

## **Metaphysical Considerations for External World Skepticism**

**By Micah Sparacio**

External world skepticism (EWS) has intuitive force. If I don't know that I'm not a brain in a vat (BIV), then how can I possibly know that I have two hands? By the principle of closure, it seems that not knowing that I'm not a BIV implies not knowing that I have two hands. Epistemological solutions to the problem have typically relied on one of three strategies: reject the skeptical scenario (e.g. I know that I'm not a BIV), deny the principle of closure, or contextualize knowledge. Each of these solutions seems to deny something intuitive regarding the force of EWS. Drawing on Chalmers recent work<sup>1</sup>, I will argue that the problem of EWS is intimately tied to how we conceive of the *proper* relationship between our ontology and our concepts, and that a metaphysical analysis can provide clarity on the issue. However, I reject Chalmers' thesis that metaphysics can *solve* global skepticism about the external world. Rather, I make a weaker claim; namely that the bite of EWS is significantly weakened if we consider our ordinary epistemic response to the micro-macro relationships of the physical world. I conclude with a dilemma: either contemporary physical theories should leave us skeptical about our knowledge of the macro physical world or external world skepticism is philosophically benign.

### **External World Skepticism (EWS)**

It is important to note at the outset that EWS is not a challenge to the truth of propositions about the world. For all we know, such propositions *could be* true. Rather, the skeptic is concerned with the **justification** of knowledge claims about such propositions. Thus, skepticism makes no metaphysical claim about the nature of reality; it is an epistemic critique of the justification for beliefs (i.e. propositions referring to the external world).

The intuition behind external world skepticism initially seems right. If I don't know (1) that I'm not a BIV in some distant universe, then how can I possibly know that (2) I have two hands? By implication, it appears that my not knowing the first premise *entails* my not knowing the second. Similarly, knowing the second premise should entail the first.

---

<sup>1</sup> A talk given at the *Towards a Science of Consciousness* 2004 Conference and currently an online paper, *The Matrix as Metaphysics*. <http://jamaica.u.arizona.edu/~chalmers/papers/matrix.html>

The strength of EWS can be illustrated by a consideration of the following three seemingly incompatible propositions, which together make up what has been called the skeptical paradox:

1. I know that I am in Princeton, New Jersey.
2. I do not know that I am not a brain in a vat.
3. If I do not know that I am not a brain in a vat, then I do not know that I am in Princeton, New Jersey. (Principle of Epistemic Closure)

The tension in the skeptical paradox takes place between premise one (i.e. common knowledge) and the joint force of premise two and three (i.e. indiscernibility and closure). It seems that unless I know that I am not a BIV (and I can't because such a skeptical scenario is empirically indiscernible from the actual world), I cannot know that I am in Princeton, New Jersey. Premise two and three, together, provide a formidable challenge to premise one, and give the skeptic his ammo.

So how is the non-skeptic to respond to this uncomfortable situation? Most philosophers would like to preserve the first premise. Obvious solutions for doing so are to deny either of the second two premises. To deny the second premise, we might maintain the principle of closure, but conclude that the closure principle, combined with premise one (which is more probable than premise two), can yield a rejection of premise two. This response is typically associated with Moore. Another response, exemplified by Nozick's tracking solution to justification, has been to accept premise two but to deny premise three (the principle of epistemic closure). But such a strategy seems plagued by the high cost of rejecting the intuition that our knowledge is closed under implication. Yet another method for dealing with skepticism, currently popular among philosophers, is to contextualize knowledge such that the standards of ascribing knowledge change according to epistemic context. This view has gained favor because it seems to accomplish the impossible (i.e. grant skepticism its due, maintain the principle of closure and uphold the legitimacy of everyday knowledge).

### **The Principle of Epistemic Closure**

The principle of epistemic closure is critical to the power of EWS. The principle gets us to acknowledge both that, "If I know that I am a BIV, then I do not know that I have two hands" and "If I do not know that I am not a BIV, then I can not know that I have two hands." In both cases the skeptic urges us that given the first premise we are able, through

implication, to know the second. The principle of closure formalizes the intuition that we can know extra things that follow from what we already know. Generally stated, the principle goes as follows: If person  $S$  knows  $p$ , and  $S$  knows that  $p$  entails  $q$ , then  $S$  knows  $q$ <sup>2</sup>.

So why is this principle central to EWS? Because it is the cognitive mechanism by which we feel the force of the skeptical intuition. We think, initially, that given a bizarre skeptical scenario, our normal knowledge would be undermined. The bizarre skeptical scenario (and our inability to rule out such a scenario) *implies* that our macro-world concepts are either false (given a skeptical scenario) or unjustified (given indiscernibility) or both. At least this is the skeptical contention.

### **Denying Epistemic Closure**

As mentioned earlier, one method for defeating EWS is to deny the principle of epistemic closure. This solution has been tried in a variety of ways, most notably by Dretske and Nozick. Dretske contends that my not knowing that I'm not a BIV does not matter because it is not a relevant alternative to consider in my belief that I have two hands. For Dretske, justification comes from ruling out *relevant* alternatives to your belief. Using Dretske's own example, I can know that an animal at a zoo is a zebra without considering the alternative possibility that the animal is actually a painted mule. In that environment, the "painted mule" hypotheses is not a relevant alternative. Dretske's theory of justification leads him to deny the principle of closure, because an irrelevant alternative need not impact belief.

Robert Nozick's subjunctive conditional account (SCA) of knowledge<sup>3</sup>, and especially its tracking condition, lead to a unique strategy for dealing with EWS. Under Nozick's account, the subject's belief in a proposition tracks the truth of the proposition, such that if the proposition were not true, the subject would not believe it. An SCA analysis is performed by considering other proximate possible worlds where the proposition is false. In such worlds, if sufficiently close to our world, the subject will not believe the proposition. So, in a sufficiently close world, if the subject does not have

---

<sup>2</sup> There are a variety of formulations of the principle, dealing with various criticisms, but this one will do for my purposes. See: <http://plato.stanford.edu/entries/closure-epistemic/> for a fuller discussion.

<sup>3</sup> "Knowledge and Skepticism." *Epistemology: an anthology*. Ed. Ernest Sosa, and Jaegwon Kim. Malden: Blackwell Publishers, 2000. 79-101.

hands, she will not believe that she has hands and if she does have hands she will believe that she has hands. It is this tracking condition that turns a true belief into knowledge.

Nozick's SCA account allows him to accept premise two: I do not know that I am not a brain in a vat. He can accept this premise because the tracking condition is not met. If I was a brain in a vat (in this and similar worlds), I still wouldn't believe it. Thus, I can't know that I'm not a brain in a vat because my belief does not track the truth of the proposition.

Having conceded the second premise to the skeptic, Nozick is left to deal with skepticism by denying the principle of closure. Under this view I can both (a) know that I am in Princeton, New Jersey and (b) not know that I'm not a brain in a vat in some distant universe. This seems counterintuitive. Being in Princeton and being a brain in a vat in some distant universe seem like mutually exclusive options. Nozick's solution opts for one difficult position (a denial of epistemic closure) over another (the skeptic's denial of everyday knowledge). But is the rejection of epistemic closure really worth the sacrifice?

### **Moore's Common Sense Solution**

The common sense approach to skepticism, employed famously by G.E. Moore, can accept the principle of epistemic closure, and can even use the principle to his advantage against the skeptic<sup>4</sup>. Moore tells us that any rational person has more confidence in some propositions (e.g. that he has two hands) than he has in other propositions (e.g. the premises in the sceptical argument)<sup>5</sup>. If this is the case, then the rational choice is to choose the most likely of competing propositions.

Rather than using what we don't know as the antecedent to the closure principle (e.g. that I am not a brain in a vat), Moore's approach suggests that we use what we do in fact know (e.g. that we have two hands). We can end up with something like the following:

1. I have two hands (derived from a previous argument in which he raises his hands).<sup>6</sup>
2. If I have two hands, then I know that I am not a brain in a vat.
3. I know that I am not a brain in a vat.

For Moore, the important thing is to ask whether I am more certain that (a) I have two hands or (b) that I don't know that I'm not a BIV. If I am more certain that I have two

---

<sup>4</sup> "Certainty." *Epistemology: an anthology*. Ed. Ernest Sosa, and Jaegwon Kim. Malden: Blackwell Publishers, 2000. 29-33.

<sup>5</sup> "Four Forms of Skepticism." In Sosa and Kim. 27-29.

<sup>6</sup> "Proof of an External World." In Sosa and Kim . 24-27.

hands, then the principle of closure tells me that I am not a BIV. If we begin our philosophy with common sense principles, as Moore does, then we are forced to conclude that we do, indeed, have hands and are not BIVs. At the very least, it seems irrational to deny knowledge that I have two hands on the less certain grounds that I don't know that I'm not a brain in a vat. I'll have more to say about Moore's solution shortly.

### **Contextualism**

Many philosophers have embraced the contextualist solution to EWS offered by DeRose<sup>7</sup>, Lewis<sup>8</sup> and others. Contextualism, as seen by its advocates, has the virtue of maneuvering through EWS without having to throw away some of the intuitive principles or ideas involved in the skeptical paradox. Contextualists maintain a qualified version of epistemic closure where knowledge, in a given epistemic environment under particular epistemic standards, is closed under implication. They also acknowledge the force of skepticism and defer to EWS under skeptical epistemic contexts. And, despite the odds against them, the contextualists still manage to save everyday knowledge under normal epistemic standards. Notice the strategy: for each of the seemingly incompatible propositions involved in the skeptical paradox, the contextualist offers a qualification regarding the epistemic context<sup>9</sup>.

For the contextualist, attributions of knowledge are context-sensitive: the meaning of the word changes in different conversational situations. Thus, the word "know" acts very much like an indexical, where its meaning changes in direct relation to changing epistemic standards. This allows the contextualist to do his trick. The epistemic standards under normal conversational contexts are much lower than the standards employed in skeptical contexts. Thus, attribution of the word "know" in low-standard contexts is perfectly legitimate, whereas it ceases to be appropriate in skeptical contexts. So, our ordinary knowledge is saved, under contextualism, as long as we are in the appropriate epistemic context with low standards of knowledge. But this leaves us in the unfortunate position of having to accept the following proposition: "I know that I have two hands in everyday contexts while I do not know that I have hands in BIV contexts." The

---

<sup>7</sup> "Solving the Skeptical Problem." *Epistemology: an anthology*. Ed. Ernest Sosa, and Jaegwon Kim. Malden: Blackwell Publishers, 2000. 482-502.

<sup>8</sup> "Elusive Knowledge." In Sosa and Kim, 503-516.

<sup>9</sup> For a more detailed discussion of contextualism, see: <http://www.iep.utm.edu/c/contextu.htm>

contextualist thinks that this is just fine; that knowledge indexes in relation to the epistemic context and that the meaning of the word “know” does not transfer across contexts. We therefore do not run into a contradiction between uses of the word in different scenarios.

### **The Skeptical Hunch Lingers**

In each of the considered candidate solutions, the force of EWS remains. Moore’s solution begs the question, and neglects the indiscernibility of various metaphysical scenarios. Moore’s premises bias one metaphysical scenario over another and do not help to clarify the skeptical intuition that we have no way of discerning between various experiential sources; between (a) my two hands causing my experience of my two hands and (b) an evil scientist in some distant universe poking my brain to cause my experience of my two hands. In either case, I’d have the same experience. The force of EWS lies in the indiscernibility of various metaphysical scenarios, and the impact of bizarre scenarios upon our knowing. Additionally, Moore seems to concede to the skeptic that if BIV were true, I would not know that I have two hands. Metaphysical scenarios remain potent for EWS under a Moorean critique.

Nozick’s strategy against EWS falls into a similar trap. In rejecting the closure principle, the theory already has a good deal standing against it (Dretske faces the same problem).<sup>10</sup> However, the primary flaw, as I see it, is that any modal theory of knowledge requires, from the outset, that we already know what *this* world is like. But the force of EWS comes from the implications of an ontology in which our perceptions have a very different metaphysical nature (i.e. foundation and source) than we expect. It is the very nature and source of the perceived world that EWS calls into question. Nozick’s theory requires that we assess the similarity of possible worlds in relation to this world, and thereby assumes that we know the nature of this world.

Contextualist solutions to EWS seem to provide a way for EWS to be correct under skeptical standards, in a skeptical environment, without undermining the knowledge we ascribe in normal conversational contexts. But is this really a solution to EWS? I don’t think so. *Given* the skeptical context, this position does not do proper justice to the anti-skeptical intuitions that remain within that context. Such intuitions lead me to think that I can indeed know (even in the skeptical context) that I have two hands, even though I do not

---

<sup>10</sup> Namely, that the causal principle seems both intuitive and eminently plausible.

know that I'm not a BIV. Or, at least minimally, that if I cannot know that I have two hands *in that context*, then quite simply, I cannot know that I have two hands *at all*. In either case, the contextualist artificially divorces the propositions of one context from others, so that they have no bearing on one another. This leaves our epistemic life in a state of schizophrenia. Indeed, it seems quite odd to say "yes, yesterday I knew that I had a silver car but, here in philosophy class, now that you've explained the BIV hypotheses, I do not know that I have a silver car."

An additional point is in order. Let us grant that the contextualist is right. Given skeptical contexts, we cannot know that we are not BIVs. But does this imply that given the same skeptical contexts, we do not know the things we think we know in everyday contexts. EWS is much stronger on the first point: that we cannot know that we are not BIVs (indiscernability). On the second point, regarding our everyday knowledge, EWS relies on an often unstated, and unclear metaphysical implication. I will argue that one can maintain the closure principle without coming to the conclusion that skeptical scenarios imply that we cannot know everyday propositions. I think that such knowledge can be preserved *even within* a skeptical framework (despite the contextualists). I will try to show why in the following sections.

### **The Critical Issues**

Before moving on, I'd like to make a few issues salient. First, EWS has bite because it seems to challenge our everyday knowledge of the external world. Such knowledge is metaphysically loaded. Our propositions about the world imply beliefs about the nature of the world. I know that I have a silver car. The proposition "I have a silver car" means something more than "On most nights I dream that I have a silver car" or "I have a silver car but I can only see it on my computer screen." Saying that I have a silver car implies that I own some macro-level object that appears silver to me, has four tires, has a particular physical shape, provides transportation, etc. There are a limited number of conditions under which I can accurately say, "I know that I have a silver car." There are temporal conditions, physical conditions, etc. This brings up two important questions for EWS that I'd like to address. First, what sort of metaphysical scenarios fall outside the conditions under which I can appropriately be said to know that I have a silver car? And second, do skeptical scenarios fall outside the appropriate conditions? These questions

seem fundamental to an analysis of EWS because the force of this brand of skepticism lies in **the implication that certain far-out metaphysical scenarios *should* cause us skepticism about everyday knowledge**. In other words, the strength of this implication: “if I’m a BIV, then I don’t know that I have a silver car” provides the force for this implication: “if I don’t know that I’m not a BIV, then I can’t know that I have a silver car.”

### **Metaphysical Considerations**

An interesting thing about EWS is its location at the intersection of epistemology and metaphysics. EWS finds its force in positing bizarre metaphysical scenarios which are supposed to undermine everyday knowledge by implication. The details of this implication are not usually spelled out. In my view, a proper analysis of EWS should deal specifically with the epistemological effect that such metaphysical considerations have on our view of normal knowledge. The contextualist concedes that upon introducing the skeptical scenario to an epistemic context, EWS takes on an almost indestructible force. As such, the contextualist grants that it is the epistemic context of particular metaphysical considerations that lends EWS its power. My interest here is to consider whether the force of EWS can be weakened by a *different sort* of metaphysical consideration.

Here’s the intuition: consider the real world, as we understand it. Think about the macro-objects that you make knowledge claims about. I know that I have a silver car. I know that I am a graduate student at Temple University. I know that Temple University is in Philadelphia, PA. External world skepticism is designed to make you doubt whether you really know these things, by proposing that you might just be a BIV in some far off universe getting poked by an eccentric scientist. Consider the following proposition: If I were a BIV in some other universe, then I certainly wouldn't know that I was in Princeton, NJ, nor that I owned a silver car. This is where the force of EWS emerges and it is a premise that most in the EWS debate have agreed upon<sup>11</sup>. If I am right and this is the intuition that drives EWS, then it deserves some further attention. My goal is to challenge this commonly held intuition and, in doing so, to soften the force of EWS.

The best way to start is by considering how our contemporary views of modern physical theories affect our knowledge. In particular, it is important to look at the basic

---

<sup>11</sup> As far as I can tell, Moore, Nozick, Dretske and the contextualists all concede that if I could be certain that I was a BIV, then all my normal knowledge would be destroyed.

levels of physical reality. Modern physical theories provide us with an account of the world that fundamentally consists of basic particles, waves, fields, etc. We can forget, for a moment, about various skeptical scenarios such as the BIV. Modern physical theory is the picture of reality that we *really* have. Our ontology is informed by modern science, and it is a hierarchical ontology where the *really, real* things at the foundation of our world, are completely foreign to us at a macro level. But does this cause us skepticism about the macro-objects of which we claim to have knowledge? Does the fact that my car (or even my body) fundamentally consists of invisible, physical particles (or energy, etc.) cause me skepticism about macro-level objects? Most of us would answer no to this question. Our macro-level concepts refer to the objects as they appear to us.

From this analysis, the critical question becomes a metaphysical one: what is the source and fundamental nature of the objects to which our macro-level concepts refer and what are the implications (principle of closure) of these metaphysical considerations? Even in our world, one might have trouble giving an answer to this question, especially when we consider the way in which macro-level objects are constructed from invisible particles, and the fact that our scientific view of reality is largely alien to our macro-level conceptions. Given this consideration, the skeptic needs to give us a reason for seeing the skeptical scenarios (e.g. BIV) as fundamentally different from the real world in order to justify the normally employed skeptical implication. In both cases we encounter sources or foundations for our macro-level concepts that are foreign to the concepts themselves.

Macro-micro considerations of the real world make salient the need for the skeptic to more fully spell out the implication involved in the principle of closure. The difference between a skeptical scenario (e.g. BIV) as the foundation of our perception of the macro-world and the level of particle physics needs to be such that it is sufficient to enforce the implication: If I'm a BIV, then I don't know that I have a silver car. We don't, after all, say, "If the world fundamentally consists of invisible physical particles (or even some lower level), then I don't know that I have a silver car." What is it about the metaphysical considerations of the skeptic that enforces the implication against macro-world knowledge?

### **Skeptical Responses**

The skeptic may respond to the above considerations by asserting that the force of the EWS comes from *our not knowing* the source and foundation of our experienced

macro-world (i.e. indiscernability) and not a particular metaphysical scenario (e.g. BIV). Yet, given what we know about the history of science, there is a good chance that *even in this world*, there may be additional layers of physical reality that go beneath our currently perceived “bedrock.” Does such a possibility cause us skepticism about macro-world knowledge?

The skeptic might go on to distinguish between causally appropriate and gratuitous sorts of metaphysical sources/foundations for our perceived world. Under such a view, the only type of knowledge of the external world that is worth having is the sort that is caused by the appropriate objects. In the BIV scenario, an eccentric scientist causes my perception of the silver car, not the silver car itself. This sort of perception is gratuitous. However, we should then question the skeptic on what the fundamental difference is between a BIV cause and a causal sequence involving millions of photons, atoms, neurons, etc. In what sense is the car in the BIV scenario any less real than a car in the physical world? The skeptic probably has a response.

In order to answer the last question, the skeptic might now turn to a consideration of the nature of the perceived world. In the BIV scenario, the perceived world exists in your mind (and perhaps in the mind of the eccentric scientist). In the physical world, the perceived world refers to external objects. But what is the nature of these objects? They are made of physical particles, whose intrinsic nature we simply don’t know; modern physical theory doesn’t tell us. In fact, some scientists have even posited that the physical world is itself, in some sense, an aspect of something else, such as information. Conceivably, Berkeley could be right, and the information that grounds reality could be information in the mind of God (or some evil machines). If this were the case, in what sense would the “real” world be any different than the BIV world?

### **Drawing boundaries on a legitimate ontology**

I’ve tried to make clear that in order to maintain the strong force of EWS the skeptic needs to draw sharp metaphysical lines between skeptical scenarios and our beliefs about the real micro-physical world. I don’t think that the skeptic can achieve this task, and I think that there is a reason. Our concepts refer to a macro-level world that is ***robust, consistent, comprehensive and coherent***. As such, we have knowledge of this consistent and comprehensive world without making reference to the fundamental nature of the

perceived world. It is in *this macro-world* that we know. It seems that the boundaries, or metaphysical conditions, that the skeptic must draw are not between this world and various bizarre metaphysical scenarios, but between worlds that are both consistent and comprehensive, and those that are not. If this is right, then proper skeptical scenarios are those that either lack coherence from moment to moment or are insufficiently comprehensive.

Under this view, knowledge does not require that we have access to the “bedrock” of reality. Knowledge is “level” independent, and propositions merely need to make reference in perceived worlds that are both robust and coherent. This allows our concepts to persist in reference, regardless of their fundamental base. Borrowing from the philosophy of mind, one might call this view the multiple-realizability of macro-world concepts, where macro-world concepts are not necessarily affected by the nature of their source.

### **Conclusion**

The force of external world skepticism results from the epistemic impact of bizarre metaphysical scenarios on our everyday, macro-level concepts. There is a common intuition that our macro-level concepts would be false and unjustified given the truth of something like a BIV scenario. However, this common intuition can be called into question through another metaphysical consideration: that of our own micro-physical world. It seems difficult to see a critical difference between a BIV scenario and the discrepancy between our macro-level concepts and our (surely flawed and incomplete) fundamental view of physical reality. Yet, the skeptical hunch arises in the first case and not the second. This discrepancy highlights a dilemma. Either we should be skeptical without skeptical scenarios (i.e. the real world is a skeptical context), or this analysis should soften the force of EWS. The important question for EWS becomes a metaphysical one: to what degree should the source and fundamental nature of the macro-level objects of our perception affect our epistemology? On the one hand, perceptual anomalies (i.e. illusions) should not count as knowledge. On the other hand, large, consistent and robust perceptual worlds seem well suited for forming real knowledge *in reference to that world*.