly to mind what deception and such deceiving would be like. Such deception involves illusions and such deceiving involves the creation of illusions... [I]t is assumed that the illusion is of such a kind that no seeing, no touching, no smelling, are relevant to detecting the illusion. Nevertheless the evil genius sees, touches, smells, and does detect the illusion. He made the illusion; so, of course, he must know it. How then does he know it? The evil genius has a sense denied to men. He senses the flower-in-itself, Milly-in-herself, etc [...] He has certainly created his own illusions, though he has not himself been deceived. But neither has anyone else been deceived. For human beings do not use the word "illusion" by relation to a sense with which only the evil genius is blessed. (Bouwsma 1949, pp. 150-51)

Bouwsma's idea can be put by saying that an illusion requires the possibility of discovery. If the evil demon fashions apparent flowers made out of paper, then we can discover this and know it to be illusion. But if the evil demon fashions an illusion that we can never know about, it is no illusion at all, at least as we use the term. Bouwsma's version of this idea seems to driven in large part by ordinary language philosophy—we use the term "illusion" for illusions we can discover—but the idea can also be motivated through verificationism (claims about illusion require the possibility of verification) and coherentism (if all our beliefs are coherent, they are correct and there is no illusion).

I am very much sympathetic with a veridicalist line on skeptical hypotheses, but I do not think Bouwsma's argument for this line succeeds. Even in ordinary language, we can make perfect sense of the idea of an undiscoverable illusion. More generally, I do not think that anti-realist

views such as idealism, phenomenism, coherentism, or verificationism provide much reason to accept a veridicalist line, as I think these views fail for well-known reasons.¹

In *Reason, Truth, and History*, Hilary Putnam discusses skeptical hypotheses such as the brain-in-vat hypothesis using quite different considerations grounded in semantic externalism and the causal theory of reference. Most of Putnam's discussion is in support of an incoherentist line on which I can rule out the hypothesis that I am a brain in a vat (by reductio: if I am a brain in a vat, externalism entails that my 'brain' refers either to nothing or to an external kind such as brains-in-the-image; so 'I am a brain in a vat' is not true; so I am not a brain in a vat). However, one paragraph in Putnam's discussion supports a veridicalist line on which brains in vats have mostly true beliefs.

By what was just said, when the brain in a vat (in the world where every sentient being is and always was a brain in a vat) thinks 'There is a tree in front of me', his thought does not refer to actual trees. On some theories that we shall discuss it might refer to trees in the image, or to the electronic impulses that cause tree experiences, or to the features of the program that are responsible for those electronic impulses. These theories are not ruled out by what was just said, for there is a close causal connection between the use of the word 'tree' in vat-English and the presence of trees in the image, the presence of electronic impulses of a certain kind, and the presence of certain features in the machine's program. On these theories the brain is right, not wrong in thinking 'There is a tree in front of me.' Given what 'tree' refers to in vat-English and what 'in front of' refers to, assuming one of these theories is correct, then the truth conditions for 'There is a tree in front of me' when it occurs in vat-English

¹An additional worry is that each of these views threatens to slide into incoherentism rather than veridicalism. Certainly the first-person evil-demon hypothesis ('My experiences as of an ordinary world are produced by an evil demon') seems meaningless by verificationist lights, because it is unverifiable. It is likewise hard to see how the hypothesis could be true by phenomenalist and coherentist lights, or by Bouwsma's own lights, if there is no first-person evidence for it. There is no analogous problem for a third-person evil-demon situation in which someone else's experiences are produced by a third party. We might have good evidence that this scenario obtains. We can then ask whether the subject's beliefs are true. It seems that the advocate of these views should say that they are, leading to a sort of veridicalism about the third-person case. All this tends to lead to the view that our claim 'They are the victim of an evil demon' is true while their claim 'I am the victim of an evil demon' is false. In effect, this view is incoherentist about the first-person case but veridicalist about the third-person case. Putnam's externalist route to veridicalism also has something like this consequence, while my structuralist route does not: for most Cartesian hypotheses, it is veridicalist about both the first-person and the third-person case.

are simply that a tree in the image be 'in front of' the 'me' in question in the image or, perhaps, that the kind of electronic impulse that normally produces this experience be coming from the automatic machinery, or, perhaps, that the feature of the machinery that is supposed to produce the 'tree in front of one' experience be operating. And these truth conditions are certainly fulfilled. (*Reason, Truth, and History*, p. 14)

Elements of both the coherentist and the externalist line are developed by Donald Davidson in "A Coherence Theory of Truth and Knowledge" (1989). As a coherentist Davidson holds that truth of a belief consists in coherence with other beliefs, which results in the thesis that it is impossible for most or all of a subject's beliefs to be false. In defending the thesis, Davidson advocates an externalist variety of interpretivism on which we interpret subject's beliefs according to their external cases:

What stands in the way of global skepticism of the senses is, in my view, the fact that we must, in the plainest and methodologically most basic cases, take the objects of a belief to be the causes of that belief. And what we, as interpreters, must take them to be is what they in fact are. (Davidson 1989, p.)

Davidson does not address global skeptical hypotheses directly, but in expounding Davidson's view in "Davidson versus Descartes", Richard Rorty takes the extra step for him:

That brain too is reacting to features of its environment. But its environment is the computer's data bank. The only way you can translate the noises it makes is to correlate them with the bits of data that the computer is feeding in. So the noises that sounds like 'It's Tuesday the 7th of October 2003, and I am eating tofu' must mean something like 'Now I am hooked up to sector 43762 of the hard drive'. For most the envatted brain's beliefs, like most of ours, must be true. It is not as easy to delude a brain as the evil scientist thinks. (Rorty 1993, p.)

I think there is much insight in the Putnam/Davidson veridicalist line on skeptical hypotheses, but I do not think that standard forms of externalism give compelling reasons to accept veridicalism. There are at least two reasons.

First: On standard causal theories of reference, a causal connection between a word and an item in the world is necessary for reference, but it is not sufficient. For example, 'phlogiston' may be caused by deoxygenization, but it does not refer to deoxygenization; rather, it fails to refer,

because deoxygenization does not satisfy the complex causal conditions required for reference. It is natural for a causal theorist with standard Cartesian intuitions to hold that 'tree' for a brain in a vat fails to refer for similar reasons. Externalism and the causal theory alone give little reason to think otherwise.

Second: The standard arguments for externalism and the causal theory of reference apply to some words but not to all. Perhaps the central argument is Putnam's Twin Earth thought-experiment, which suggests that twins can competently use the term 'water' with different referents. Let us put this by saying that 'water' is Twin-Earthable. One can run Twin Earth thought-experiments for numerous other terms, including 'tree', and 'brain', but it is much harder to run them for many other words such as 'plus', 'three', 'person', 'believe', 'see', 'object', 'philosopher', or 'cause'.² Certainly the standard arguments do not support the claim that these words are Twin-Earthable, and I think it is more plausible that they are not. For this reason, Putnam's externalist argument will not apply to claims cast in non-Twin-Earthable language, such as 'I can see three objects' or 'I am talking to three philosophers', and cannot be used to argue that these claims are true for a brain in a vat. And if claims like these are false in Cartesian scenarios, it seems likely that other claims in Twin-Earthable language (for example, 'I can see an apple, a banana, and an orange') are false too.

There are stronger versions of externalism to which these objections do not apply. If we accept a crude causal theory of reference on which a simple causal connection suffices for reference, the first objection will not apply. If we accept a global externalism on which every use of every word is Twin-Earthable, the second objection will not apply. But the standard arguments for a causal theory of reference give little reason to accept a crude causal theory, and the standard arguments for externalism give little reason to accept global externalism.³ Furthermore, I think there is good reason to think these views are false. Perhaps a more sophisticated externalist view combined with further considerations might overcome these objections, but as things stand, the externalist case for veridicalism seems weak.

³Arguably, every term is subject to Burge-style social externalism. For example, one can construct a Burge-style Twin Earth case in which two duplicate non-expert subjects use 'zero' with deference to their linguistic community, and in which the two communities use 'zero' for zero or for one respectively. Then these twins use 'zero' with different meanings and different referents. However, Burge-style deference does not lead to a veridicalist line on skeptical scenarios, unless Putnam-style or other considerations have already gotten us there at the level of community use of a term. Furthermore, it is plausible that Burge-style externalism is not truly global, in that it applies only to certain uses of these words: roughly, to deferential uses, where a subject defers to the community. It is possible to use terms like 'zero' nondeferentially, so that even if one's community uses the term to refer to 1, one's own use refers to 0. So in order

In “The Matrix as Metaphysics” (2003), I develop another veridicalist response to global skepticism, holding that subjects in a computer simulation have mostly true beliefs, as do subjects in evil-demon scenarios, global dream scenarios, and so on. In that paper I argue that the simulation hypothesis is equivalent to a certain metaphysical hypothesis, one in which standard physics is grounded in a deeper level of computational physics. Given the further claim that this metaphysical hypothesis is one in which our beliefs are mostly true, a sort of veridicalism follows. That argument (which I will call the “matrix argument”) does not take any specific philosophical view such as structuralism or externalism as a premise. Instead it argues for veridicalism from first-order considerations about various scenarios. However, certain key ideas in the matrix argument are structuralist in character. A crucial part of the argument holds, for example, that only structural considerations matter for a theory in computational physics to be true, so that given that a simulation scenario has the right structure, it will be one in which computational physics (and therefore standard physics) is correct.

In the remainder of this paper I will develop a structuralist response to global skepticism explicitly, making explicit a structuralist thesis and showing how it leads to a veridicalist conclusion that undermines global skepticism.⁴ This structuralist argument is distinct from the matrix argument, but the two are complementary. The structuralist argument here in effect explains why the other argument is effective, if indeed it is, and the matrix argument can be used to support certain key claims in mounting the structuralist argument.

Preliminaries

Before proceeding, some bookkeeping. In this paper, a hypothesis (such as the Cartesian hypothesis that one is a brain in a vat) can be construed as a sentence, a proposition, or a mental item such as a thought. A hypothesis is epistemically possible (in the sense relevant here) when it is not conclusively ruled out a priori. A scenario is a highly specific epistemic possibility, akin to an epistemically possible world: for example, a certain specific situation in which one’s experiences

for a skeptical hypothesis to be immune to Burge-style externalism, one need only stipulate that subject entertaining the hypothesis uses the relevant words nondeferentially. In order for the definition of Twin-Earthability in the text to exclude Burge-style cases, the requirement of competence should be understood as requiring nondeferential usage.

⁴I develop a related structuralist response to skepticism briefly in Excursus 15 of *Constructing the World* (Chalmers 2012, pp. 431-40). That response presupposes a good deal of material from earlier in the book and in particular works within the book’s framework of scrutability. The response in this article is largely independent of the scrutability framework and is intended to be free-standing.

are being produced by an evil demon is an epistemically possible scenario, at least if incoherentism is false. I will talk of hypotheses being true or false in scenarios. Roughly speaking, a hypothesis *H* is true in scenario *S* iff accepting that we are in *S* should lead us to accept *H* (or better, conditionally supposing that we are in *S* should lead us to conditionally accept *H*). For example, the hypothesis that I have hands is true in a certain evil demon scenario if accepting that I am in that scenario should lead me to accept ‘I have hands’. This is roughly the sort of criterion suggested by the Ramsey test for evaluating an indicative conditional: if I am in the evil-demon scenario, do I have hands?

It should be noted that epistemically possible scenarios work somewhat different from metaphysically possible worlds. For example, it is epistemically possible (at least at the start of inquiry) that water is XYZ, even though this is not metaphysically possible. There is an epistemically possible Twin Earth scenario in which the liquid in the oceans and lakes (and so on) is XYZ, and the hypothesis that water is XYZ is true in that scenario. If we accepted that we were in that scenario, we should accept ‘Water is XYZ’. Note that all this is quite consistent with the claim that it is metaphysically impossible that water is XYZ. It is also consistent with the claim that ‘Water is XYZ’ is false in a metaphysically possible Twin Earth world, and with the counterfactual claim that if a Twin Earth world obtained, water would be XYZ. The issues in this article are largely epistemic, not metaphysical.⁵

Because of this, it can be misleading to think of Cartesian scenarios in counterfactual or third-personal terms. Veridicalism does not require the counterfactual claim that if I were in a simulation I would have hands, or the third-person claim that someone being stimulated by evil demons to have visual experiences as of tables is really seeing a table. The most fundamental issues here concern first-person epistemic possibilities as above. Most straightforwardly: if I am in a simulation, do I have hands? Here the simulation hypothesis is considered as a hypothesis about my own situation, and the conditional is evaluated according to the epistemological Ramsey test as above.

That said, if one is uncomfortable or unfamiliar with this style of analyzing epistemic possibilities, many of the issues can be cast somewhat more conservatively in metalinguistic terms. For example, one can ask: if I am in a simulation, is what I express with ‘I have hands’ (in the language I use in the simulation) true? If I am in a simulation, should I accept ‘I have hands’? One can even use a metalinguistic counterfactual: if I were in a simulation, would what I express with ‘I have

⁵See “The Nature of Epistemic Space” for a more rigorous treatment of truth in epistemically possible scenarios.

hands' be true? Or one can use a metalinguistic third-person claim: if someone is being stimulated by an evil demon, are their beliefs (such as the belief they express with 'I see a table') true? These counterfactual and third-person claims are not quite the same as the epistemological claims, but they are close enough that establishing them would at least do much of the work required for the veridicalist anti-skeptical strategy.

Another bookkeeping issue: Veridicalist responses to skepticism hold that subjects in Cartesian scenarios have mostly true beliefs. An immediate objection may be that for certain Cartesian hypotheses, veridicalism is ruled out by definition. For example, under the Cartesian hypothesis that one is entirely deceived by an evil demon, it follows that most or all of one's beliefs are false. Under the Cartesian hypothesis that one is a handless brain in a vat, then it follows that one does not have hands. And so on.

To finesse this issue, I stipulate that Cartesian hypotheses and Cartesian scenarios are specified in a way that makes no explicit commitment on the veridicality of ordinary beliefs. For example, the evil-demon hypothesis is the hypothesis that all of our sensory experiences have been produced by an evil demon, and the brain in a vat hypothesis is the hypothesis that we are brains in vats receiving inputs from computers. Similarly, evil-demon scenarios will be certain specific scenarios in which our sensory experiences are produced by an evil demon, and so on. This leaves it a substantive issue whether our beliefs are true in these scenarios. The skeptic typically finds it obvious that even in these neutrally specified scenarios, our beliefs are false, but the veridicalist thinks that in these scenarios, most of our beliefs are true. As for the hypothesis that one is entirely deceived by an evil demon, and other non-neutral hypotheses: a veridicalist is likely to say that these hypotheses are incoherent, or least that they are not vindicated by the scenarios that one normally has in mind in considering the hypotheses. I will return to those hypotheses later.

Structuralism

Structuralism about a domain says roughly that all claims (or all truths) about those domains are equivalent to structural claims. I will focus on structuralism about the physical, in a version that says all physical claims are equivalent to structural claims. Four notions need explication: claim, physical, equivalent, structural.

What is a claim? In principle a claim (like a hypothesis, above) can be an assertive sentence, a proposition, or a thought. Ideally one would do things in terms of propositions, but controversies about the individuation of propositions (whether they are Russellian or Fregean, for example) tend

to confuse important issues here. For present purposes we can focus initially on sentences (or to avoid worries about context-dependence, utterances of sentences or sentences in contexts), for which the issues are clearest and most concrete.

What is a physical claim? This is a claim about physical objects using ordinary physical vocabulary: roughly the vocabulary of space (location, shape, size, and so on), time (including motion and so on), and mass, plus logical and mathematical expressions. We could also allow a more theoretical physical vocabulary to be used here, but this will not make much difference for the purposes of the argument.

What is equivalence? For present purposes this can be understood as a priori equivalence: A and B are equivalent when “A iff B” is knowable a priori. For our purposes we could also understand it as a conceptual equivalence, or semantic equivalence or truth-conditional equivalence, as long as these relations are epistemologically constrained (roughly so that when two claims are equivalent, we can know a priori that they are equivalent).

What is a structural claim? For present purposes this can be understood as a claim using only logical and mathematical expressions and a limited vocabulary of auxiliary structural expressions (to be characterized shortly), perhaps along with phenomenal expressions (to capture the role of observation).

The structuralist thesis here is closely related to structural realism in the philosophy of science. Standardly, structural realism is understood as a claim about scientific theories, including especially physical theories, holding that roughly that those theories are equivalent to structural theories. Epistemological structural realism says roughly that all knowable scientific truths are equivalent to structural truths, while ontological structural realism says that the reality that scientific theories describe is wholly structural. We can also formulate domain-relative versions of epistemological and ontological structural realism, where the truths and theories in question are restricted to specific domains such as the physical domain.

The version of structural realism that is most relevant for my purposes is a close relative of these two. What we might call *conceptual structural realism* holds that scientific theories are equivalent to structural theories, and that scientific claims are equivalent to structural claims. Conceptual structural realism is not the same as epistemological or ontological structural realism, since it does not make claims about knowledge or ontology. It entails epistemological structural realism, but not vice versa: the two views may come apart if some scientific claims are nonstructural but unknowable. It does not entail ontological structural realism, or vice versa: conceptual structural realism is consistent with there being nonstructural properties in reality (for example, a

nonstructural mass property that plays the structural mass role).

Conceptual structural realism has been a central form of scientific structuralism throughout the 20th century. Certainly versions of it were advocated by Russell, Carnap, and Lewis in their cases for structuralism. It is also conceptual structural realism that falls most naturally out of the Ramsey-Carnap-Lewis method for regimenting scientific theories into Ramsey sentences.⁶

The Ramsey-Carnap-Lewis method is a method for eliminating theoretical terms from a theory, moving from a statement of the theory that uses a theoretical term to an equivalent statement without it. The basic idea is that one starts with a complete statement T of a theory, including both principles that connect theoretical notions to each other and principles connecting them to observation. For any given theoretical term that occurs in this statement ('charge', say), one can represent this theory as saying $T'(\text{charge})$, where T' is a complex one-place predicate. The corresponding Ramsey sentence is then $\exists\phi T'(\phi)$. Advocates of the Ramsey-Carnap-Lewis method typically hold that the original theory T , using the word 'charge', is equivalent to the Ramsified theory $\exists\phi T'(\phi)$ that does not use the word.

If we repeat this process for every theoretical term in a theory, we are left with a Ramsey sentence R that uses only non-theoretical expressions. For example, if the original theory T says $T^*(\text{mass, charge, space, time})$, the Ramsified theory R will say $\exists\phi\exists\psi\exists\mu\exists\nu T^*(\phi, \psi, \mu, \nu)$. The non-theoretical expressions left in R are usually construed to include logical and mathematical expressions and perhaps certain auxiliary structural expressions (to be discussed shortly), as well as observational expressions. The standard claim is that the Ramsey sentence R encapsulates the content of the original theory T , so that R and T are conceptually equivalent.

What undergirds this conceptual equivalence is the idea that theoretical terms such as 'charge' get their entire meaning from the role they play within the theory. In effect, 'charge' just means 'the property that plays the charge role', where the charge role is specified by the theory. On one very simple view, the charge role is given by the entire theory, so that 'charge' is conceptually equivalent to 'the property ϕ such that $T(\phi)$ '. This view is almost certainly false, since it is usually the case that some parts of a theory can turn out to be false without us saying that charge (or another theoretical entity) does not exist. More plausible is that there is some subset C of the theory (perhaps a weighted set of principles) that encapsulates the "core" charge role, such that 'charge' is conceptually equivalent to 'the property ϕ such that $C'(\phi)$ ' (where C' is a one-place predicate encapsulating what C says about charge). Then the theory T as a whole, which says

⁶See Carnap 1927; Lewis 1972; Ramsey 1931; Russell 1927.

T' (charge), is equivalent to $\exists\phi(C'(\phi)\&T'(\phi))$. This is in turn equivalent to $\exists\phi T'(\phi)$, given that $T'(\phi)$ entails $C'(\phi)$ (because C is a subset of T).

In moving from theoretical terms to ordinary physical terms, a complication is that many key scientific notions, such as ‘mass’, occur in more than one theory, so their content is not plausibly given by their role within a single theory. In these cases one can usually stipulate a technical notion (‘gravitational mass’, for example) tied to a single theory, but for the broader notion that occurs across theories, then a core associated role may well involve elements of multiple theories. Furthermore, many of these notions occur prescientifically, in which case a full accounting of the associated role for everyday physical concepts may include elements of folk physics, including connections to observation. Still, it remains plausible that ‘mass’ can be understood as ‘the property that plays the mass role’, where the mass role may involve multiple theories and also connections to observation.

Given this equivalence for arbitrary physical terms such as ‘mass’, ‘space, and ‘time’, we have the resources to translate not just physical theories but arbitrary physical sentences into structural sentences. Take a physical sentence such as ‘There is an object with mass two kilograms three meters from me’. To translate the sentence, we need only Ramsify the physical expressions included. Mass can be represented as mass-in-kilograms, a function from objects to numbers, which can then be Ramsified as something like ‘the function ϕ such that $R_1(\phi)$ ’. Likewise distance can be represented as distance in meters, a function from pairs of objects to numbers, which can be Ramsified as something like ‘the function ψ such that $R_2(\psi)$ ’. Applying these equivalences to the original sentence yields something like: $\exists x\exists\phi\exists\psi(R_1(\phi)\&R_2(\psi)\&(\phi(x) = 2)\&(\psi(x, me) = 3))$. The result will be a structural sentence as before.

What expressions will be left in these structural sentences, after theoretical expressions are eliminated? One class of expressions left in are observational expressions. On some views these will include expressions for externally observed qualities such as colors, shapes, and the like. However, we can use the Ramsey method to eliminate these expressions in favor of phenomenal expressions (concerning conscious experiences) and possibly indexical expressions, for example replacing ‘red’ by something like ‘the property that normally causes red experiences in me’ (or perhaps ‘in my community’). I will discuss this move further later in the paper (where the application to space in particular raises some crucial issues), but for now I will assume that phenomenal and indexical expressions are all the observational expressions that we need in structural sentences.

Observational expressions aside, what other non-Ramsified terms will be left in a structure sentence? The early structuralists (Carnap and Russell) hoped that these might include only logical

and mathematical expressions. Around 1927 the mathematician Max Newman as well as Carnap himself independently noted that structure sentences like this are too weak to capture the content of our scientific theories: they are satisfied by any world of the right cardinality (and with the right phenomenal states, if phenomenal expressions are included). To remove Newman's problem, one needs to include additional expressions. One popular addition is expressions constraining the classes of properties involved, such as 'natural' or 'fundamental', capturing the idea that theories should involve relatively natural properties rather than wildly disjunctive properties (which play a key role in Newman's demonstration). Another popular addition is nomic and causal expressions such as 'law' and 'cause', capturing the idea that theories are not just capturing mathematical structure but nomic and causal structure. For now I will adopt a view on which expressions of both sorts can be included, perhaps along with mereological notions such as 'part' and categorical notions such as 'object', and 'property'.⁷

Under these constraints, a typical Ramsey sentence for a physical theory may take the form: "There exist natural properties and relations such that these properties and relations stand in such-and-such nomic/causal relations to each other and to experience". A typical structure sentence for a physical sentence will involve both a Ramsey-like part like this along with a part specifying a particular distribution of the properties and relations in question, along the lines given for 'There is an object with mass two kilograms three meters from me' above.

In this way, application of the Ramsey method supports conceptual structuralism about the physical. Of course like every philosophical thesis, this thesis is not uncontroversial. Perhaps the most controversial step involves the existence of a core role associated with physical notions such that it is a priori or analytic that (for example) mass is whatever plays the mass role. This involves a sort of descriptivism about physical terms that some will dispute. It is possible to develop a structuralist response to skepticism without this descriptivist version of structuralism,⁸ but to avoid complications I will stay with the simpler version here. There is also an important question about whether every physical notion, including space and time, can be eliminated via the Ramsey method.

It should be noted that conceptual structuralism about the physical does not entail a metaphys-

⁷The problem for logical structuralism is pointed out by Carnap 1927 and Newman 1928. Carnap invokes naturalness to solve the problem, as does Lewis (1984). For discussion of just what it is to be structural, and of which expressions should not be Ramsified away, see Chalmers (2012, pp. 362-5 and 409-22.)

⁸See *Constructing the World* (excursus 15), in which a non-descriptivist version of the response is developed using the scrutability framework.

ical structuralism on which all physical properties are identical to structural properties (roughly, the property of playing a certain role, or of being disposed to play that role, or of having a property that plays that role). The view I have set out here identifies physical properties with those that play a certain structural role. These properties may themselves be fine-grained structural properties, or nonstructural properties. They might even be hidden “quiddities”. There are also versions of conceptual structuralism where physical properties are defined as structural properties: for example, where ‘mass’ is defined as a disposition or a power to resist acceleration. I discuss all this in more depth in the section on realizer structuralism and role structuralism.⁹

A terminological point: structural realism is usually understood as a claim about scientific theories, whereas the claim that ordinary physical concepts such as ‘mass’ and ‘space’ are structural is not exactly a claim about theories. I will speak of it more generally as a form of *structuralism* about the physical domain. This term also has the advantage that it may seem question-begging to assume structural *realism* in arguing against a skeptic who questions realism about the physical world. In fact, the conceptual structural realism discussed here is realist about the content of scientific theories (these make a substantive claim about the world beyond appearances) but not about their truth. Structuralism about physical concepts is realist about content in a similar way. Realism about content does not beg any questions against the skeptic, and in fact tends to be more congenial to skeptical arguments than anti-realism about content. But calling the thesis “structuralism” rather than “structural realism” will help avoid confusion on this score.

The Structuralist Argument Against Global Skepticism

We are now in a position to mount a structuralist argument for veridicalism.

Let us say that a non-Cartesian scenario is one in which experiences are not produced by dreams, simulations, or other models of an environment. It is plausible that for every non-Cartesian scenario, there are various corresponding Cartesian scenarios, in which the same experiences are had by the subject at the center of the scenario but in which that subject’s environment involves a model of the non-Cartesian scenario.¹⁰ There is a simulation scenario in which the non-Cartesian scenario is perfectly simulated; there is a demon scenario in which the demon keeps track of all elements of the non-Cartesian scenario; there is a dream scenario in which the dream involves a

⁹For metaphysical structuralism where physical properties are identical to structural properties, see (for example) Hawthorne 2002 and Shoemaker 1980. For an epistemological structuralism involving hidden quiddities, see (for example) Langton 1998 and Lewis 2009.

complete model of the non-Cartesian scenario; and so on.

It is not obvious that for every Cartesian scenario, there is a corresponding non-Cartesian scenario. But plausibly this is the case for those Cartesian scenarios that involve a complete model of a non-Cartesian scenario as above. I will call these *full-scale Cartesian scenarios*. For example, for a non-Cartesian scenario reflecting the universe as we normally think of it, there is a corresponding full-scale simulation scenario in which physics throughout the universe is simulated, and similarly for full-scale demon and dream scenarios.

I will start with an argument that applies to full-scale Cartesian scenarios. In the next section I will generalize to other Cartesian scenarios.

(1) All physical claims P are equivalent to structural claims $S(P)$.

(2) For all positive physical claims P , if $S(P)$ is true in a non-Cartesian scenario, it is true in any corresponding Cartesian scenario.

—

(3) For all positive physical claims P , if P is true in a non-Cartesian scenario, it is true in any corresponding Cartesian scenario.

Premise 1 is a statement of the structuralist thesis from the previous section. Premise 2 is the other key premise. Here the idea is roughly that all the structure present in an ordinary scenario is present in a corresponding Cartesian scenario. The physical claims are restricted to positive claims in order to accommodate the fact that Cartesian scenarios may involve extra structure not present in the original non-Cartesian scenario. Given the additional background premise that equivalent claims are true in the same scenarios, the conclusion follows. The conclusion suggests that veridicalism is true for full-scale Cartesian scenarios, at least given that it is true for corresponding non-Cartesian scenarios. For example, as long as most of our beliefs are true in the physical world as we ordinarily conceive of it, they will be true in a corresponding Cartesian scenario. If this is right, full-scale Cartesian scenarios do no work to support global skepticism, and their epistemic possibility does not undercut our knowledge of the external world.

¹⁰In assuming that Cartesian scenarios are epistemically possible (not ruled out a priori), I here assume the falsity of incoherentism. In assuming that Cartesian and non-Cartesian scenarios can involve the same experiences, I assume the falsity of certain forms of phenomenal externalism, on which experiences constitutively depend on the environment (by principles that are knowable a priori). These forms of phenomenal externalism might be developed into a distinctive incoherentist response to skepticism; see McDowell 2008 and Peacocke 2014 for discussion.

Admittedly the conclusion concerns only positive physical claims and not external-world claims in general, but even this limited conclusion is enough to undercut the argument for global skepticism. There is a genuine skeptical issue about knowledge of negative claims (how can we know that there is not more to the world than what we find here?), but this goes beyond standard external-world skepticism and at best leads to partial skepticism, not global skepticism. Setting these aside, it is easy enough to generalize from positive physical claims to positive external-world claims more generally. In Cartesian scenarios, the main reasons for doubting an external-world claim such as ‘There’s a tree in front of me’ are also reasons for doubting physical claims such as ‘There’s an object of a certain size and shape in front of me’. Once reasons for doubting physical claims have been defeated, there remains little reason to doubt other external-world claims. One could also give a more comprehensive version of the argument applying to any positive external-world claims grounded in structuralism about all or most external-world notions, but this goes beyond what is needed to defeat global skepticism.

It remains to defend premise 2: for all positive physical claims P, if S(P) is true in a non-Cartesian scenario, it is true in a corresponding Cartesian scenario. The central idea is that the positive causal structure in a non-Cartesian scenario is also present in a corresponding Cartesian scenario, so that the positive structure sentences true in the former are true in the latter.¹¹

It is easiest to see this in the case of a simulation scenario. Simulations work by replicating the causal structure of a system in detail.¹² For every physical object in the actual world, there will be a corresponding virtual object in the simulation scenario. Virtual objects are real objects realized within computer processes: roughly, they are data structures or entities constituted by data structures. For every physical property of physical objects in the actual world, there will be a corresponding virtual property of virtual objects in the simulation: this will be a property of data structures, such as a value in a data array. If two properties have a systematic causal connection tying them together in the actual world, the corresponding virtual properties will have a systematic causal structure in the simulation. Furthermore, because we have stipulated that Cartesian scenarios should be experientially identical to ordinary scenarios, where a physical property brings

¹¹Sklar (1985, pp. 60-61) and Vogel (1990, p. 661) both make the point that full-scale Cartesian hypotheses are structurally isomorphic to non-Cartesian hypotheses. Vogel suggests that under a “computer skeptical hypothesis”, portions of a computer disk and certain data properties could stand in the same pattern of causal relations as objects and properties in the real world, although he rejects the claim that this structural isomorphism yields equivalence in a stronger sense (see Vogel forthcoming). Sklar (p. 60) notes that in principle a structuralist about theoretical equivalence could “bite the bullet” and use this structural isomorphism to argue that certain Cartesian and non-Cartesian hypotheses are equivalent.

about experiences in our world, the corresponding virtual property will bring about the same sort of experiences in the simulation.

This structural similarity gives good initial reason to think that the relevant structure sentences satisfied in the ordinary scenario are also satisfied in a corresponding simulation scenario. The Ramsey sentence for physical properties requires mainly causal structure: that there exist properties and relations that stand in certain nomic/causal relations to each other and experience. For the reasons above, it is plausible that if this structure is present in a non-Cartesian scenario, it is present in a corresponding simulation scenario. Furthermore, given that the Ramsey sentence for the physical is satisfied in the simulation, it is just a small step to the satisfaction of specific structure sentences corresponding to specific positive physical claims, such as ‘There is an object with mass two kilograms three meters from me’. These specific structure sentences will conjoin the Ramsey sentence, or enough of it to say that there are properties ϕ and ψ playing the mass role and the distance role, with a claim about how those properties are instantiated. If the original claim is true in the actual world, and the simulation is accurate, there will certainly be a virtual object with the right virtual properties (for example, standing in a virtual distance of three meters from my virtual body) to satisfy that claim.

What goes for a simulation scenario also goes for other full-scale Cartesian scenarios. If the brain-in-vat scenario involves a computer, as it is usually held to, then the discussion of the simulation scenario applies directly. As for the evil-demon scenario: supposing the demon is modeling classical physics, then for every classical particle there will presumably be some sort of representation in the demon’s mind, and these representations will interact very much as virtual objects interact in the simulation scenario. Something similar goes for the dream scenario, except here the representations will be in part of one’s own mind. In both cases, these representations and their properties will have causal structure that enables the relevant Ramsey sentence to be satisfied.

This makes an initial case for premise 2. There are a number of immediate possible objections. Perhaps the most important objections are tied to the fact that the non-Cartesian scenarios are not completely structurally isomorphic to the original scenario, so that not all of the same structural sentences are true there.¹³ We can illustrate with the case of simulation scenarios.

¹²See “On Implementing a Computation” (1994) for an analysis on which implementing a computation requires a system to have appropriate causal structure, and “The Virtual and the Real” (2017) for a discussion of the metaphysics of virtual objects.

¹³A version of this strategy is suggested by the discussion in Vogel (forthcoming). Vogel argues that there are important structural differences involving fundamentality, modality, and causation (including differences on the lines of (i)-(iii) below) between space in non-Cartesian scenarios and its counterparts in simulation scenarios, and likewise

Some potential structural differences between non-Cartesian scenarios and simulation scenarios include: (i) the virtual properties involved in the simulation scenario are somewhat less fundamental and less natural than their counterparts in the original scenario; (ii) these virtual properties are not connected by iron-clad laws of nature; (iii) causal connections among virtual objects and properties are heavily mediated, for example by a central processing unit; (iv) more generally, the simulation will have a great deal of extra fine-grained causal structure, for example in the processes that realize the simulation; (v) the simulation is presumably embedded in a broader world which contains much additional structure; (vi) all this is designed by someone who might change or shut down the simulation at any time.

These observations are all correct, and they entail that not every sentence in a structural vocabulary that is true in a non-Cartesian scenario is true in a Cartesian scenario. However, I do not think they undermine premise 2. An initial observation is that many of them (especially (iii)-(vi)) involve *extra* structure that is present in the Cartesian but not the non-Cartesian scenario. In cases like this, it is arguable that any *positive* structure sentence that is true in the non-Cartesian scenario will also be true in the Cartesian scenario. If so, and if positive physical sentences are all equivalent to positive structure sentences, then cases like (iii)-(vi) do not pose a threat to premise 2. There remains a threat from cases (i) and (ii), however, insofar as it is arguable that positive structure statements about fundamentality and laws of nature can be true in non-Cartesian scenarios but not the corresponding Cartesian scenarios.

The key cases here include structural claims such as “a fundamental property plays the mass role”, “the properties playing the mass role and the distance role are connected by an iron-clad law”, and “other subjects have experiences”. These plausibly can be true in a non-Cartesian scenario without being true in the corresponding non-Cartesian scenario. *If* structural claims like these are involved in structure sentences associated with ordinary positive physical claims (for example, if ‘mass’ was equivalent to ‘that fundamental property that plays the mass role’), then ordinary physical claims would also be false in a simulation scenario. However, I will argue that these structural claims play no role in the structure sentences associated with ordinary physical claims, so they do not threaten the truth of these physical claims in simulation scenarios.

To make this case, we need only consider non-Cartesian scenarios in which the relevant constraints are violated. I will focus on constraints involving mass, the what I say applies equally to space and other crucial physical properties. For example, we could certainly discover that mass

between material objects in non-Cartesian scenarios and their counterparts in simulation scenarios.

is not fundamental and is instead a derivative property. This suggests that fundamentality is not required for there to be mass in a scenario, and that fundamentality is not part of the structure sentence associated with mass. Similarly, we could discover that current laws such as laws of gravitation are not iron-clad but apply only in certain special circumstances that our cosmos happens to be in. In fact, something like this applies to all the cases (i)-(vi). We could discover that what we thought were simple causal connections in our world are mediated. We could discover that there are a number of levels of physical structure, with enormous causal complexity, below the level of mass. We could discover that our universe is embedded in a much broader cosmos containing many such universes. We could discover that there is a god who created the physical world and who sustains the laws of nature and mediates all causal connections. None of these things would lead us to say that there is no mass. In fact, for every structural constraint violated in the simulation scenario, it is easy to come up with a non-Cartesian scenario in which the same constraints are violated, and in which there is still mass. The same goes for space and other physical properties.

All this suggests that these structural constraints do not play an essential role in the application of physical terms or in the truth of ordinary physical claims about the world. So they are no obstacle to the truth of ordinary physical claims in a simulation scenario. Of course some highly theoretical claims, for example involving fundamentality and iron-clad laws, may be true in an ordinary scenario and false in a Cartesian scenario. If so, structuralism alone does not rebut a Cartesian skeptical challenge to our knowledge of these theoretical claims. But these claims are not ordinary physical claims; and even if they were, the objection would threaten only a small subset of physical claims, not physical claims generally.

A related issue involves possible structural differences involving the roles of minds in these scenarios. We have allowed that structure sentences can include claims about connections to experiences. But some Cartesian scenarios differ from corresponding non-Cartesian scenarios. For example, we have stipulated that a Cartesian scenario should involve the same stream of experiences for a subject at the center of the scenario, but experiences in others may be different or absent. For example, there is no guarantee that a simulation scenario will replicate the experiences of others; at least, this would require the truth of certain functionalist theses about the mind, which it would seem odd to assume here. I think we can respond to this sort of worry as with those above, however. The existence of other minds is not required for our ordinary physical claims to be true. Even if it turned out that everyone else was a zombie, my claims about the existence of a square object in front of me could still be true. So this potential difference between Cartesian and

non-Cartesian scenarios does not threaten the truth of ordinary physical claims.

A deeper issue about the mind arises from the demon and dream scenarios. In these cases, it seems that entities in the subject's environment will be mind-dependent: in one case they depend on the mind of the demon, and in the other case on the mind of the dreamer. Since mind-independence of the external world is often used as a criterion for realism, this mind-dependence may limit the sorts of realism that is true if we are in these scenarios. Some may even think that mind-dependence of this sort will entail that ordinary physical claims are not true.

While mind-independence may be a criterion for certain sorts of realism, I do not think it is a criterion for the truth of physical claims. For example, it could turn out that a certain sort of panpsychism is true, on which fundamental physics is realized at the microphysical level by conscious subjects with phenomenal properties interacting according to the structure and dynamics of physics. On this view, there would still be physical entities such as electrons— they would just be realized mentally. It could even turn out that a certain sort of cosmopsychism is true, and the physics of the universe as a whole is realized by the structure and dynamics within a single mind. Again, there would still be physical entities on this view, even though they are mentally realized. The demon and dream scenarios are structurally analogous: in these scenarios, physical entities really exist, playing the roles they are supposed to play; it is just that these entities turn out to be mental. So the possibility of these scenarios does not threaten knowledge of ordinary physical claims; at most, they threaten knowledge of more theoretical claims, such as the claim that the physical world exists mind-independently.

Extending the argument to other Cartesian scenarios

We have so far considered full-scale Cartesian scenarios: those that involve some sort of complete model of a universe. Not all Cartesian scenarios are full-scale Cartesian scenarios. There are also limited Cartesian scenarios, in which only part of the universe is modeled in detail (just a single room, say, or just the solar system), or in which it is modeled only at a limited macroscopic scale. There may also be Cartesian scenarios with no model at all: for example, those in which the stream of experiences develops as it does through random coincidence. What should we say about these scenarios?

Limited Cartesian scenarios require fewer resources to construct than full-scale scenarios, but they are also likely to be imperfect. With a limited model, it will be difficult to properly model interactions with parts of the universe not included in the model. For example, what happens if

someone in the room or the solar system carefully observes the world outside, or attempts to travel there? What if the beings in a macroscopic simulation perform repeated scientific observations on the microscopic level? Perhaps there may be ways to do this in an approximate way and to construct more detailed models when needed, but it will be highly nontrivial to do so in a way that yields a fully convincing and coherent stream of experiences. Still, I will set aside these worries and address the status of limited Cartesian scenarios, focusing on the case of a simulation.

One sort of limited Cartesian scenario is a local simulation, in which only part of the world is modeled in detail, and in which the rest is modeled much more coarsely. This case is broadly analogous to a physical scenario in which only the solar system exists and in which observations of the universe beyond are grounded in a coarse-grained simulation. In this case, we would almost certainly say that ordinary physical claims about our environment on Earth are true (there is a computer on my desk, and so on), although various other things we believe, for example about other stars and galaxies, may be false. This suggests that the core structural role associated with physical terms need only apply locally, not globally. If this is right, then in a local simulation, the Ramsey sentence for our physical terms will still be satisfied, and most ordinary physical claims (except those about the distant universe) will be true.

Another sort of limited Cartesian scenario is a macroscopic simulation, in which only macroscopic physical processes are simulated in detail, and microscopic processes will be simulated just well enough to account for certain macroscopic effects. This case is broadly analogous to a physical scenario in which macroscopic objects really exist but in which microscopic objects may not. In this case, we would still allow that there are physical objects around us, and that for example there is an object with a certain mass a certain distance in front of me. At the same time, we might well deny the existence of microscopic objects and microphysical processes, as well as denying various laws that govern the microphysical level (which turn out only to be instrumentally valid). We would say that objects still have mass and perhaps charge, but perhaps not microphysics-specific properties such as spin. This suggests that for the application of ordinary physical terms such as ‘mass’ and ‘distance’, fine-grained microphysical theories involving these terms do not play an essential role, with macrophysical principles and connections to experience playing a more central role. If the same applies in a high-level simulation, these terms will pick out our virtual properties that play these high-level roles, and most ordinary physical claims at the macroscopic level will remain true.

We can also combine the two sorts of scenarios, yielding a local and macroscopic scenario. In this case many claims about the faraway and the very small may be false, but claims about

one's immediate physical environment still seem likely to be true, undercutting global skepticism. All this applies equally to other styles of Cartesian scenarios, including local and/or macroscopic versions of the dream and demon hypotheses. So none of these support global skepticism.

Another relevant class of scenarios include those where one is *currently* dreaming, or where one has recently entered a simulation, or when an evil demon took over one's stream of experiences only in the recent past. On the most natural versions of these scenarios, one's previous life was normal, with experiences not generated by dreams, demons, or simulations. In these cases, some of one's current beliefs based on perception and short-term memory may be false, but one will at least have correct memories of the past (that one was born in a certain location, say) and many correct standing beliefs (that Paris is in France, say). In another version of the case, one was only recently created, with all one's memories in place. In that case, some sort of model will be required to create one's beliefs and one's memories, and the model will have the same sort of structure present in other Cartesian models, so I am inclined to think that many of one's beliefs will be true in virtue of the model. Either way, many of one's beliefs will be true, and global skepticism does not threaten.

One can devise disjunctive scenarios in which one's experiences have many sorts of cause—perhaps a demon one day, a dream the next day, a simulation the next day, and so on. If one sort of cause is dominant here, then perhaps as in the recent envatment case, one's beliefs grounded in the dominant source will be true, and the others false. If no sort of cause is dominant, then one can ask what explains the great similarity between the different causes, so that they generate a highly regular stream of experiences. Perhaps the most natural explanation is that there is a single model generating all of these specific causes (a master model used by the demon, the dream, and so on). If so, then the relevant structure can be found in that model, and it will plausibly make one's beliefs about the world mostly correct.

Something similar to this applies to any scenario in which there is a systematic explanation of the regularities in our experience. As long as there is some reasonable explanation, there must be relevant structure in the environment. Then much of the structure required for the truth of our physical beliefs will be present, and many of those beliefs will be true. So no such scenario will be a global skeptical scenario.

What if there is no good explanation of the regularities in our experience? For example, what if there is random physical stimulation that just happens to produce our highly coherent stream of consciousness? As an extreme case, what if there is no external world at all, and just random dynamics for experience, which by complete coincidence yields a regular stream of experience?

I think we should concede that these no-explanation scenarios are scenarios in which one's ordinary physical beliefs are false. These beliefs require some structure in the environment and some normal causes of our experiences, but in this case there is no such structure and no such normal causes. So veridicalism does not apply to these scenarios. Instead, I think that these scenarios can be ruled out by inference to the best explanation. Precisely because these scenarios have no good explanation of our experiences, and many better explanations are available, they can be ruled out by abduction.¹⁴

Abductive replies to Cartesian skepticism are already familiar from Russell and others, who argue that demon hypotheses, dream hypotheses, and other hypotheses give less good explanations of the regularities in our experiences than non-Cartesian hypotheses. One recurrent problem is that it is far from clear that these explanations are so much worse; the simulation explanation especially can seem quite good. The current use of abduction requires a far weaker use of the principle, to rule out extreme cases in which no systematic explanation is available.

In effect, structuralism and abductivism combine to deal with Cartesian skepticism. Explanatory scenarios—those in which there is some explanation of one's experiences—are handled by structuralism, which yields a veridicalist line on these scenarios. Non-explanatory scenarios—those in which there is no explanation of the regularities in our experience—are handled by abductivism, which allows us to exclude these scenarios. We can sum things up in an argument as follows:

- (1) In explanatory Cartesian scenarios, many of our physical beliefs are true.
- (2) Non-explanatory Cartesian scenarios are ruled out by abduction.
-
- (3) All Cartesian scenarios are either ruled out by abduction or are scenarios where most of our physical beliefs are true.

¹⁴What about Boltzmann brains? These are cases where one's brain is assembled randomly, so there is no good explanation except randomness for one's experience. Certain physical theories predict that in an infinite universe, Boltzmann brains just like mine will appear an infinite number of times. If so, one might suggest that abduction does not rule out the Boltzmann brain hypothesis, but may even support it: at least, abduction from our experience may support a physical theory which supports the Boltzmann brain hypothesis. In response, I think that the chain of reasoning that assigns the Boltzmann brain high probability should be rejected, not least because this assignment is unstable (if we are Boltzmann brains our inferences to physical theories cannot be trusted). More generally, in Boltzmann brain universes the vast majority of experiencers will have irregular experiences, so the fact that we have regular experiences tends to disconfirm the Boltzmann brain hypothesis and any theory that supports it.

In this way structuralism and abductivism can come together to yield a comprehensive response to global skepticism.

Role Structuralism and Realizer Structuralism

It may be objected that the argument collapses the distinction between simulation and reality. For example, in reality the Empire State Building has a mass of 331 million kilograms. In principle we could build a perfect computational simulation of the Empire State Building, but the simulation need not have a mass of 331 million kilograms. It might be implemented in computer circuitry weighing only a few grams. It is not entirely clear what if any mass properties the virtual objects involved in the simulation have, but it seems clear that none of them has mass of hundreds of millions of kilograms. But by the arguments in the previous section, the simulation will have all the relevant structural properties of reality, and a virtual object will have the same relevant structural properties as the Empire State Building. So something must have gone wrong: either mass is not a structural property, or structural properties are not duplicated in simulations.

In response, we have already seen that conceptual structuralism about the physical is not committed to the claim that physical properties are structural properties. The Ramsey method as developed above identifies physical properties such as mass as the properties that play a certain structural role. On this approach, mass is not itself the structural property of playing that role, or of having a property of playing that role. Rather, it is the *realizer* of the role: the property that in fact plays that role. The realizer property will typically be entirely distinct from the role property. In some cases the realizer property may be a distinct structural property, while in other cases it may be a nonstructural property: the structuralism I have developed is consistent with both possibilities. But either way, we cannot infer from sameness of role property to sameness of realizer property.

The role/realizer distinction enables us to defuse the objection. When we say a simulation has the same structure as a non-simulation, we are saying the two have the same structural role properties. For example, where mass is concerned, both will instantiate the higher-order role property of having a realizer property that plays the mass role, where this role is specified structurally. But the simulation and the non-simulation will not have the same realizer properties. In a simulation, the realizer property for mass will be some sort of complex computational property. In a non-simulation, the realizer property for mass will be a noncomputational physical property (at least assuming physics is not computational).

On this approach, mass will be identified as the property that plays the mass role in our situation in the actual world. If we are not in a simulation or a world with computational physics, the property of having a mass of 331 million kg will be a noncomputational property. If we build a perfect simulation of the Empire State Building with a virtual object corresponding to the building, the virtual object will not have this noncomputational property, so it will not have mass 331 million kg. It will have a different computational property that plays the role in the simulation, a property we might call virtual mass. But virtual mass will be quite different from mass.

On the other hand, if we are in fact in a simulation, then the property that plays the mass role will itself be a computational property, so it will be correct to say that mass is a computational property. More generally, we can say that in the simulation scenario (considered as an epistemic possibility), mass is a computational property. This is because in any scenario, ‘mass’ picks out whatever plays the mass role in that scenario.

All this is analogous to what goes on with a term such as ‘heat’, which we can think of as picking out whatever plays the heat role (which involves causing certain sensations in us). As it happens, the motion of molecules plays the heat role. But it could have turned out that something else, say PQR, played the heat role. If we are in fact in a scenario where PQR plays the heat role, then it will be correct to say that heat is PQR. So in the PQR scenario, heat is PQR. But given that we are actually in a situation in which molecular motion plays the heat role, heat is the motion of molecules. As Kripke put things, it is a priori that heat is whatever plays the heat role, and given that molecular motion plays the heat role, it is necessary that heat is molecular motion. In the current framework, it is a priori that mass plays the mass role, and given that a certain property plays that role, it is necessary that mass is that property.

This phenomenon gives rise to a certain sort of Necker-cube phenomenon in thinking about skeptical scenarios. When we think about the building simulation scenario as a distant part of our world, or as a counterfactual possible world, we describe it as one in which mass is a nonvirtual property and in which nothing weighs 331 million tons. On the other hand, when we think of it as an epistemic possibility concerning the situation we may actually be in, we describe it as a scenario in which mass is a virtual property and in which the building weighs 331 million tons. This is again analogous to the Kripkean cases: if we consider Putnam’s Twin Earth as a distant part of our world, or a counterfactual metaphysically possible world, we describe it as one in which water is still H₂O and the lakes do not contain water. But if we consider it as an epistemic possibility concerning the situation we may actually be in (it may turn out that the lakes contain XYZ), we describe it as a scenario in which water is XYZ and the lakes contain water. After all, if we accept

that we have been on Twin Earth all along, we should accept that water is XYZ. Likewise, if we accept that we have been in a simulation all along, we should accept that mass is a virtual property.

One subtlety concerns what goes on if there are multiple properties playing the same structural role in the actual world. For example, if we are in a nonsimulated world that also contains simulations, there may be two properties playing the mass role, and so two candidates for mass. The problem might be all the worse if we are in a simulation and there are multiple other simulations as well as a non-simulated world containing us. To handle situations like this, it is important that out structuralism includes an indexical component: mass is picked out as the property that plays the mass role for *us* (or perhaps for *me*), in that it is the property that causes relevant experiences in us, that governs interactions in the objects I observe, and so on. Then if we are not in a simulation, ‘mass’ will pick out a noncomputational property, and if we are in a simulation, ‘mass’ will pick out a computational property instantiated in that simulation.

We might call this framework *realizer structuralism*: physical properties are picked as those realizer properties that play a certain role. This view is analogous with realizer functionalism in the philosophy of mind, where mental states are identified as those properties (typically physical properties) that play a certain functional role. Realizer functionalism contrasts with role functionalism, where mental states are identified with role properties, such as the property of having some realizer property that plays a certain functional role. By analogy, we can define *role structuralism* as the view that physical properties are identified with role properties, such as the property of having some realizer property that plays a certain structural role.¹⁵

In the philosophy of mind, role functionalism is somewhat more popular than realizer functionalism. The main reason is that there is a widely shared intuition that two functional duplicates (a brain and its silicon functional isomorph, for example) may have the same mental states. In these duplicates the role properties are the same, but the realizer properties are different. Given that the mental properties are the same, it makes sense to identify mental properties with role properties. In the case of the physical, the situation is somewhat reversed. In this case, as we have seen, there is a strong intuition that structural duplicates (a simulation and a non-simulation, for

¹⁵Why “structuralism” in the physical case and “functionalism” in the mental case? One reason is just that different terms are used in the different traditions. A somewhat more substantive reason is that functionalism is typically understood to involve a commitment to the roles in question being causal roles. Structuralism about the physical can be developed with causal roles, but it need not be: there are forms of structuralism where the roles are defined without the use of causal notions, and even without the use of nomic notions. The more general term “structuralism” serves better to accommodate views such as these. That said, the terms can be used in various ways, and it is not entirely unreasonable to understand structuralism as a sort of functionalism about the physical.

example) have quite different physical properties. To accommodate this intuition, it makes sense to identify physical properties with realizer properties.

That said, there are certainly some concepts for which something like role structuralism is correct. One plausible case, as we have already seen, involves mental concepts such as belief, say, which arguably picks out a role property. Another case is the concept of a computer, or of computation, which is plausibly definable as a role property involving the existence of certain realizer properties that play certain complex roles. It may be that various social properties, such as that of being a philosopher, are also best understood as role properties.

Importantly, the anti-skeptical argument I have given above works equally well given both role structuralism and realizer structuralism. All that is required is conceptual equivalence between physical terms and structural sentences. We will get this equivalence whether mass is equivalent to “the property that actually plays the mass role” or “the property of having some property that plays the mass role”. If mass was identical to the role property, the anti-skeptical argument would be even more straightforward: in an evil-demon situation, the same role property would be present, and the same mass facts would be present. But realizer structuralism is good enough, since this has the consequence that if we are actually in the evil-demon situation, objects still have mass.

Anti-Structuralism about Space

There remains a strong intuition that Cartesian scenarios are skeptical scenarios, and that ordinary physical claims are false in Cartesian scenarios. Many people think it is obvious that even if we are in a simulation or being stimulated by an evil demon, most of our physical beliefs are false. Sometimes the intuition is expressed in terms of notions such as solidity and heaviness: if I am in a Cartesian scenario, objects are simply not solid or heavy as they seem to be. Sometimes it is expressed in spatial terms: if I am in a Cartesian scenario, there are not the objects out there in space in front of me as there seem to be.

The intuitions here are distinct from and run deeper than those in the previous section. They are not addressed simply by distinguishing role and realizer structuralism, and they arise just as strongly when we think of the Cartesian scenario explicitly as a first-person epistemic possibility. These intuitions reflect anti-structuralist intuitions about physical notions such as mass and space. The underlying thought is that we have a nonstructural grasp of what it is to have mass or to occupy space. We do not just grasp mass as what plays the mass role, and we do not just grasp space as what plays the space role. Rather, we have some sort of grasp of mass and space as they are in

themselves.

In the current framework, this comes out by saying that notions such as mass and space should not be Ramsified. Certainly it is formally possible to Ramsify mass and space, finding a mass role and a space role, and formulating Ramsey sentences about what plays these roles. But on the anti-structuralist view, “the property that plays the mass role” will not be equivalent to “mass”, and “the property that plays the space role” equivalent to “space”, precisely because we have a nonstructural grasp on mass and space.

This move can be made at the theoretical level, holding that mass and space should not be Ramsified away within theories. It can also be made at the observational level, holding that mass and space are observational concepts that should not be Ramsified away. I will discuss the observational version of the objection, but most of what I say about the case of mass and space below will apply to the theoretical version as well.

A move like this could in principle be made for any observational concept, including that of color. An opponent might hold that we have a direct and nonstructural grasp of redness as an intrinsic quality of things in the world, and that ‘red’ is not equivalent to anything like ‘the normal cause of experiences as of redness’. If this is right, then the structure sentences that are equivalent to physical sentences may include ‘red’ as an ineliminable observational term. If so, then premise 2 will be implausible or at least unjustified, since it is implausible or at least so far unjustified that Cartesian scenarios contain things with the same colors as corresponding non-Cartesian scenarios. On the face of it the relevant virtual objects need not be red, for example. So it is at least far from obvious that structure sentences involving redness will be satisfied in the scenario.

In the case of color, there is a natural response. Perhaps there is some intuition that we have a nonstructural concept of redness as a property of things in the world: for example, we arguably grasp as a certain primitive quality that I have called Edenic redness in experience. But there is little reason to think that this quality is instantiated, or that it is required for our ordinary color claims to be correct. Science suggests that colors experiences are caused by complex physical properties with little resemblance to Edenic redness. Even so, we still correctly say that apples are red. So our ordinary concept of redness seems not to be the nonstructural concept. It is something much more like the structural concept that picks out whatever property normally causes a certain sort of experience. Correspondingly, it is plausibly the structural concept of redness that plays a role in observationally constraining our conception of the physical world.

Something similar applies to mass, and to related observational concepts such as heaviness and solidity. Here again there may be some intuition that we have a nonstructural concept of intrinsic

massiveness or solidity. But even when it turns out that objects are mostly empty space, we still say they are solid. The reason is that our central concept of solidity is structural: it is roughly the concept of what resists penetration in certain ways. The same goes for mass. Our central concept of mass is structural: it is roughly the concept of what is hard to accelerate, what is subject to gravity, and what normally causes certain experiences of heaviness. There is not much reason to think that something beyond this is required to satisfy our claims about mass.

The hardest case is that of space. Here, as in the case of color, there can be a strong intuition that we have a direct and nonstructural grasp of properties like shape, size, and relative location. Our grasp on what it is to be square, for example, is not reducible to a grasp on what causes square experiences. And it may be held that perhaps unlike the case of color, there is a reasonable case that objects in the world really have these special spatial properties, and that they are required to do so in order for our ordinary claims about the physical world to be true.¹⁶

I think the case of space is broadly analogous to the case of color. Perhaps ordinary experience affords us a nonstructural grasp of primitive spatial properties involving what we might call Edenic space. But there is little reason to think these properties are instantiated in the world. For a start, relativity theory suggests that spatial properties are not absolute but instead are relative to reference frames, while widely held versions of quantum mechanics and string theory suggests that space may not be primitive at all. These developments suggest that our world may be quite unlike a world with Edenic space, but they do little to suggest that our world is not a spatial world. Even if we are in a string-theoretic, quantum mechanical, relativistic world without Edenic space, there still can be people who are six feet tall. If so, these nonstructural properties are not required for our ordinary claims about space to be correct.

In other work, I have argued for *spatial functionalism*: the view that spatial concepts are functional concepts, picking out spatial properties in virtue of their functional role. In particular, spatial properties are picked out by their role in motion (roughly, objects tend to change their spatial properties continuously), interaction (roughly, objects close together tend to have stronger interactions), and observation (roughly, spatial properties are those that normally bring about certain

¹⁶Jonathan Vogel (forthcoming) argues against structuralism about space and the associated response to skepticism. He focuses largely on role structuralism (where spatial properties are identical to structural properties) rather than realizer structuralism, but some of his arguments may also bear on realizer structuralism. Christopher Peacocke (2014) argues against spatial functionalism from a phenomenal externalist viewpoint on which spatial experience is directly constituted by external properties. I respond to these arguments in Chalmers 2014. That said, Vogel's abductivism and Peacocke's phenomenal externalism have their own resources as non-structuralist responses to skepticism.

spatial experiences). This view can be supported by an examination of the role of space in various physical theories, as well as by appeal to thought-experiments in which different properties play the relevant roles. If this is right, then our core grasp on space is structural. Like mass and color, spatial terms can be Ramsified out of our physical theories.

Perhaps the best way for an opponent to resist the arguments of this paper is to deny spatial functionalism and to argue for spatial primitivism, on which we have a primitive nonstructural grasp of space. My own view is that as in the case of color, we may well have a nonstructural concept of space (Edenic space), and that beliefs involving that concept may well be false in Cartesian scenarios. However, beliefs involving that concept are also false in the actual world, since the world of modern physics does not seem to be a world of Edenic space. So if we think our ordinary spatial beliefs involve this nonstructural concept, we should be prepared for the consequence that our ordinary spatial beliefs are false in the actual world. The consequence then is not so much skepticism but a sort of negative dogmatism, on which we have good reason to think that our ordinary physical beliefs are false.

In any case, all this affords us something of a diagnosis of our Cartesian skeptical intuitions. There is a strong intuition that we have a nonstructural grasp of space. This intuition explains our intuition that in Cartesian scenarios, our spatial beliefs are massively false: there are not objects where there seem to be. That intuition in turn explains much of our intuition that in Cartesian scenarios, we are massively wrong about the world. The view I have advocated here accommodates an element of that intuition: in Cartesian scenarios, objects do not have the primitive spatial properties that they seem to have. But in this respect Cartesian scenarios are no worse than the actual world. And as in the case of color, the fact that objects do not have the primitive spatial properties does not entail that our ordinary spatial claims and beliefs are false.

It may be objected that the case for spatial functionalism is partly empirical. Insofar as the argument turns on certain physical theories supported by empirical evidence, it has limited force against a skeptic, since the skeptic will not accept the theories in question. However, the case for spatial functionalism does not turn on accepting these theories. Our reaction to these theories (that even if they are true, our spatial beliefs are still true) brings out something about our spatial *concepts*: roughly that their application does not require the existence of primitive space. This conceptual moral holds whether the physical theories in question are true or false, and whether we are in a Cartesian scenario or not.

It may also be objected that our spatial concepts may have changed or developed under empirical pressure, so that even if now our spatial concepts are functional concepts, our initial concepts

were not. I do not find this diagnosis especially plausible: I think that even using our pretheoretical concepts, we were plausibly disposed to react in this way to the relevant empirical situations all along. But even if this claim is correct, it does not do anything to support skepticism about the beliefs we have now, which is the sort of skepticism that is the most concerning. At most, it supports skepticism about beliefs we had in the past, which is a much weaker form of skepticism.

That said, one can at least raise the question of the epistemic status of our nonstructural beliefs in light of skeptical considerations. I have said we have good reason to accept that there are no Edenic colors and no Edenic space, but insofar as these reasons rest on accepting physical theories, a skeptic may reject them. (Perhaps an evil demon in Eden is trying to mislead us into thinking we are in a non-Edenic world?) Still, I think that even prior to accepting these physical theories, we had reasons to be doubtful of our Edenic beliefs. Relatively straightforward reflection suggests that many different sorts of properties could be causing our color experiences or our spatial experiences, and there is not much strong positive reason to think that the causal basis of these experiences should be Edenic color or Edenic space. So I think that antecedently, some sort of skeptical doubt is the right attitude where nonstructural beliefs are concerned.

One might think of this as yielding a mixed verdict on skepticism that the nonstructuralist and I can share: skepticism about nonstructural beliefs, nonskepticism about structural beliefs. The main difference is that I think our ordinary beliefs about the external world are best understood as largely structural beliefs. The situation roughly parallels two different structural realist lines about scientific theories. One line (taken by Worrall and others) holds that scientific theories have both nonstructural and structural elements (and that the structural elements are more secure and robust than the nonstructural elements), while another line (taken by Russell and others) holds that scientific theories are structural through and through. The two lines can at least agree that structural claims about the physical world are better justified than corresponding nonstructural claims.

All this suggests that even if I am wrong that our ordinary physical beliefs are equivalent to structural beliefs, the structuralist response to skepticism still has some force. Insofar as the second non-structuralist premise of the argument is correct, so that structural beliefs are true in Cartesian scenarios, then at least structural beliefs form a core set of beliefs about the external world (namely structural beliefs) that are not threatened by Cartesian arguments. Of course these beliefs may still be subject to other skeptical worries (including local skeptical scenarios and the like), but this line gives at least some purchase against the global skeptic.

Other Anti-Structuralist Objections¹⁷

Other objections can be made to the structuralism I have relied on. One objection is that in arguing that ordinary physical concepts such as mass and space are equivalent to structural concepts, I have relied on more than standard structural realism. Standard structural realism makes claims about physical theories (relativity or quantum mechanics, say) but not about ordinary physical concepts.

An obvious reply is that these physical theories use notions like mass and space, so that it is not easy to see how structural realism could be true of the theories without these being structural concepts. An opponent might reply, though, by saying that the theories use a technical concept (relativistic space or gravitational mass, say) and that structuralism about this concept does not imply structuralism about the ordinary concept. Now, insofar as physical theories are structural and the rest of my argument works, this line will presumably concede that the theories are true at least in full-scale Cartesian scenarios. This seems already a large concession for a Cartesian skeptic, and it will not be easy to maintain a view where physical theories are true in these scenarios but ordinary physical claims are not: in these scenarios, there will be gravitational mass but no mass, and relativistic space but no space? Alternatively, the opponent can argue that even the theories use nonstructural concepts (perhaps because they are tied to observation), but then we are back to the line of objection in the previous section. In any case, even if structuralism about ordinary physical concepts makes a stronger claim than structural realism about physical theories, I think the claim is still highly plausible and can be defended against alternatives among the lines above.

A connected objection is that my structuralism about physical concepts requires descriptivism of a sort that has been shown to be false. According to structuralism, 'mass' is equivalent to a descriptive concept 'the property that plays the mass role', and so on. According to the objection, this view has been ruled out by anti-descriptivist arguments such as Kripke's modal and epistemic arguments. Now, it is not obvious how Kripke's arguments apply to cases such as mass and space, and in any case it is clear that the modal argument does not have much force against realizer structuralism, which does not claim that it is necessary that mass plays the mass role. But perhaps a version of Kripke's epistemic argument might be mounted to suggest that for any specific finite description, we can find scenarios in which there is something satisfying the description but no mass, or vice versa. If so, mass is not a priori equivalent to 'the property that plays the mass role'.

In response: structuralism does not require that 'mass' be equivalent to a finite description. Even if it can be only approximated by descriptions, successive approximations will serve the

¹⁷This section was removed from the published version for reasons of space.

purpose of the argument. Or better, we can do without descriptions altogether. All we need is the weaker claim that truths about mass (and about space and so on) are *scrutable* from a full enough specification of the world in structural terms: roughly, given knowledge of the world in structural terms, we can determine the truths about mass. I argue in *Constructing the World* that scrutability can do much of the work of descriptivism without many of the costs: for example, it is not threatened by Kripke's modal and epistemic arguments. In excursus 15 of that book, I mount a version of the structuralist response to skepticism in terms of scrutability alone.

A still broader line of objection is that the argument relies on claims about the a priori: if not the claim that physical sentences are a priori equivalent to structure sentences, then the claim that they are a priori scrutable from structure sentences. I have argued at length that these claims are plausible, but it is interesting to see how far one can get without any appeal to the a priori. For example, one could still argue that we can know on philosophical (albeit empirical) grounds that mass is what plays the mass role and that our claims about the external world are not globally false. Of course if these claims are known only empirically, their basis may be questioned by the skeptic. The same goes for other core philosophical claims behind the argument. Still, a case against the skeptic that rests only on plausible (if empirical) philosophical background principles is better than nothing. For some naturalists, this may count as at least some purchase against the skeptic.

Another objection is that structuralism does not allow the possibility of local illusions.¹⁸ Say that our structuralism says that 'red' is equivalent to 'what normally causes reddish experiences'. Then we can find a disjunctive property that *always* causes reddish experiences: just take the disjunction of every property that ever causes such an experience. Such a property will presumably be the most normal cause of reddish experiences, so it will be redness. The result is that every object that ever causes a reddish experiences will be red, so there will be no color illusions.

This is in effect a version of Newman's problem for structuralism, and can be responded to the same way. At least part of the redness role requires being a reasonably natural property, and the disjunctive property in question will not be at all natural. In some cases there can be some pressure toward disjunctiveness, but the pressure is limited. In cases like this, we seek a balance between normalcy and naturalness. The highly disjunctive view privileges normalcy at the expense of naturalness and should be rejected. That said, I do not have any explicit analysis of just how to analyze structural concepts in terms of naturalness and normalcy. Articulating the specific

¹⁸Thanks to Nick diBella, Melissa Ebbers, and Alex Sandgren for pressing this objection.

conditions for reference in cases like this is a challenge for theories of reference quite generally, and I am not offering an answer to it, but I do not think the challenge is harder for structuralism than for other theories of reference.

A final objection is that the structuralist response to skepticism is overly intellectualist, since it requires subjects to know the truth of structuralism and of veridicalism in order to have knowledge of the external world.¹⁹ In response: there is no such requirement. To undermine the traditional Cartesian argument at the start of the paper, it suffices to undermine the case for the second premise (if we cannot know that the evil-demon hypothesis is false, we cannot know we have hands). The case for this premise rests on a presumed entailment from the evil-demon hypothesis to our having hands. If structuralism is true, there is no such entailment—at least for the neutrally stated evil-demon hypothesis—and the traditional argument is undermined. This removes the corresponding threat to knowledge even for subjects who have never entertained structuralism or skepticism.

A refined version of the objection appeals to a non-neutrally stated evil-demon hypothesis according to which the demon is globally deceiving us (we do not have hands, and so on). As discussed earlier, a structuralist will hold that hypotheses like this are incoherent. The objector responds that if a subject does not know that structuralism or incoherentism is true, their mere truth will not enable them to exclude the hypothesis in question. If they cannot exclude it, they are subject to a version of the original argument involving this hypothesis, leading to global skepticism.

At this point the opponent is raising a general issue for incoherentist responses for skepticism: do they vindicate knowledge of the external world only for subjects who know that the putative skeptical hypotheses are incoherent? I think not. The closure principles needed to justify the second premise here—for example, if one cannot know q , and $\neg q$ entails $\neg p$, then one cannot know p —are certainly false when the hypothesis q is incoherent. If it were true, our inability to know certain highly complex logical theorems would entail that we have no knowledge at all, since the negation of these theorems (which we cannot exclude) entails the negation of all propositions. *Perhaps* for subjects who actually entertain the skeptical argument and believe the incoherent skeptical hypothesis to be coherent, there is a remaining threat to their ordinary knowledge that requires more than structuralism to overcome. But for ordinary subjects who are not thinking about these matters, it is hard to see a threat to their knowledge here.

¹⁹Thanks to Michelle Dyke and Stephen Schiffer for raising this objection.

Structuralist and Externalist Anti-Skepticism

The structuralist reply to global skepticism shares some elements with the externalist response discussed earlier. Both views hold that brains in vats have true beliefs about their environment, and refer to causally connected properties in their environment that are distinct from the properties that their nonenvatted duplicates refer to. At the same time, there are also important differences. We can see this by first tracing the connections between structuralism and externalism more generally.

Structuralism is consistent with both externalism and internalism about reference. To a first approximation, realizer structuralism goes naturally with externalism, while role structuralism goes naturally with internalism.

Where realizer structuralism applies to a given expression, the expression will pick out the realizer that plays a certain role in the speaker's situation. This realizer will typically vary in different environments even if the speakers are identical. For example, 'mass' may pick out a certain physical property for me and a certain virtual property for my brain-in-vat twin, who inhabits an environment where the virtual property plays the relevant structural role. So realizer structuralism leads naturally to externalism about reference. By contrast, role structuralism leads naturally to internalism, since my twin and I will pick out the same role properties, at least as far as the framework described so far is concerned. For example, 'computer' may pick out the same role property for both me and my twin.

The mapping between role/realizer structuralism and internalism/externalism is not entirely cut and dried. There could be cases or views where realizer structuralist is internalist, perhaps because only one property can realize the relevant role in any environment. There could also be cases or views where role structuralism is externalist, because the role itself depends on environmental factors (perhaps because one is externalist about structural concepts such as causation). Still, there is at least a natural correlation between the two.

Importantly, because structuralism is consistent with both internalism and externalism, the structuralist reply to skepticism does not turn on externalism. We have seen that whether role structuralism (internalist) or realizer structuralist (externalist) is true, the anti-skeptical strategy works equally well either way. In numerous cases (such as 'mass'), the realizer structuralist view is more plausible than the role structuralist view, but if we had spoken a different language in which 'mass' had picked out the role property rather than the associated realizer property, the argument against skepticism would still have gone through. So it is really the structuralism that is doing the central work in responding to the skeptic, not the externalism.

The differences between structuralism and externalism also allow the structuralist to reply easily to the objections to externalist replies to skepticism raised earlier. The first objection, that Putnam's externalist route to veridicalism requires a crude and implausible causal theory of reference on which causal connections suffice for reference, does not arise. The structuralist reply invokes a much less crude and more plausible structuralist theory of reference. The second objection, saying that the externalist reply cannot handle skepticism about non-Twin-Earthable claims such as 'I am talking with four philosophers', is handled straightforwardly by structuralism. The structuralist will understand non-Twin-Earthable claims by the role structuralist framework, which allows twins to be referring to the same role properties which are instantiated in both ordinary situations and Cartesian situations.

Another difference is that where externalism tends to lead to an incoherentist reply to skepticism, as with Putnam's argument that we are not brains in vats, structuralism has much less pressure in this direction. For skeptical hypotheses expressed using terms for which externalism is correct, such as 'I am a brain in a vat', we can give a realizer structuralist analysis and it may turn out that the hypothesis can be excluded. But there will be many other Cartesian hypotheses expressible in internalist terms, such as 'I am in a computer simulation'. Structuralism will give a role structuralist analysis to these hypotheses, and there is no pressure toward incoherentism. Beings in computer simulations can coherently entertain the hypothesis that they are in computer simulations, and so can we. The more important work in replying to the skeptic on this view is done by veridicalism, not by incoherentism.

All that said, it is certainly possible for externalists and structuralists to form an alliance here, mounting a unified structuralist/externalist treatment of skeptical cases. We have seen that the structuralist needs an element of externalism to handle cases in which it seems that simulations have different properties from non-simulations. The externalist can also appeal to structuralism to help spell out the needed constraints on reference, and also to handle those cases where externalism does not seem to apply. On my view it is structuralism that is doing the most work, but externalism is certainly playing a role, and the two views are perhaps strongest together.

Conclusion

I have argued that structuralism about the physical world, helped along by elements of abductivism and externalism, can form the basis of a veridicalist response to global skeptical arguments about the external world. Of course the anti-skeptical force of this response is limited. I have said

nothing to address local skeptical hypotheses, which arguably suffice for the skeptic to cast doubt on all of our beliefs, one at a time. I have said nothing to address skepticism about other minds, or about induction, or about one's own reasoning. And I have said nothing to address forms of skepticism that arises from sources other than the consideration of skeptical hypotheses. It may be that structuralist considerations will help address some of these skeptical worries, but this remains to be seen. Still, this response to global skepticism at least addresses skepticism in one of its most virulent forms and offers a small foothold in replying to skepticism more generally.

References

- Bouwsma, O.K. 1949. Descartes' evil genius. *Philosophical Review* 58:141-51.
- Burge, T. 1979. Individualism and the mental. *Midwest Studies in Psychology* 4: 73-122
- Carnap, R. 1927. *The Logical Structure of the World*. Open Court.
- Chalmers, D.J. 1994. On implementing a computation. *Minds and Machines*.
- Chalmers, D.J. 2003. The matrix as metaphysics. Thematrix.com website. Reprinted in (C. Grau, ed) *Philosophers Explore the Matrix*. Oxford University Press, 2005.
- Chalmers, D.J. 2006. Perception and the fall from Eden. In (T. Gendler and J. Hawthorne, eds.) *Perceptual Experience*. Oxford University Press.
- Chalmers, D.J. 2011. The nature of epistemic space. In (A. Egan and B. Weatherson, eds.) *Epistemic Modality*. Oxford University Press.
- Chalmers, D.J. 2012. *Constructing the World*. Oxford University Press.
- Chalmers, D.J. 2014. The contents of consciousness: Reply to Hellie, Peacocke, and Siegel. *Analysis*.
- Chalmers, D.J. 2017. The virtual and the real. *Disputatio*.
- Davidson, D. 1986. A coherence theory of truth and knowledge. In (E. LePore, eds.) *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Blackwell.
- Frances, B. 2008. *Scepticism comes alive*. Oxford University Press.
- Ladyman, J. 1998. What is structural realism? *Studies in History and Philosophy of Science Part A* 29:409-424.
- Langton, R. 1998. *Kantian Humility: Our Ignorance of Things in Themselves*. Oxford University Press.
- Lewis, D. 1972. Psychophysical and theoretical identifications. *Australasian Journal of Philosophy* 50:249-258.

- Lewis, D. 2009. Ramseyan humility. In (D. Braddon-Mitchell & Robert Nola, eds.) *Conceptual Analysis and Philosophical Naturalism*. MIT Press.
- McDowell, J. 2008. The disjunctive conception of experience as material for a transcendental argument. In Fiona Macpherson & Adrian Haddock (eds.), *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press. pp. 376-389.
- Newman, M. 1928. Mr. Russell's causal theory of perception. *Mind* 5:26-43.
- Peacocke, C. 2014. Phenomenal content, space, and the subject of consciousness. *Analysis*.
- Putnam, H. 1975. The meaning of "meaning". *Minnesota Studies in the Philosophy of Science* 7:131-193.
- Putnam, H. 1981. *Reason, Truth, and History*. Cambridge University Press.
- Ramsey, F. 1931. Theories. In *The Foundations of Mathematics and Other Essays*. Routledge and Kegan Paul.
- Rorty, R. 2003. Davidson versus Descartes. In (J. Malpas, ed.) *Dialogues with Davidson: Acting, Interpreting, Understanding*. MIT Press.
- Russell, B. 1927. *The Analysis of Matter*. Kegan Paul.
- Shoemaker, S. 1980. Causality and properties. In (P. van Inwagen, ed.) *Time and Cause*. Reidel.
- Sklar, L. 1985. Saving the noumena. *Philosophical Topics* 13:89-110.
- Vogel, J. 1990. Cartesian skepticism and inference to the best explanation. *Journal of Philosophy* Vol. 87:658-666.
- Vogel, J. forthcoming. Space, structuralism, and skepticism. *Oxford Studies in Epistemology*, volume 6, 2019.
- Worrall, J. 1989. Structural realism: The best of both worlds? *Dialectica* 43:99-124.