On Being a Lonely Brain-in-a-Vat:
Structuralism, Solipsism, and The Threat from External World Skepticism

Grace Helton

Word Count. 12,390 (incl. all notes)

Abstract. David Chalmers has recently developed a novel strategy of refuting external world skepticism, one he dubs *the structuralist solution*. In this paper, I make three primary claims: First, structuralism does not vindicate knowledge of other minds, even if it is combined with a functionalist approach to the metaphysics of minds. Second, because structuralism does not vindicate knowledge of other minds, the structuralist solution vindicates far less worldly knowledge than we would hope for from a solution to skepticism. For, solipsism threatens surprisingly vast swathes of worldly knowledge across multiple domains, including at least some knowledge about: political affairs, religious practices, artistic movements, historical events, and cultural trends. Third, the significance of these results exceeds their implications for the structuralist solution; these results suggest that *no* solution to external world skepticism which does not also solve the problem of other minds will ultimately yield the kind of solution we might have hoped for. Relatedly, these results suggest that the problem of external world skepticism should perhaps be construed as two different problems, since the problem might turn out to require two substantively different solutions, one for knowledge of the kind that is not dependent on other minds and one for knowledge that is.

Key words. external world skepticism, structuralism, solipsism, social ontology, problem of other minds, Turing Test, blockheads, zombies, faux-folk, David Chalmers

Suppose you are a brain in a vat. Just a pink and wrinkled mound, you bathe in electrical fluids, tethered to pulsating wires which permit a perfectly convincing simulation of an ordinary life. But you don't know that you are a brain-in-a-vat. You think you are embodied in the ordinary way, with limbs and a torso, eyes and glasses, arms and legs. All of your experiences are subjectively indistinguishable from some you might have had if you were embodied in the usual way: You (seem to) see your train pulling into the station, so you (seem to) run towards it. After (seeming to) board the train, you (seem to) sit down onto one of the familiar molded orange-and-beige benches. Feeling relieved, you (seem to) set down your bags and (seem to) open *The Chimpanzees of Gombe* to read. Nothing hints at your envatted state, no clues suggest that you are

not embodied in the normal way, and no evidence could help you confirm or disconfirm the hypothesis that you are a brain-in-a-vat, even if you should try to find some.

Scenarios like the brain-in-a-vat scenario play a central role in the most prominent arguments for *external world skepticism*, the view that vast swathes of your beliefs about the external world fail to amount to knowledge. One such argument runs this way: If you know that you're on the train, then you know that you're not a brain-in-a-vat. But, you don't know that you're not a brain-in-a-vat. So, you don't know that you're on the train. What goes for your belief about being on the train goes for many of your other worldly beliefs, so these other beliefs also fail to amount to knowledge.

David Chalmers has recently developed a novel strategy of refuting arguments like the brain-in-a-vat argument. This is *the structuralist solution to skepticism*.[1] According to the structuralist solution, ordinary concepts, like IS A TRAIN and IS A MOLECULE, pick out either certain third-personally observable behavioral dispositions or else the realizers of those dispositions. In very rough form, the structuralist says that in certain brain-in-a-vat scenarios, things like trains exist because IS A TRAIN—our concept in English—picks out those entities which move on tracks, transport cargo or passengers, and the like. And, in the brain-in-a-vat scenario, there are entities which play these roles. So, in the brain-in-a-vat scenario, there are trains.

Notice that this solution is distinct from the much-discussed semantic externalist solution to skepticism. For, while the semantic externalist maintains that the *brain-in-a-vat's* concept IS A TRAIN has a referent in her environment, the structuralist maintains that the *English speaker's* concept IS A TRAIN has a referent in the brain-in-a-vat scenario. Put otherwise, the semantic externalist argues that 'I am on a train,' the sentence in the brain-in-a-vat's language—call it *BIV-ese*—is true. The structuralist has the resources to argue that 'I am on a train,' the sentence in English, is true. (I further discuss the differences between structuralism and its superficially similar cousin, semantic externalism, in §1.1).

Chalmers' own ambitions for the structuralist solution are to claim that it can vindicate at least some worldly knowledge, not that the structuralist approach can vindicate all worldly knowledge. In this way, Chalmers' ambitions for the view are avowedly modest, in that he does not take the view to be usable to exclude all possible challenges to worldly knowledge.

This paper explores the limitations of structuralism and in particular, whether structuralism might vindicate more worldly knowledge than Chalmers himself suggests that it might. Specifically, it asks whether structuralism might vindicate knowledge of other minds, for instance, if structuralism is combined with a functionalist approach to other minds.[2] And it further asks whether structuralism—if it cannot vindicate knowledge of other minds—might vindicate

---

[1] Here, I focus on Chalmers' (2018, 2022) defense of the view. For distinct but complementary precursors to the structuralist solution, see Chalmers (2003b, 2010: ch. 13, 2012: 431-40). An important antecedent of the structuralist view is Bouwsma (1949).

[2] Chalmers (2018: see, e.g., 660, 2010: 473-78).

enough worldly knowledge so as to constitute a wholly satisfying solution to skepticism (contra Chalmers' comparatively limited ambitions for the view).

In taking up these lines of inquiry, *Chalmers himself is not the target of this paper.* For, Chalmers himself does not take structuralism to vindicate knowledge of other minds, nor does he suppose the view to constitute a wholly satisfying solution to skepticism. Nevertheless, these questions are philosophically significant for three main reasons.

First, one might think, as Chalmers himself suggests, that the structuralist approach would vindicate knowledge of other minds *if* it were combined with a functionalist approach to other minds. While Chalmers himself does not endorse functionalism, many philosophers of mind adopt this approach, and thus, this question is one of considerable interest.[3]

Second, the structuralist approach is a novel and important proposed solution to a perennial and pressing question. As already alluded to, structuralism is distinct from the superficially similar semantic externalist approach, and Chalmers argues that structuralism outperforms semantic externalism in at least some respects. In §1, I describe an additional respect in which structuralism is preferable to its externalist counterpart.[4] Thus, attempting to map more precisely the outside limitations of a novel and promising approach to a significant and seemingly insoluble philosophical problem is intrinsically valuable in its own right. This is so even though the chief defender of this view himself employs the view for comparatively modest aims.

Third, if, as I will argue, structuralism cannot vindicate knowledge of other minds even at the same time that it might vindicate knowledge of non-minded entities, such as tables and quarks, then this points to a surprising possibility about the best total solution to worldly skepticism. Namely, this result opens up the possibility that skepticism about non-minded entities might require a different solution than skepticism about minded entities. Moreover, the points of this paper clear room for the view that skepticism about non-minded entities might require a different solution than skepticism about minded entities *even if we accept a naturalistic metaphysics of mental entities.*

At the heart of this paper is the suggestion that structuralism does not vindicate knowledge of others' mental states *even if* structuralism vindicates knowledge of things like molecules and trains. Very roughly, this is because structuralism appeals to *third-personally* observable dispositions (or that which underlies them), whereas plausible variants of functionalism appeal in part to mental dispositions not reducible to third-personally observable dispositions. This immediately raises the question: Why are things like molecules and trains plausibly reducible to third-personally observable dispositions (or that which underlies them), whereas mental states are not?

---

[3] Chalmers (2018: 645, 660)

[4] Notably, Chalmers himself eschews the thought that his total approach can evade the charge of topic-changing, since his total approach is a hybrid one which contains both structuralist and semantic externalist elements (Chalmers 2018, 2022, personal communication). This is yet another way in which my target is not Chalmers himself but structuralism, understood as a potential total solution to skepticism.

There is a shallow answer to the question and a deep one. The shallow answer is that while it is at least initially feasible that concepts like IS A TRAIN and IS A MOLECULE might pick out some third-personally observable disposition or whatever realizes that disposition, it is not ultimately feasible that concepts like FEARS DYING or FEELS JOY might pick out some third-personally observable disposition or whatever realizes that disposition. Put otherwise, a kind of crude behaviorism is simply implausible as a metaphysics of mind, whatever its plausibility when it comes to non-mental entities, such as molecules and trains.

This shallow answer cries out for a deeper one, one which would suggest *why* concepts like IS A TRAIN and IS A MOLECULE might pick out some third-personally observable disposition or else whatever realizes that disposition, even if it is not ultimately plausible that concepts like FEARS DYING or FEELS JOY might pick out some third-personally observable disposition or whatever realizes that disposition. To answer this question would be to answer a perennial, fraught, and central question of philosophy of mind: What makes a mind a mind and not a molecule, a train, or some other non-mental thing? Fortunately, for present purposes, I needn't attempt a positive answer to this question, one which lies at the very bedrock of theorizing about the mind. I need only defend the 'negative' claim that whatever minds are, they are not (merely) third-personally observable dispositions or else whatever realizes those dispositions. The sophisticated functionalist will deny this, since she will maintain that some aspect of mental states is irreducibly mental and thus, not equivalent to any third-personally observable disposition.

The paper proceeds this way: In §1, I describe the structuralist solution to skepticism, suggesting that, on at least one way this solution might be developed, it has an important advantage over its semantic externalist rival. This is the advantage that structuralism is not unacceptably topic-changing in the way in which semantic externalism is sometimes claimed to be. In §2, I argue that structuralist considerations do not exclude epistemological solipsism, even when combined with a functionalist view of the mental. Very roughly, this is because plausible variants of functionalism posit that the relevant functional roles are at least partly constituted by an irreducible mental role, one not capturable by the kinds of third-personally observable dispositions which figure in a structuralist semantics. In §3, I argue that epistemological solipsism compromises vast swathes of empirical knowledge, including at least some knowledge about political affairs, artistic movements, and cultural events, with the result that solipsism

threatens structuralism's ability to provide the kind of substantive solution to skepticism we would hope for.[5] In §4, I draw out some morals.[6]

1 The Structuralist Solution to Skepticism

External world skepticism is, roughly, the view that vast swathes of your worldly beliefs do not amount to knowledge, whether or not those beliefs are in fact true.[7] Worldly beliefs are *a posteriori* beliefs about the world outside your mind. Going forward, I will sometimes refer to worldly beliefs as 'empirical beliefs,' and I will sometimes abbreviate 'external world skepticism' to 'skepticism.' Since external world skepticism threatens only worldly knowledge, the view leaves open that you might enjoy wide swathes of other sorts of knowledge, such as knowledge about the *a priori* (e.g., *2+2=4*) or knowledge about the introspectable aspects of your own psychology, (e.g., *I feel joy*).

In order to explain the structuralist solution to skepticism, let's return to the brain-in-a-vat argument for skepticism, this time developing it a little more fully, in terms of a closure principle:

> THE BRAIN-IN-A-VAT ARGUMENT
> (1) If: I know p and I know that p entails q, then: I know q.
> (2) I know that: If I'm on the train (and not in a vat), then I'm not a brain-in-a-vat.
> (C1) If I know that I'm on the train (and not in a vat), then I know that I'm not a brain-in-a-vat.
> (3) I don't know that I'm not a brain-in-a-vat.
> (C2) I don't know that I'm on the train (and not in a vat).

---

[5] This complaint is structurally similar to the 'recent envatment' complaint against semantic externalism of the kind defended by Putnam (1981), Davidson (1986), and Rorty (2011). According to this complaint, the externalist approach fails to vindicate knowledge of the kind that is threatened by 'recent envatment' scenarios, with the result that externalism cannot offer a wholly satisfying solution to external world skepticism. See, e.g., Brueckner (1999: 237), Christensen (1993: 314-5), Farrell (1986: 150), Forbes (1995: 207), Glymour (1982: 173-5), Smith (1984:117), Tymoczko (1989: 294-5), and Wright (1992: 86-90). For a recent rival perspective, see Thorpe (2018).

[6] In a companion piece (Helton 2021), I argue from some of the same considerations presented here that epistemological solipsism just *is*, for creatures like ourselves, a genuine form of external world skepticism (Cf. Chalmers 2022: 500-501). One needn't accept that claim to accept the points I make here. In fact, the points of each paper mutually dissociate. For, one might accept my claim that structuralism cannot refute solipsism, with the result that structuralism is not as substantive a solution to skepticism as we might have hoped for, whilst denying that solipsism is itself a form of skepticism. Conversely, one might accept my claim that solipsism is a form of skepticism, whilst denying that structuralism cannot rule solipsism out.

[7] See Helton (2021) for a discussion of various ways of characterizing skepticism.

Since iterations of this argument can be constructed for many, perhaps most, of your worldly views, conjoined iterations of this argument can be used to establish that vast swathes of your worldly beliefs fail to amount to knowledge, even if those beliefs happen to be accurate.

Chalmers rejects (2), the claim that you know that your being on the train entails that you're not a brain-in-a-vat. On Chalmers' view, you don't know this because it's not *true* that if you're on the train (and not a vat), then you're not a brain-in-a-vat. Rather, for Chalmers, you might *both* be on the train *and* a brain-in-a-vat. What would appear to be, at first blush, an impossibility, is not an impossibility. You can both be on a train and embodied in a non-envatted way (in one form of your embodiment) and also a brain-in-a-vat (in another form of your embodiment). For, Chalmers argues, the brain-in-a-vat scenario is populated with things like trains, philosophers, bodies, and triangular things. Chalmers derives this result from *conceptual structuralism* (in abbreviation: *structuralism*).

In this section, I describe structuralism and show how Chalmers uses it in the service of refuting (2). To anticipate: conceptual structuralism is roughly the view that concepts like IS A TRAIN, IS A TRIANGULAR THING, and IS A PHILOSOPHER pick out whatever plays some relevant role or else whatever realizes that role. For example, IS A TRAIN picks out whatever plays the train-role or else what realizes that role. Since the brain-in-a-vat scenario contains entities which play the *is-a-train*-role (and entities which play the *is-a-triangular-thing*-role and entities which play *the is-a-philosopher*-role), beliefs such as *a philosopher is riding the train with the triangular thing on it* turn out to be true in that world.

If Chalmers is right that the brain-in-a-vat scenario is populated with things like trains, then (2) is false. So, even if we accept (1), which is a kind of closure principle, we will never reach (C1), the interim conclusion that if you know you're on the train (and not in a vat), then you know that you're not a brain-in-a-vat. As a result, even if (3) is true, that is, even if you don't know that you're not a brain-in-a-vat, there is no reason to accept the conclusion that you don't that know you're on the train and, on that train, embodied in a normal way.

I first describe conceptual structuralism (§1.1) and then explain how Chalmers employs structuralism against skepticism (§1.2).

1.1 Conceptual Structuralism

In broad form, conceptual structuralism about some concept is the view that that concept's meaning is determined by that concept's role in some relevant theory or set of theories. Specifically, structuralism is consistent with both of the following approaches to reference, here illustrated with respect to the concept HAS MASS:

- The referent of HAS MASS is whatever plays the relevant mass-having-role in the world of evaluation

- The referent of HAS MASS is the realizer of the mass-having-role in whatever world is considered as actual.[8]

On the first theory, HAS MASS picks out whatever plays the mass-having role in some world of evaluation, regardless of whether that entity also plays that role in whatever world is treated as actual. Call this the 'role' approach to reference. On the second theory, HAS MASS picks out whatever plays the mass-having role in whatever world is treated as actual. Call this the 'realizer' approach to reference. Note that the 'realizer' approach permits that the referent of HAS MASS might have a non-structuralist *nature*. That is, this entity might, for all conceptual structuralism says, have a property-less 'core' or *quiddity* which individuates that entity from others. Thus, conceptual structuralism does not entail *ontological structuralism*, the view that all entities are equivalent to or exhausted by some role they play.[9]

Conceptual structuralism is consistent with a variety of competing positions about how the referents of particular concepts are fixed. For instance, conceptual structuralism is consistent with the view that the referents of all concepts are picked out via the 'role' way, the view that the referents of all concepts are picked out via the 'realizer' way, or that some are picked out via the 'role' way and some via the 'realizer' way. Thus, the view can accommodate the widely accepted view that natural kind terms, such as IS WATER and IS GOLD, pick out their realizers in the actual world. But structuralism is *also* consistent with—though it does not entail—the view that other terms, such as IS A CUP or IS A TRAIN, pick out whatever satisfies some relevant role, without any regard for whether the realizers of that role are the same as that in the actual world. We might see it as a virtue, then, of the structuralist approach that the semantics it appeals to is sufficiently schematic so as to appeal to theorists of many stripes.

As mentioned in the paper's introduction, the structuralist solution is superficially similar to the much-discussed semantic externalist solution to skepticism.[10] Like the structuralist, the semantic externalist posits that the brain-in-a-vat's beliefs are largely true. But, the structuralist and the semantic externalist reach this conclusion via very different paths. The semantic externalist draws on a particular view of how concepts' meanings are fixed to derive the result that *in the brain-in-a-vat's language*—call it *BIV-ese*—concepts like IS A TRAIN and IS A MOLECULE have different referents than their counterparts in English. For the externalist, when the brain-in-a-vat has the thought, couched in her concepts, *I'm on the train*, 'train' picks out roughly whatever is the typical cause of relevant train-utterances, so for the BIV, it is true that she is on the train.

Like the semantic externalist, the structuralist arrives at the result that the BIV's beliefs are largely true, but the structuralist does not appeal to BIV-ese to make this claim. She can

---

[8] Chalmers (2018: 649-853).
[9] Chalmers (2018: 638-39).
[10] Locus classici of this view include: Putnam (1981), Davidson (1986), and Rorty (2011).

rather make her claim with respect to whatever English concepts pick out. The structuralist has the resources to claim that *even construed in English*, and not in BIV-ese, the brain-in-a-vat's beliefs will turn out to be largely true. For, conceptual structuralism is the view that *our* ordinary concepts, like IS A TRAIN, have referents in the brain-in-a-vat scenario. For the structuralist, this is because our concept of IS A TRAIN picks out some third-personally observable role or what realizes it, and both that role and its realizers obtain in the brain-in-a-vat scenario. Thus, the structuralist's claim is that, in the BIV world, there are things which are *trains*—where these are the things picked out by our concept IS A TRAIN—and not merely things which 'are trains' where these are whatever is picked out by the brain-in-a-vat's concept IS A TRAIN.

Adjudicating between semantic externalist and structuralist solutions is not the aim of the present paper. However, I would suggest that the following is a reason to prefer structuralism to semantic externalism, one that is distinct from those Chalmers himself describes: A common complaint about externalism is that it is unacceptably topic-changing, in that it answers the question, 'do you know that you are on a train?' with the somewhat peculiar answer 'You know that you are on something which 'is a train," where 'is a train' is the referent of a concept in BIV-ese.' But what you wanted to know was whether you know that you're on something which is a *train*, where this is the referent of the concept in English, not whether you know you're on something which 'is a train,' where 'train' is the referent of a concept in BIV-ese.[11]

Whatever the merit of the charge that externalism is unacceptably topic-changing, notice that structuralism is not necessarily committed to it. For, the structuralist, in her most ambitious mode, aims to show that in the brain-in-a-vat world, you can know that you are on something which is a train, where this is a thing picked out by the English speaker's concept IS A TRAIN, not (merely) that you can know that you are on something which 'is a train,' where this is something picked out by the concept IS A TRAIN in BIV-ese. This is because, for the structuralist, the realizers of the English speaker's concept IS A TRAIN obtain in the brain-in-a-vat scenario. For in the brain-in-a-vat scenario, there is something which moves on tracks and transports people and goods. The structuralist therefore can claim that *the English speaker's* concept IS A TRAIN picks out things with these qualities.[12]

1.2 The Structuralist Solution to Skepticism

With structuralism in hand, Chalmers develops an argument against (2), the claim that you know that if you're on the train (and not in a vat), then you're not a brain-in-a-vat. To do this, Chalmers argues that whether referents of IS A TRAIN are determined in the 'role' way or the 'realizer' way, there are trains—that is, referents of the English speaker's concept IS A

---

[11] For the concern that the externalist solution is unacceptably topic-changing, see Brueckner (1986:164-167) and Forbes (1995); Cf. Brueckner (1992, 2003). Thanks to Harvey Lederman for discussion on this point.

[12] For different reasons to prefer structuralism to semantic externalism, see Chalmers (2018: 629-31, 658-9).

TRAIN—in the brain-in-a-vat scenario. Chalmers further argues that you are embodied in two different ways, in a 'normal' way and in an envatted way. For the sake of concrete discussion, in what follows, I will focus on Chalmers' argument that there is a train in the brain-in-a-vat scenario, setting to one side his more contentious views concerning dual embodiment.

First, suppose that the referent of IS A TRAIN is fixed in the 'role' way; IS A TRAIN picks out whatever (if anything) plays some *is-a-train*-role. To play this role, some entity must presumably do things like run on tracks, transport passengers or cargo, and so on. In the brain-in-a-vat scenario, entities play these roles, so trains exist in this world.

Second, suppose that the referent of IS A TRAIN is fixed in the 'realizer' way. On this view, IS A TRAIN picks out whatever realizes the *is-a-train*-role in some world considered as actual. On the presumption that something realizes the *is-a-train*-role in the brain-in-a-vat world, this view also predicts that IS A TRAIN has a referent in this world.

Note that in order to identify what the referents of our concepts are in the brain-in-a-vat world, we should *not* take the brain-in-a-vat world to be a counterfactual version of the actual world, a way the world might have been but is not. Rather, we should instead consider what would be the case if the *actual* world is the brain-in-a-vat world. (Compare: If, contrary to what we take to be the case, the actual world is one in which the water-role is realized by XYZ, then according to the 'realizer' view of reference, IS WATER refers to those things which are XYZ, not those things which are H_2O.) So, if the brain-in-a-vat world is the actual world, then IS A TRAIN picks out whatever realizes the *is-a-train*-role in that world. Presumably this realizer will be some sort of computational structure.

Chalmers maintains that what goes for IS A TRAIN goes for a very wide range of other concepts, including IS WATER and HAS MASS. Either their referents are determined in the 'role' way, in which case they have referents in the brain-in-a-vat scenario, since this scenario contains entities which play the relevant roles. Or else their referents are determined in the 'realizer' way, in which case they also have referents in the brain-in-a-vat scenario. For, in the case in which the brain-in-a-vat scenario is the actual scenario, the realizer of the relevant concept is fixed relative to the brain-in-a-vat scenario. Thus, Chalmers concludes that in all global skeptical scenarios, many of your beliefs involving ordinary concepts will turn out to be *true*. This isn't to say that all such beliefs will turn out to be true; some, such as those beliefs about the fundamental nature of reality, will turn out to be false. But, vast swathes of your worldly beliefs will turn out to be true, and the most radical skeptical conclusions are forestalled.

I take it, then, that structuralism is a well-motivated, cogent and novel view, one which enjoys at least some virtues over its semantic externalist cousin. I thus take it to be a worthy project to test the outer limits of the view, to see how far the view might stretch in the service of dissolving skepticism. In what follows, I will argue that the view does not vindicate knowledge of other minds, even if combined with a naturalistic functionalist view of the metaphysics of mind. This is because it appeals to *third-personally* observable dispositions, that is, to dispositions viewable 'from the outside.' The view must appeal to such dispositions to vindicate knowledge

of entities such as trains—it is because we can be sure that there are things which move on tracks, for instance, that we can be sure that there are things such as trains. At the same time, this approach does not vindicate knowledge of other minds precisely because the most attractive metaphysics of the mental will invariably posit some mental element—such as a mind-involved element of a functional role—which is not reducible to any third-personal disposition. This is so even on non-reductive naturalistic approaches, such as functionalism; the positing of this unreducible mental element is necessary if such views are to stave off worries of the kind which plagued their behaviorist predecessors.[13]

2 Structuralism Cannot Be Used to Rule Out Solipsism

In this section, I turn to developing the view that structuralism does not refute solipsism, even if we combine structuralism with an independently plausible variant of functionalism about the mental.

2.1 The Lonely Brain-in-a-Vat Argument for Solipsism

To establish structuralism's limitations vis-à-vis solipsism, I will present a traditional argument for solipsism, one that is roughly isomorphic to the traditional argument for skepticism and will explain why structuralism does not offer a way to rebut this argument.[14] The argument in question draws in part on the claim that it is at least in principle consistent with your evidence about the world that you are the only minded creature in the universe. For our purposes, we can focus in particular on the following way the world might be:

> 'BRUTE FORCE' PROGRAMS
> Suppose you are a brain in a vat. Perhaps you were envatted as part of a wide-scale effort by post-human AI to study human psychology, or perhaps you are the haphazardly created side project of a single post-human AI, the equivalent of an amateur scientist's basement experiment. But whoever your creators were, they and their community are no longer around. Maybe their community was destroyed by a meteor or by mutually destructive warfare or by some dramatic shift in the earth's climate, but whatever the cause, your creators have been gone for a very long time. Still, your simulation hums along, a perfectly convincing rendering of an ordinary life.

---

[13] For discussions of other objections to structuralism, see, e.g., Vogel (2019) and Chalmers (2010, esp. 473) and Chalmers (2018: 654-658). For the concern that if reality is virtual, the simulation might not be stable enough to accommodate Chalmers' epistemological optimism, see Schwitzgebel (2017, 2019).

[14] For other routes to arguing for epistemological solipsism see, e.g., Gomes (2009, 2011, forthcoming) and Smith (2010). See Dogramaci (2020) and Schwitzgebel (2017: 276-77) for routes to skepticism derived from reflections on Boltzmann brains and Schwitzgebel (2017, 2019) for discussion of other routes to skepticism.

In your simulation, there appear to be nearly eight billion humans on the planet, sprawled across the habitable face of the earth. You seem to have friends, family, colleagues, co-nationals, neighbors. Perhaps you think you have a partner or children. Perhaps you think you have enemies or rivals. You presume that these others are conscious. But, these others aren't themselves brains-in-vats. Nor are they simulated beings with the same physical-functional features of yourself. Despite their indistinguishability from you on all third-personally observable measures, these others are the result of a highly sophisticated 'brute force' program, one which employs what is in effect a very large look-up table to offer seemingly sophisticated behavioral and linguistic responses across all relevant domains.

In a moment, I will suggest, following an influential argument from Ned Block, that the 'brute force' programs lack thought, and I will extend this argument to suggest that they lack anything worth calling a mental life, where a mental life requires, very roughly and minimally, well-integrated and sophisticated mental processes.[15] (In the literature, these 'brute force' programs are sometimes called *blockheads*, owing to Block's influential argument.) Lacking a mental life, these programs are also not conscious. These beings are rather what I call *faux-folk*. They are seemingly sentient entities which lack mental lives. They no more burn with desire than do toasters; no more have points of view than do books on a shelf; no more control their actions than does "Siri," the iPhone's talking assistant.[16] On this, contentious interpretation of the described scenario, you are a brain-in-a-vat who is also alone in the universe; I thus sometimes refer to this scenario, tongue firmly in cheek, as one in which you are a lonely brain-in-a-vat.

Importantly, to claim that the 'brute force' programs in the scenario are faux-folk is not to claim that they are *zombies*, where zombies are physical-functional duplicates of conscious creatures which themselves lack consciousness.[17] There are two main differences between faux-folk and zombies. First, while faux-folk lack anything worth calling a psychology, zombies merely lack phenomenal consciousness. That is, there is nothing it is like to be a zombie, but this is not to say that zombies lack mental states. Many philosophers suppose that zombies necessarily have very sophisticated and integrated mental states. It's just that these states aren't

---

[15] To my knowledge, Ada Lovelace (1842) is the earliest example of a philosopher arguing that some facet of a machine's architecture can prevent that machine from possessing 'intelligence.' Speaking of a hypothetical computing machine, she writes, "Thus, although it is not itself the being that reflects, it may yet be considered as the being which executes the conceptions of intelligence" (p. 15) and "[it] has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform. It can follow analysis; but it has no power of anticipating any analytical relations or truths." (p. 49). Turing (1950: 450-1, 454-60) responds to Lovelace's view at length.

[16] I am presuming, though it is not necessary for my argument, that 'brute force' programs are merely one of many potential kinds of faux-folk. Any relevantly architecturally deformed entity that can evince the right behavioral dispositions will be a member of the faux-folk; it needn't be a very large look-up table.

[17] Chalmers (1996: 93–171).

experienced for zombies in the way they are for conscious beings.[18] Second, while zombies are, by stipulation, physical-functional duplicates of conscious beings, faux-folk needn't be physical-functional duplicates of conscious beings. In fact, if we accept the widely held view that a physical-functional duplicate of a minded being must itself be minded, then faux-folk are necessarily not physical-functional duplicates of minded beings. Putting these claims together, we have reason to think that: no faux-folk are zombies and no zombies are faux-folk.
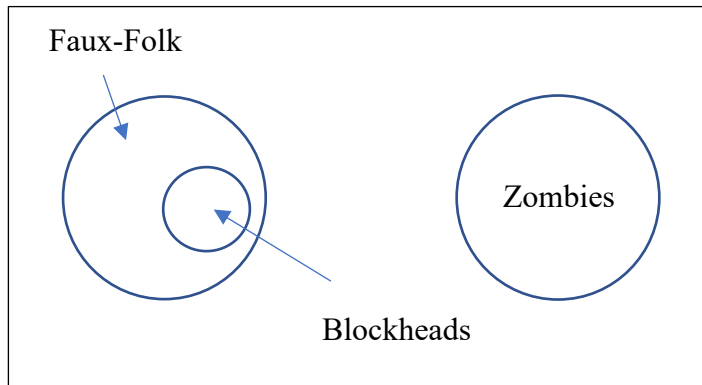
Figure 1. An Opinionated Guide to Blockheads, Faux-Folk, & Zombies

Here is why it matters that faux-folk and zombies form wholly distinct classes. First, even if zombies are metaphysically impossible, as some theorists maintain, this fact would not suggest that faux-folk are impossible.[19] Second, it will become important, in the second half of the paper, that faux-folk do not (merely) lack consciousness; they lack *anything worth calling a mental life*. For it is only because faux-folk lack psychologies—and not merely consciousness— that a world in which others are faux-folk is a world which threatens vast swathes of worldly knowledge.

The claim that 'brute force' programs lack minds can be used as part of an argument for epistemological solipsism, as follows. The reader can interpret the 'I' in the argument as herself. The argument is developed with respect to some particular friend who seems to be currently extremely happy, joyful even:

THE "BRUTE FORCE' PROGRAMS' ARGUMENT FOR EPISTEMOLOGICAL SOLIPSISM

(1) If: I know p and I know that p entails q, then: I know q.

(2) I know that: if my friend feels joy, then I am not in the "brute force' programs' scenario.

---

[18] For discussion, see Chalmers (1996, 2003a), Lyons (2009: ch. 4), and Smithies (2012).
[19] See, e.g., Levin (1985), McLaughlin (2005), Cf. Stoljar (2000).

(C1) If I know that my friend feels joy, then I know that I'm not in the "brute force' programs' scenario. (from 1, 2)

(3) I don't know that I'm not in the "brute force' programs' scenario.

(C2) I don't know that my friend feels joy (even if my friend in fact feels joy).

Like the classical argument for external world skepticism, the argument for solipsism draws on the claim that our evidence is consistent with a skeptical scenario and on the claim that knowledge is closed under a kind of entailment. The argument says that if you know that your friend feels joy, then you know the "brute force' programs' scenario doesn't obtain. But, you don't know that the "brute force' programs' scenario doesn't obtain. So, you don't know that your friend feels joy.

Since iterations of this argument can be constructed for many, perhaps most, of your beliefs concerning others' mental states, conjoined iterations of this argument can be used to establish that either solipsism or *near-solipsism* is true, where near-solipsism is the view that for at least many seemingly sentient creatures, you do not know whether they enjoy particular mental states or indeed, any mental states whatsoever. You do not know this even if the others around you do in fact have mental lives.

The structuralist who wishes to defuse this argument must reject (2), the claim that if you know that your friend feels joy, then you know that you are not in the "brute force' programs' scenario.

Recall how the structuralist countered a similar premise in the argument for skepticism. The structuralist rejected the claim that you know that: if you're on the train (and not in a vat), then you're not a brain-in-a-vat. For the structuralist, being on a train and being a brain-in-a-vat are mutually consistent, so you do not know that your being on a train precludes your being a brain-in-a-vat; you cannot know what is false. To make a similar move in response to the argument for solipsism, the structuralist must argue that it's false that you know that: if your friend feels joy, then you are not in the "brute force' programs' scenario. For, the structuralist must argue, there is no inconsistency between your friend feeling joy and your friend being realized by what is in effect a very large look-up table.

I will argue that there are independent grounds for thinking structuralism does not extend to concepts like FEELS JOY, *even if* structuralist does extend to concepts like IS A TRAIN. Thus, I suggest that the structuralist lacks cogent reasons to reject (2), even though it is coherent that (2) is false. Thus, I will suggest that there is a kind of asymmetry in the kinds of concepts for which conceptual structuralism is *prima facie* suitable; even if structuralism is prima facie attractive as a theory of concepts such as IS A TRAIN, it is not prima facie attractive as a theory of relevant psychological terms such as FEELS JOY. The argument I develop draws on and extends considerations familiar from 20th century philosophy of mind. I will first argue that

structuralism is ill-suited as a theory of relevant psychological concepts, such as FEELS JOY, before further suggesting that structurally similar reasons do not tell against structuralism as a theory of non-psychological concepts, such as IS A TRAIN.

To begin, here is an argument that structuralism should be rejected as a theory of at least some psychological concepts, such as FEELS JOY:

THE ARGUMENT AGAINST CONCEPTUAL STRUCTURALISM ABOUT PSYCHOLOGICAL CONCEPTS
  (1) If conceptual structuralism applies to an English speaker's concept FEELS JOY, then this concept picks out 'brute force' programs.
  (2) An English speaker's concept FEELS JOY does not pick out 'brute force' programs.
  (C1) Conceptual structuralism does not extend to the English speaker's concept FEELS JOY.

Notice that, throughout, the argument is couched in terms of the concept a speaker of *English* will have. It is not about the lonely brain-in-a-vat's concept FEELS JOY. Perhaps the lonely brain-in-a-vat's concept FEELS JOY, rendered in 'solipsese,' *does* pick out 'brute force' programs. But notice that, even if this is right, this fact would do nothing to assist the structuralist strategy. For, recall that structuralism, in its most powerful form, claims that *our* concepts, those had by an English speaker, have referents in the brain-in-a-vat world. So, for instance, when the structuralist claims that there are trains in the lonely-brain-in-a-vat world, she claims that there are referents of an English speaker's concept IS A TRAIN, not (merely) that there are 'trains,' where these are the referents of a BIV-speaker's concept IS A TRAIN. Thus, in assessing whether structuralism can provide a solution to solipsism, the only relevant question is whether an English speaker's concept FEELS JOY can pick out 'brute force' programs, not whether a 'solipsese' speaker's concept FEELS JOY can pick out such programs.

Consider the first premise. Here is why conceptual structuralism predicts that the concept FEELS JOY extends to 'brute force' programs. Recall that structuralism about some concept is the view that that concept's *meaning* is somehow determined by its role in some theory. But this characterization alone does not specify what the *referent* of that concept is. There are two different approaches the conceptual structuralist might take to the referent of FEELS JOY:

> THE 'ROLE' VIEW OF REFERENCE
> The referent of FEELS JOY is whatever third-personally observable set of dispositions plays the relevant *feels-joy*-role in the world of evaluation.

> THE 'REALIZER' VIEW OF REFERENCE
> The referent of FEELS JOY is the realizer of whatever third-personally observable set of dispositions plays the relevant *feels-joy*-role in whatever world is considered as actual.

On either of these views of reference, conceptual structuralism will predict that FEELS JOY, the concept in English, picks out 'brute force' programs.

First, consider the 'role' view of reference. By stipulation, the relevant 'brute force' programs play whatever third-personally observable set of dispositions is relevant to the *feels-joy*-role, so according to conceptual structuralism, these programs are included in the extension of FEELS JOY. For the structuralist, the relevant third-personally observable dispositions are those describable in terms of physics, observation, and terms to do with some 'centered' observer's subjective experience, in this case *your* subjective experience, so the experiences—or perhaps lack thereof—of the 'brute force' programs, are not relevant to the question of whether the relevant dispositions can play the relevant role.

Second, consider the 'realizer' view of reference, on which the referent of FEELS JOY is the realizer of whatever third-personally observable set of dispositions plays the relevant *feels-joy*-role in whatever world is considered as actual. If the "brute force' programs' scenario is the actual world, then 'brute force' programs realize the third-personally observable set of dispositions in the actual world.[20] Given that our world—that is, the actual world—is the "brute force' programs' scenario, then our concept FEELS JOY extends to 'brute force' programs.

So, conceptual structuralism predicts that the English concept FEELS JOY picks out 'brute force' programs. As I will now suggest, this result is problematic, since considerations familiar from 20th century philosophy of mind provide at least some reason to reject it. To do so, I draw on an argument from Ned Block concerning 'brute force' programs.

In the paper in which he introduces 'brute force' programs, Block argues that due to their unsophisticated structure, these programs lack 'intelligence,' or—as he puts it in one place—that these programs have 'the intelligence of a toaster.' By 'intelligence,' Block has in mind something like the ability to process or to understand information. On Block's view, this result is intuitive; how could something which is in effect a very large look-up table be a thinking being? Moreover, we can trust this intuition even if we do not have at the ready a positive proposal about what distinguishes 'brute force' programs from genuinely thinking beings. As Block puts the point:

> My machine [(a 'brute force' program)] lacks the kind of "richness" of information pro-cessing requisite for intelligence. Perhaps this richness has something to do with the application of abstract principles of problem solving, learning, etc. I wish I could say more about just what this sort of richness comes to. But I have chosen a much less ambitious task: to give a clear case of something that lacks that richness, but nonetheless behaves as if it were intelligent. (Block, 1981: 28)

---

[20] Or at least, in the "brute force' programs' scenario, such programs realize the vast preponderance of minds. They presumably do not realize your own.

Block's work and other work in its vein has convinced a surprisingly broad swathe of theorists to abandon the view that there is any third-personally observable disposition or set of dispositions an entity might have which would guarantee that that entity is a thinking being. At the same time, there is widespread disagreement about why 'brute force' programs are not thinking beings. What are they missing which distinguishes them from thinking beings? Some theorists take up Block's tentative suggestion that these 'brute force' programs are missing the ability to compress information or employ abstraction. Others maintain that these programs are missing causal sensitivity to semantic features. Still others suggest that these programs lack proper counterfactual dependence on their merely possible states.[21]

For our purposes, we needn't settle the fraught and longstanding dispute about why 'brute force' programs are not thinking beings. But, it *is* important for our purposes that the fact that 'brute force' programs are designed by some other intelligent being is *not* the reason they are not thinking beings. Nor is it do with the fact that these programs are, in some sense, wholly digital objects. It is rather to do with the fact that the *kind* of program employed seems too architecturally thin to sustain mental life. Put otherwise, one might accept that a wholly digital, programmed entity could be a thinking being whilst denying that the 'brute force' program in particular is such an entity.

Notice that Block's point that a 'brute force' program would not be a thinking being is consistent with a functionalist approach to the metaphysics of thinking. In general, functionalism is the view that some entity is a thinking being (or, extended to other psychological states, is a sentient being or is conscious being) just in case its mental states or processes satisfy certain functional roles, where a functional role is exhausted by the following three types of dispositions: world-to-mind dispositions, mind-to-mind dispositions, and mind-to-world dispositions.[22] Since the world-involving functional roles do not determine the mind-to-mind functional role, a 'brute force' program might lack the mind-to-mind functional role of a thinking being, even though, by stipulation, the world-involved aspects of this entity's dispositions are the same as that of a thinking being.

I would suggest that Block's point that 'brute force' programs are not thinking beings can be readily extended to argue that 'brute force' programs are not sentient, in the sense that they lack mental lives altogether. On this view, not only are these entities incapable of

---

[21] For recent criticisms of Block's argument, see Ben-Yami (2005) and McDermott (2014). For the view that semantic features are relevant for sentience, see, e.g., Haugeland (1981). For the view that information compression is relevant to sentience, see Dowe & Hajek (1997). For the view that a kind of counterfactual dependence on merely possible states is relevant to sentience, see O'Rourke's (2018) interpretation of Jackson (1993) and Braddon-Mitchell & Jackson (2007). Note that if 'brute force' programs are not sentient, then *interpretationism* is false, where this is the view that sentience is made up of interpretability (see, e.g., Dennett 1979, 1983). For the 'brute force' program is rationally interpretable and yet, intuitively not sentient. See Schwitzgebel (2015) for a discussion of the implications of broadly functionalist views for group minds.

[22] For a classic discussion, see Block (1978). For a recent, helpful overview, see Levin (2018).

understanding language or concepts, they also lack well-integrated and sophisticated mental states of the kind that are required to count as having a mind. Put otherwise, 'brute force' programs not only have "the intelligence of a toaster.' Intuitively, they have the *mental life* of a toaster as well, which is to say, none at all. They lack sufficiently integrated or sophisticated perceptual states, beliefs, wishes, urges, or feelings of the kind that are required in order to have a psychology. So, these programs are faux-folk, entities which appear to have a mental life but which have none. For instance, when they seem to feel joy, they do not feel joy, for they have no feelings whatsoever.

Notably, to suggest that 'brute force' programs are faux-folk is consistent with a functionalist approach to sentience. Since the world-involving aspects of the functional role needn't determine or fix the mind-to-mind disposition of the functional role, an entity that is 'externally' just like a sentient being needn't itself be sentient. Such an entity might have different mind-to-mind roles than a sentient creature.

So far, I have suggested that conceptual structuralism predicts that FEELS JOY, our concept in English, picks out 'brute force' programs and that, for reasons drawn from 20th century philosophy of mind, this is an unacceptable result. Moreover, this result holds even if we accept functionalism about the metaphysics of minds. I conclude that we should reject structuralism as a semantics of FEELS JOY and similar concepts. None of this rules out that structuralism might succeed as a theory of other concepts, such as HAS MASS or IS A PHILOSOPHER.

At this point, one might offer the following objection on behalf of the structuralist: Suppose that the referent of the concept FEELS JOY is fixed in the 'realizer' way, such that this concept picks out whatever realizes the *feels-joy*-role in whatever world is actual. If the 'brute force' programs world is the actual world, then everything, except for yourself, which plays the *feels-joy*-role is a 'brute force' program. Would this not undermine Block's suggestion that FEELS JOY, our concept in English, does not pick out 'brute force' programs? Consider that, in the case described, most employments of the concept are made with respect to 'brute force' programs.

I would suggest that the kind of intuition Block draws out is not contingent on the nature of the actual world. In particular, this intuition is not hostage to ordinary usage. Rather, if this intuition has any force at all, it is because it tracks a *substantive constraint on what can count as having a mind*, a constraint that cannot be overridden (merely) by surprising discoveries about the hidden nature of those entities to whom we have tended to apply the concept FEELS JOY or other psychological concepts. It is a prediction of this suggestion that, were we to discover that we were in fact in the 'brute force' programs scenario, we would *not* happily accept the result that the concept FEELS JOY extends to 'brute force' programs. We would rather conclude— surely with something like shock—that those around us don't feel joy because they can't. I take it that these predictions are borne out. (In contrast: I presume that were we to discover that ours is a world in which water-y stuff is realized by XYZ, not H_2O, we would revise our view about what our concept IS WATER picks out. We would *not* conclude that water does not exist.

This difference illustrates a deep difference between our concepts of natural kinds and our psychological concepts.)

If, as I have argued, conceptual structuralism does not extend to FEELS JOY, the structuralist solution to external world skepticism does not vindicate knowledge of other minds. This is so even if structuralism does vindicate knowledge of inanimate entities, such as trains. For the structuralist might plausibly maintain that our concept in English of (say) IS A TRAIN picks out whatever plays the *is-a-train*-role in some scenario—or else, whatever realizes that role in whatever world is taken to be actual—with the result that, in skeptical scenarios, there are trains. She *might* maintain this even if this approach is implausible when it comes to FEELS JOY.

At this point, the advocate of structuralism might reply as follows: Surely, *if* structuralism extends to concepts such as IS A TRAIN, it also extends to concepts like FEELS JOY. On this view, concepts like IS A TRAIN and psychological concepts like FEELS JOY demand symmetrical treatment when it comes to theorizing their meanings. For, why would structuralism hold of one but not the other? Extending this thought, the structuralist might argue that because there is reason to think structuralism extends to concepts like IS A TRAIN, there is also reason to think that structuralism extends to concepts like FEELS JOY.[23]

The preceding line of thought, on which there is a kind of symmetry between psychological concepts and non-psychological ones, such that structuralism either holds of both kinds of concepts or neither kind, is an extremely intuitive one. Nevertheless, I submit that there is some reason to reject this symmetry thesis. The reason is this: While reflection on 'brute force' programs provides some reason to reject structuralism as a theory of psychological concepts, analogous thought experiments to do *not* provide a reason to reject structuralism as a theory of non-psychological concepts. This result alone should at least cause us to *doubt* the symmetry thesis, whether or not it constitutes a decisive reason against it.

Recall that the ''brute force' programs' thought experiment elicited the intuition that an entity realized by an extremely large, 'brute force' program—what is in effect a very large look-up table—needn't be sentient, even if it is behaviorally indistinguishable from a sentient entity. This intuition in turn formed part of the argument that structuralism does not extend to psychological concepts, such as FEELS JOY. For, if it did, then FEELS JOY and other concepts would extend to 'brute force' programs, and this is an intuitively unpalatable result.

But consider now whether a relevantly similar line of thought might be used to cast doubt on the view that structuralism holds of a theory of non-psychological concepts, such as IS A TRAIN. Consider an entity that is, in terms of third-personally observable behavior, identical to those entities we take to be trains. This entity transports cargo and passengers, runs on tracks, and the like. Functionally speaking, it is *just* like a train. Were we to learn that this entity is somehow realized by a highly sophisticated look-up table or 'brute force' program, would we immediately reject the view that this entity is a train? To my mind, it is at least not clear what we

---

[23] I thank an anonymous referee for developing and pressing me on this point.

should say of such an entity. Intuitions alone do not seem to settle the issue. At most, learning that this functional duplicate of a train is realized by a 'brute force' program might *raise the question* of whether this entity is a train, but intuition alone does not seem to settle that question. If I'm right about this, then the intuitive response to the 'train' variant of the "brute force' programs' thought experiment is rather different than the standard response to the "brute force' programs' thought experiment, since many take the latter to provide a probative reason to reject wholly behavioral analyses of mental states.

If the preceding considerations are right, then we have some reason to doubt the symmetry claim that structuralism holds of both psychological and non-psychological concepts, if it holds at all. For, intuitions about whether such concepts might be realized by 'brute force' programs differ in the case of psychological and non-psychological concepts. The structuralist who wishes to endorse the symmetry thesis in the face of these considerations should offer some debunking explanation of the intuitions at hand or else explain why these intuitions are irrelevant to the matter.

If, as I have suggested, conceptual structuralism is ill-suited for concepts such as FEELS JOY, what would follow? I will next suggest that what follows is that vast swathes of our worldly knowledge are compromised, with the result that structuralism cannot, on its own, provide the kind of substantive solution to external world skepticism that we might have hoped for. Again, this suggestion does not run counter to Chalmers' own, comparatively modest ambitions for the view, but it does suggest an outside limitation to the view and one that derives from a rather surprising source—its failure to solve the problem of other minds.

3 Solipsism Constrains the Ambitions of the Structuralist Solution

Proposed solutions to radical skepticism can vary in how much worldly knowledge they vindicate. We can say that some proposed solution *m* is a more substantive solution than some competitor solution *n* just in case *m* vindicates more worldly knowledge than does *n*. All else being equal, more substantive solutions are preferable to less substantive solutions. I will suggest that the structuralist solution is less substantive than what we would hope for from a solution to skepticism. So, even if we grant that structuralism offers some shelter from skepticism, structuralism falls short of the kind of substantive solution which, once offered, would quell our skeptical anxieties.[24]

How troubling is it for the structuralist solution if it cannot vindicate knowledge of mental states, such as knowledge about a friend's joy? I will argue that surprisingly vast and varied swathes of empirical knowledge depend on knowledge of others' particular mental states, including at least some knowledge of political affairs, cultural traditions, and religious practices. This supports the paper's second thesis:

---

[24] See Schwitzgebel (2017: 286-287) and Helton (2021: 241) for the suggestion that solipsism is of special concern for an independent reason, namely that it threatens our most valuable beliefs.

SOLIPSISM THREATENS THE STRUCTURALIST SOLUTION

If epistemological solipsism is true, then the structuralist solution to skepticism is far less substantive than we might have hoped for from a solution to skepticism.

Suppose that epistemological solipsism holds; you do not know that other minds exist because, for all you know, the seemingly sentient others in your environment are *faux-folk.* They are seemingly in possession of psychologies but lack mental lives. I will suggest reasons to think that in this situation, at least some beliefs which attribute specific mental states to others and some beliefs which attribute actions to others will turn out to be false.

First, consider beliefs which attribute specific mental states to others. In the faux-folk scenario, all or virtually all such beliefs are false. For instance, virtually all beliefs are false which attribute opinions, urges, wishes, imaginings, hopes, intentions, auditory experiences, itches, twinges, or the like to others.

For instance, suppose you see your colleague in the lounge of your building. He greets you warmly and asks you about your weekend. You infer he's interested in hearing about your life. This encounter will likely result in your having belief such as the following (among many others):

(1) My colleague *is happy* to see me.
(2) My colleague *wants* to know how my weekend was.

If your colleague is not sentient but is rather a member of the faux-folk, then, given the points argued for at length in the previous section, none of these beliefs are true. Your colleague isn't glad to see you, because he isn't capable of gladness. He doesn't want to know about your weekend, because he has no desires whatsoever. These points about (1)-(2) extend to any of your beliefs which attribute particular mental states to members of the faux-folk.

As it turns out, many of our beliefs about the mental states of others are beliefs about those not known to us personally. Some such beliefs are beliefs about political affairs, sociological facts, historical events, or aesthetic movements. Consider, for instance, the following:

(3) Many people go to bars because they *enjoy* being around other people.
(4) Most Sandinistas *aimed* to end imperialism.
(5) Some Hindus *recognize* many deities.
(6) The originators of free jazz *wanted* to challenge the conventions of bebop.
(7) Many people *love* dancing.[25]

---

[25] The argument does not draw on *generic* beliefs, where generics are claims of the form Fs are a, since generics might turn out to be true even where very few or perhaps none of the relevant Fs are a. See, e.g., Leslie (2007, 2008, 2017).

Next, consider beliefs which attribute actions to others. In the solipsistic scenario, at least many of these beliefs will be false, since at least many actions are partly made up of relevant proximal intentions, such that someone cannot count as performing that action if they lack the relevant proximal intention.[26]

At least many actions are partly constituted by relevant proximal intentions. Consider, for instance, creating. Creating requires more than being causally connected to some object. The group of crabs whose wanderings in the sand happen, though some fluke, to leave in their wake something that resembles a sand castle did not *create* a sand castle, even though they causally contributed to the impressions in sand. I submit that what the crabs are missing is a relevant *proximal intention*, which is roughly an intention which guides movements 'in real time.' For instance, consider the following:

(8) Cezanne *experimented* with form in his Mont Sainte-Victoire paintings.
(9) That child *is building* a sand castle.
(10) The Normans *created* a feudal system in England.
(11) The judges *are* carefully *deliberating* about the case.

To claim that in the solipsistic scenario, (8)-(11) will be false is not to claim that all beliefs in the vicinity of (8)-(11) will be false. For instance, suppose you see a member of the faux-folk seemingly building a castle in the sand. Since building involves a proximal intention, they aren't building, but they *are* moving their hands, and their hands are in turn moving sand. For, moving oneself or another entity is not an activity which requires having a mind. Drops of water move when they plummet to the ground as part of a rainstorm. A snowball rolling down a mountain, gathering up snow as it goes, moves other snow. So, the claim that in the solipsistic scenario, (9) is false is *not* to claim that all beliefs in the vicinity are false.

I suspect that many other kinds of beliefs would turn out to be false in a solipsistic scenario, but for present purposes, it suffices to have argued that many beliefs concerning others' particular mental states and intention-involving actions will be false. For, notice that even just these two kinds of beliefs include beliefs which cross intuitive domains and beliefs about a wide swathe of beliefs about political affairs, aesthetic movements, and social and religious practices. Intuitively, we would not rest content with any solution to skepticism which did not vindicate these swathes of knowledge; rather, presented with such a solution, we would presumably eagerly continue our search for one which would (also) vindicate these varied and numerous beliefs.[27]

---

[26] See Helton (2018: 249-51) for a fuller development of this argument.
[27] Indeed, there is some reason to think that if knowledge of other minds were lost, at least some knowledge of artifacts, natural kinds, and even moral knowledge would also be false in a solipsistic scenario, insofar as there is some reason to think that artifacts, scientific kinds, and moral facts require other

4 Some Broader Morals

I have argued that structuralism, whatever its other virtues, cannot provide a solution to skepticism of the kind we might have hoped for. This is so even if we combine structuralism with a functionalist metaphysics of the mind. This is because, on any plausible version of functionalism, some aspect of the functional role is irreducibly mental and this role element in turn is not reducible to the kinds of *third-personally observable* dispositions that are crucial to the structuralist vindication of worldly knowledge. In virtue of failing to vindicate knowledge of other minds, structuralism fails to vindicate a very broad swathe of worldly knowledge, including at least some knowledge of political affairs, religious practices, cultural practices, and aesthetic events.

This result suggests some morals that go beyond the evaluation of structuralism as an approach to skepticism. For, the example of structuralism raises as a distinct possibility that the all things considered best solution to external world skepticism might require two different theoretical approaches—one to vindicate worldly knowledge of the kind that does *not* depend on other minds and one to vindicate knowledge that *does* depend on other minds. Suppose, for the sake of argument, that structuralism is the best way of vindicating worldly knowledge of the kind that is not dependent on other minds. This would mean that, in order for us to obtain a 'total' solution to external world skepticism, we must employ a different strategy to vindicate knowledge of the kind that depends on other minds. If this is right, there is a sense in which the problem of external world skepticism is perhaps best cast as two different problems, insofar as its solution might turn out to require two substantively different strategies.[28]

REFERENCES

Ásta (published under Sveinsdóttir, Ásta). (2008). Essentiality conferred, *Philosophical Studies*, 140(1): 135–148. doi:10.1007/s11098-008-9230-4

Ásta (published under Sveinsdóttir, Ásta). (2010). Siding with Euthyphro: Response-Dependence and Conferred Properties, *European Journal of Philosophy*, 18(1): 108–125. doi:10.1111/j.1468-0378.2008.00303.x

Ásta (published under Sveinsdóttir, Ásta). (2013). The Social Construction of Human Kinds. *Hypatia*, *8*(4), 716-732.

Baker, L. R. (2007). *The Metaphysics of Everyday Life: An Essay in Practical Realism*, Cambridge: Cambridge University Press.

Ben-Yami, H. (2005). Behaviorism and psychologism: Why block's argument against behaviorism is unsound. *Philosophical Psychology*, *18*(2), 179-186.

Block, N. (1978). Troubles with Functionalism, in C. W. Savage, ed., *Perception and Cognition: Issues in the Foundations of Psychology*, Univ. Minn. Press, 261-326.

Block, N. (1981). Psychologism and behaviorism. *The Philosophical Review*, *90*(1), 5-43.

Bouwsma, O. K. (1949, April). Descartes' Evil Genius (An Exercise in Language). In *Proceedings of the Tenth International Congress of Philosophy*, 2, 1122-1124.

Boyd, R. (1991). Realism, anti-foundationalism and the enthusiasm for natural kinds. Philosophical Studies, 61(1), 127–148.

Boyd, R. (1999). Kinds, complexity and multiple realization. Philosophical Studies, 95(1), 67–98.

Braddon-Mitchell, D. and Jackson, F. (2007). Philosophy of Mind and Cognition: An Introduction, 2nd edn.., Blackwell Publishing.

Brueckner, A. (1986). Brains in a Vat. *The Journal of Philosophy*, *83*(3), 148-167.

Brueckner, A. (1992). Semantic Answers to Skepticism. *Pacific Philosophical Quarterly*, 73(3): 200–219.

Brueckner, A. (2003). Trees, Computer Program Features, and Skeptical Hypotheses, in Robert Stern (ed.), *Transcendental Arguments: Problems and Prospects*, Oxford: Clarendon Press.

Chalmers, D. J. (1996). *The conscious mind: In search of a fundamental theory.* Oxford Paperbacks.

Chalmers, D. J.. (2003a). The Content and Epistemology of Phenomenal Belief. In *Consciousness: New Philosophical Perspectives*, edited by Q. Smith and A. Jokic. New York: Oxford University Press.

Chalmers, D. J. (2003b). The Matrix as metaphysics. *Science Fiction and Philosophy From Time Travel to Superintelligence*, *36*.

Chalmers, D. J. (2018). Structuralism as a response to skepticism. *The Journal of Philosophy*. Volume 115, Issue 12, pp. 625-660.

Chalmers, D. J. (2022). *Reality+: Virtual worlds and the problems of philosophy*. Allen Lane: London.

Christensen, D. (1993). Skeptical Problems, Semantical Solutions. Philosophy and Phenomenological Research 53 (2): 301–21.

Davidson, D. (1986). A Coherence Theory of Truth and Knowledge, in E. Lepore, ed., *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson* (London: Blackwell), pp. 307–19.

Dennett, D. (1979). True believers: The Intentional Strategy and Why It Works. In A. Heath (ed.), *Scientific Explanation* (Oxford: Oxford University Press), 53-75.

Dennett, D. C. (1983). Intentional systems in cognitive ethology: The "Panglossian paradigm" defended. *Behavioral and Brain Sciences*, *6*(3), 343-355.

Dipert, R., (1993). *Artifacts, Art Works, and Agency*, Philadelphia: Temple University Press.

Dogramaci, S. (2020). Does my total evidence support that I'm a Boltzmann Brain? *Philosophical Studies*, *177*(12), 3717-3723.

Dowe, D. L. & Hajek, A. R. (1997). A computational extension to the Turing Test. In *Proceedings of the 4th conference of the Australasian cognitive science society, University of Newcastle, NSW, Australia* (Vol. 1, No. 3).

Elder, C.L., (2004). *Real Natures and Familiar Objects*, Cambridge, MA: MIT Press.

Elder, C. L. (2007). On the place of artifacts in ontology, in Margolis, E. and Laurence, S *Creations of the Mind: Theories of Artifacts and Their Representations*, Oxford, Oxford University Press

Ellis, B. D. (2001). Scientific essentialism. Cambridge, MA: Cambridge Univ Press.

Epstein (2008). The Realpolitik of Reference, *Pacific Philosophical Quarterly*, 89(1): 1–20. doi:10.1111/j.1468-0114.2008.00307.x

Epstein (2014). How Many Kinds of Glue Hold the Social World Together?, in Mattia Gallotti and John Michael (eds.), *Perspectives on Social Ontology and Social Cognition*, Dordrecht: Springer Netherlands, pp. 41–55. doi:10.1007/978-94-017-9147-2_4

Epstein (2015). *The Ant Trap: Rebuilding the Foundations of the Social Sciences*, New York: Oxford University Press. doi:10.1093/acprof:oso/9780199381104.001.0001

Epstein, B. (2018). Social Ontology. *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2018/entries/social-ontology/>.

Evnine, S. J. (2016). *Making Objects and Events: A Hylomorphic Theory of Artifacts, Actions, and Organisms*. Oxford University Press.

Farrell, F. (1986) Putnam and the Vat-People. Philosophia 16 (2): pp. 147–60.

Forbes, G. (1995). Realism and skepticism: brains in a vat revisited. *The Journal of philosophy*, *92*(4), 205-222

Franklin-Hall, L. R. (2015). Natural kinds as categorical bottlenecks. *Philosophical Studies*, *172*(4), 925-948.

Gauthier, D. (1986). *Morals By Agreement*, Oxford: Oxford University Press.

Gauthier, D. (1990). *Moral Dealing: Contract, Ethics, and Reason*, Ithaca: Cornell University Press.

Glymour, C. (1982) Conceptual Scheming or Confessions of a Metaphysical Realist. *Synthese* 51 (2): pp. 169–80.

Gomes, A., (2009). Other Minds and Perceived Identity, *Dialectica*, 63: 219–230.

Gomes, A. (2011). Is There a Problem of Other Minds? *Proceedings of the Aristotelian Society*, 111: 353–373.

Gomes, A. (forthcoming). Scepticism about Other Minds. In D. Machuca & B. Reed (eds.) *Skepticism: From Antiquity to Present*.

Goodman, N. (1978). Ways of worldmaking. Indianapolis: Hackett Publishing Company.

Hacking, I. (2007). Natural kinds: Rosy dawn, scholastic twilight. *Royal Institute of Philosophy Supplement*, 61(1), 203–239.

Haugeland, J. (1981). *Semantic engines: An introduction to mind design* (pp. 1-34).

Helton, G. (2018). Visually perceiving the intentions of others. *The Philosophical Quarterly*, *68*(271), 243-264.

Helton, G. (2021). Epistemological solipsism as a route to external world skepticism. *Philosophical Perspectives*, *35*(1).

Helton, G. (forthcoming). The Simulation Hypothesis, Structure, and Social Reality. *Oxford Studies in Philosophy of Mind*.

Hilpinen, R. (1992). Artifacts and Works of Art. *Theoria*, 58(1): 58–82. doi:10.1111/j.1755-2567.1992.tb01155.x

Hilpinen, R. (2011). Artifact. *The Stanford Encyclopedia of Philosophy* (Winter 2011 Edition), Edward N. Zalta (ed.), URL = < target="other">https://plato.stanford.edu/archives/win2011/entries/artifact/>.

Jackson, F. (1993). Block's challenge, in J. Bacon, K. Campbell and L. Reinhardt (eds), *Ontology, Causality, and Mind: Essays in Honour of D.M. Armstrong*, Cambridge University Press, pp. 235-246.

Khalidi, M. A. (2013). Natural categories and human kinds. Cambridge: Cambridge University Press.

Khalidi, M. A. (2015). Three kinds of social kinds. Philosophy and Phenomenological Research, 90, 96–112.

Leslie, S. (2007). Generics and the structure of the mind. *Philosophical Perspectives*, 21(1), 375–403.

Leslie, S. (2008). Generics: cognition and acquisition. *Philosophical Review*, 117(1), 1–47.

Leslie, S. (2017). The original sin of cognition: Fear, prejudice, and generalization. *The Journal of Philosophy*, 114(8), 393-421.

Levin, J., (1985). Functionalism and the Argument from Conceivability, *Canadian Journal of Philosophy*, 11 (Supplement): 85–104.

Levin, J. (2018). Functionalism, *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2018/entries/functionalism/>.

Lovelace, A. (1842). Translator's notes to an article on Babbage's Analytical Engine. Scientific Memoirs (ed. by R. Taylor), vol. 3. 691-731.

Lyons, J. C. (2009). *Perception and basic beliefs: Zombies, modules and the problem of the external world.* OUP USA.

McDermott, D. (2014), On the Claim that a Look-Up Table Program Could Pass the Turing Test, *Minds and Machines*, 24: 143–88.

McLaughlin, B. P., (2005) *A priori* versus *a osteriori* physicalism, in *Philosophy-Science-Scientific Philosophy, Main Lectures and Colloquia of GAP 5, Fifth International Congress of the Society for Analytical Philosophy* (C. Nimtz & A. Beckermann eds.), Mentis.

Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories: New Foundations for Realism*, Cambridge, MA: MIT Press.

Millikan, R. G. (1999). Historical Kinds and the 'Special Sciences', *Philosophical Studies*, 95(1–2): 45–65. doi:10.1023/A:1004532016219

Millikan, R. G. (2005). *Language: A Biological Model*, Oxford: Oxford University Press. doi:10.1093/0199284768.001.0001

O'Rourke, J. (2018). *Property dualism as a solution to the mind-body problem.* Unpublished doctoral dissertation. Princeton University.

Preston, B. (2013). *A philosophy of material culture: Action, function, and mind.* Routledge.

Preston, B. (2018). Artifact, *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2018/entries/artifact/>.

Putnam, H. (1981). Brains in a vat. In Reason, Truth and History (pp. 1-21). Cambridge: Cambridge University Press. Doi: 10.1017/CBO9780511625398.003.

Rorty, R. (2011). Davidson versus Descartes, in J. Malpas, ed., Dialogues with Davidson: Acting, Interpreting, Understanding (Cambridge, MA: MIT Press), pp. 3–6.

Schwitzgebel, E. (2015). If Materialism is True, the United States is Probably Conscious, *Philosophical Studies*, 172: 1697–1721.

Schwitzgebel, E. (2017). 1% Skepticism. *Noûs*, *51*(2), 271-290.

Schwitzgebel, E. (2019). Kant Meets Cyberpunk. *Disputatio: International Journal of Philosophy*, *11*(55).

Sider, T. (2012). Writing the book of the world. Oxford: Oxford Univ Press.

Smith, J. (2010). The conceptual problem of other bodies. *Proceedings of the Aristotelian Society*. 110: 201–217.

Smith, P. (1984). Could we be Brains in a Vat? *Canadian Journal of Philosophy* 14 (1): pp. 115–23.

Smithies, D. (2012). The mental lives of zombies. *Philosophical Perspectives*, 26, 343-372.

Stoljar, D. (2000). Physicalism and the Necessary a Posteriori. *Journal of Philosophy*, 97: 33–54.

Thomasson, A. (2003). Realism and Human Kinds, *Philosophy and Phenomenological Research*, 67(3): 580–609. doi:10.1111/j.1933-1592.2003.tb00309.x

Thomasson, A. (2007), Artifacts and Human Concepts, in *Creations of the Mind: Essays on Artifacts and Their Representation*, edited by Eric Margolis and Stephen Laurence, Oxford: Oxford University Press, pp. 52–73.

Thomasson, A. (2014). Public Artifacts, Intentions, and Norms. In Artefact Kinds. Eds.

Thorpe, J. R. (2018). Closure scepticism and the vat argument. *Mind*, *127*(507), 667-690.

Turing, A. (1950). Computing machinery and intelligence. *Mind*, *59*(236), 433-460.

Vogel, J. (2019). Space, Structuralism, and Skepticism. *Oxford Studies in Epistemology Volume 6*, 190.

Wright, C. (1992). On Putnam's Proof that We are not Brains-in-a-Vat. Proceedings of the Aristotelian Society 92 (1): pp. 67–94.