

## **Navigating Skepticism: Cognitive Insights and Bayesian Rationality in Pinillos'**

### ***Why We Doubt***

*Chad Gonnerman*

Political Science, Public Administration, and Philosophy, University of Southern Indiana,  
Evansville, IN, United States

*cgonnerman@usi.edu*

*John Philip Waterman*

Department of Philosophy, University of New England, Biddeford, ME, United States

*jwaterman1@une.edu*

### **Abstract**

Pinillos' *Why We Doubt* presents a powerful critique of such global skeptical assertions as "I don't know I am not a brain-in-a-vat (BIV)" by introducing a cognitive mechanism that is sensitive to error possibilities and a Bayesian rule of rationality that this mechanism is designed to approximate. This multifaceted argument offers a novel counter to global skepticism, contending that our basis for believing such premises is underminable. In this work, we engage with Pinillos' adoption of Bayesianism, questioning whether the Bayesian principle that he invokes truly does fail to generate the verdict that we don't we aren't BIVs, contrary to what Pinillos asserts. Furthermore, we argue that Pinillos' empirical argument is not empirical enough; we need a lot more empirical work if we are going to counter global skepticism and win over neutral bystanders.

## **Keywords**

Global skepticism, error possibilities, Bayesianism, epistemic cognition, experimental epistemology

## **1. Introduction**

*Why We Doubt* is an outstanding book. In it, Pinillos tackles an impressive range of topics, including the role that skeptical inclinations play in practical reasoning, obsessive-compulsive disorder, and conspiratorial thinking. But, despite all the great stuff happening in the book, in our contribution to this symposium, we are going to focus on just one part of the overall project: Pinillos' response to the global skeptic. Pinillos' reply is one part cognitive and one part Bayesian. Briefly put, his thought is that, when we plug these two into an independently plausible principle of defeat, our justification for believing such global skeptical premises as "I don't know I am not a brain-in-a-vat (BIV)" is undermined. In what follows, we will make a handful of comments about the Bayesian component of Pinillos' reply and about two anti-skeptical arguments that can be extracted from his materials; our interests, however, will mostly lie in the cognitive part of the overall picture. Our central claim will be that, while Pinillos has made significant contributions in developing a compelling response to global skepticism, a richer model is needed and, to get there, we need a lot more data.

## **2. Anti-Global Skepticism 1**

Pinillos begins his book by identifying the form of skepticism in which he is interested and, by extension, the forms in which he isn't. For example, lying outside the scope of his work is the

Humean skeptic who argues that any justificatory story that we may tell on behalf of our beliefs about the external world is ultimately circular, and thus these beliefs will inevitably fall short of knowledge (2023: 31–32). What seems to matter to Pinillos are skeptical inclinations and arguments that are rooted in error possibilities. Think about the possibility that your car has been stolen, that your Powerball ticket just might win, or that you may be a BIV, and you might find yourself inclined to deny that you know where your car is, that your ticket will lose, or that you aren't a BIV. This focus on error possibilities aligns, of course, with Descartes' skeptical arguments in the First Meditation. But not all the skeptical arguments discussed by Pinillos are as destructive as Descartes' deceptions, dreams, and demons are designed to be. As such, to call the form of skepticism in which Pinillos is interested "Cartesian" can be misleading. We might dub it "Errorresianism" instead. This is a neologism that captures some of the similarities with Descartes' skeptical arguments (while also being silly enough that no one is ever likely to use it).

A commonly taken route to resisting Errorresianism is to argue that one or more of the premises in their argument is (likely) false. To illustrate, consider the following global skeptical argument (cf. Pinillos 2023: 26):

*Global Skepticism*

1. I don't know I am not a BIV.
2. If I have two physical hands, then I am not a BIV.
3. Closure.
4. Therefore, I don't know I have physical hands.

In response, some folks have pushed back against the first premise—the *global skeptical premise*. Reid (1785/1983), for example, claims that we are *a priori* justified in believing that perception is reliable, which, if true, would rule out the BIV scenario. Others have taken an *a posteriori* approach, including the Moorean “But I have two hands” gambit (Moore 1939) and Russell’s (1912) suggestion that the best explanation of our experience is that we live in a world whose nature aligns, more or less, with what we commonly take it to be. Although these responses differ in many ways, they have a common goal: to show that we do (or can) know that we are not BIVs, contra the global skeptical premise.

Pinillos, however, takes a different tack. His goal, at least to a first approximation, is to show that we are not justified in believing the global skeptical premise. The argument has three main parts. First, there is a cognitive account of the mechanisms, processes, or systems that give rise to many of our skeptical inclinations. Second, Pinillos formulates a Bayesian rule of rationality that his skeptical mechanism has the function of approximating. Third, there is a principle of defeat that entails an undermining of the justification had by our heuristic beliefs when the relevant conditions are met. We describe these three in more detail over the next few paragraphs.

The first element in Pinillos’ anti-skeptical argument aligns with a lot of research in experimental epistemology over the last couple of decades. It is an attempt to work out the psychology behind our skeptical inclinations. Pinillos’ proposal is that many (though perhaps not all) of our skeptical leanings arise out of a cognitive mechanism that implements a sensitivity principle. Importantly, it is not a principle of sensitive belief. So, Pinillos’ pitch is not simply a psychological version of Nozick (1981). His sensitivity principle is about (what is taken to be) the agent’s principal basis, which he defines as “the principal cause for the [agent’s] belief among all the bases” (2023: 66). We can formulate Pinillos’ *Principal Base Sensitivity* (PBS) principle along

these lines: S knows  $p$  only if S's principal basis,  $E$ , for believing  $p$  is sensitive to  $p$ , that is, only if  $E$ 's truth value would be different were  $\sim p$  the case (cf. 2023: 65). When it comes to the BIV scenario, the underlying PBS mechanism might enact something like the following line of reasoning: "My principal basis for my belief that I am not a BIV is that it sure seems to me that I have two hands; however, in the nearest possible worlds in which I am a mere BIV, it'd still seem to me that I have two hands. So, my principal basis is insensitive to the proposition that I am not a BIV; hence, I don't know I am not a BIV." In some such way, it is posited that we can come to feel the grip of the global skeptical premise, as well as many other knowledge denials.

The second part of Pinillos' anti-skeptical counterargument is a bit of Bayesian epistemology. It is a rule of rationality that he calls "Norm" (2023: 138, 151):

(Norm) If (a)  $C_i(E|\sim p) = 1$  and (b)  $C_i(p) < n$ , (c)  $E$  is the strongest proposition learned after  $i$  (up until and including time  $f$ ), then  $S$  does not know  $p$  at  $f$ .

The claim being made in Norm is, if (a) your rational confidence in evidence  $E$ , conditional on not- $p$ , is maximal at time  $i$ , and if (b) your rational confidence in  $p$  at  $i$  falls below an *ignorance marker*,  $n$ , entailing that you don't know  $p$  at the time, then you still won't know  $p$  at time  $f$ , if it is also true that (c)  $E$  is the strongest proposition acquired between  $i$  and  $f$  in the sense that it entails all other propositions acquired during the period. To be sure, the principle has substantial intuitive appeal. If your evidential standing wasn't strong enough to put you in a position to know  $p$  before acquiring some bit of evidence  $E$  and if you are certain that  $E$  is insensitive to  $p$ , and thus that acquiring  $E$  doesn't offer any additional reason to believe or reject  $p$ , then acquiring  $E$  won't put

you in a position to know  $p$ . So, if  $E$  was the most powerful or compelling evidence gained within the given timeframe, then your lack of knowledge will persist at the end of the timeframe.

The third piece in Pinillos' anti-skeptical rebuttal is a principle of epistemic defeat, aptly dubbed "Defeat." Pinillos formulates the principle thusly:

Defeat: If you believe heuristic  $H$  is supposed to approximate norm  $N$  and that  $H$  predicts  $p$ , but you also believe that  $N$  does not predict  $p$  (or withhold belief that  $N$  predicts  $p$  as a result of considering the issue), then you are not justified in believing  $p$  if this belief arose through  $H$ . (2023: 149)

In other words, the (non-parenthetical) claim is that a reduction in your epistemic justification will result when four beliefs are present. The first is the belief subject to the reduction. Insofar as Defeat is concerned, it needs to be a belief derived through heuristic processing. To illustrate, suppose you are an experienced gutter installer, and you believe that you'll need 20' of downspout based on an eyeball estimate of the run from the roof to the ground. The second belief is about the function of the heuristic at play. You believe its purpose is to produce results that align reasonably well with the deliverances of some certain norm. So, if you are our eyeballing gutterer, you might believe that eyeball estimates are supposed to approximate tape-measure results. The third belief is about the outcome of the heuristic as applied to your situation. Specifically, you believe that the heuristic produces what ends up as the content of your heuristic-based belief. So, as our gutterer, not only do you believe that you'll need 20' of downspout, you will also believe that your eyeball estimate says that you'll need this amount. The final belief focuses on how the norm plays out in the situation at hand. You believe that the relevant norm does not produce a result that matches the

content of your heuristically derived belief. So, after eyeballing the drop, you might climb your ladder and measure the distance with your tape measure, revealing it to be 21'. In this situation, Defeat would entail that you are no longer justified in believing that you'll need 20' of downspout. Putting PBS, Norm, and Defeat together, we get something like the following argument against Global Skepticism from above:

*Anti-Global Skepticism 1*

1. You believe that PBS is supposed to approximate Norm (thanks to the arguments of Chapter 5).
2. You also believe that PBS predicts that you know don't know you aren't a BIV (in view of the work presented in Chapter 2).
3. In addition, you believe that Norm does not predict that you don't know you aren't a BIV (because of the discussion found in Chapter 6).
4. Your belief that you don't know that you aren't a BIV arose through PBS.
5. Defeat.
6. Therefore, you are not justified in believing that you don't know you aren't a BIV.

One immediate consequence of this argument is that whatever justification you had in favor of belief in the global skeptical premise (i.e., premise 1 of Global Skepticism above) has now been undermined, leaving it insufficiently motivated to ground the skeptical conclusion that you don't know you have two hands.

Clearly, Pinillos has developed a response to global skepticism that is insightful, novel, and forceful. But, like all excellent philosophy, it prompts many questions. In the next section, we

focus on one of the arguments that Pinillos gives in connection with the third premise of Anti-Global Skepticism 1, using our discussion as an opportunity to set up a slightly different argument, viz., Anti-Global Skepticism 2. In Section 4, we address some of the more descriptive components of these two arguments.

### **3. Argument from Indirect High Priors**

One question that may occur to some readers pertains to premise 3 in Anti-Global Skepticism 1. Why think that Norm doesn't generate the verdict that I don't know I am not a BIV? On this matter, Pinillos puts forward three arguments: Uncertain Insensitivity, Indirect High Priors, and Direct High Priors. Each one merits close attention; here, however, we mostly restrict our attention to the Argument from Indirect High Priors.

#### *3.1. Second Horn*

Pinillos' Argument from Indirect High Priors has a dilemma-like structure. We start with second horn of the dilemma:

Let's suppose instead that the relevant  $E$  is the information I gather right now (today) as I am considering the issue of whether I am not a BIV. [...]

To apply Norm in this second case, I need to determine my prior credence for  $p$  in this scenario. For the principle to predict that I don't know  $p$ , my prior rational confidence must be too low to count as knowing—it must be below the ignorance markers. But what was my rational confidence in  $p$  yesterday? Presumably I had very high rational confidence on ordinary propositions such as “I have hands” and “I am drinking coffee” which are each



at or above the knowledge ignorance markers (we are not assuming that these propositions are always known since being at or above the ignorance markers is not sufficient to know). However, since these ordinary propositions each entail  $p$  (I am not a BIV), then my rational confidence in  $p$  must have also been higher than the ignorance markers. This is because probability cannot be reduced across single-premise entailments. This just means that (b) is false in this instance of Norm, and hence Norm fails to make the skeptical prediction (that I don't know  $p$ ). (2023: 152–153)

In standard form, we might express Pinillos' argument along these lines:

*Argument from Indirect High Priors (Horn 2)*

1. Suppose my evidence that I am not a BIV is that I have hands.
2. My rational confidence in the proposition I have hands, prior to considering the BIV scenario, is very high, exceeding all ignorance markers.
3. If I have hands, then I am not a BIV.
4. Rational confidence is closed under single-premise entailments.
5. So, prior to considering the BIV scenario, my rational confidence in the proposition I am not a BIV was also very high, exceeding all ignorance markers (by 2, 3, 4).
6. Condition (b) in Norm holds only if my rational confidence falls below all ignorance markers.
7. Therefore, prior to considering the BIV scenario, condition (b) of Norm doesn't hold.

If we have gotten this argument right, then its success seems to depend on what occurs to the reasoner when she considers the BIV scenario. In the quote, Pinillos depicts her as relying on statements like “I have hands” or “I am drinking coffee.” These work at preventing condition (b) in Norm from being met because they entail that the person isn’t a BIV. After all, BIVs don’t have hands, and they don’t drink coffee.

But not all propositions about external facts work at blocking condition (b) nearly as well as these two. For example, “There is a coffee cup in front of me” and “I am in a room with 20 bright students” are both logically compatible with the reasoner being a BIV. But more to the point, when first presented with the BIV error possibility, how often does it occur to us to appeal to statements like “I have two hands”? Isn’t the overwhelming tendency to take the skeptic’s bait and appeal to something internal along the lines of “It seems to me that I have two hands”? This is the move that Descartes makes in the First Meditation. For example, after raising the dream hypothesis, he considers (and rejects) the possibility that such things as stretching out his hand “would not happen with such distinctness to someone asleep” (1641/1988: 77). Here, Descartes has risen to the skeptic’s bait. And if Pinillos is right about the Moorean “I have two hands” move, we should expect anti-skeptical appeals to external facts to be rare: for some reason or other, we tend to find this move to be unsatisfactory (2023: 85–88).

The main takeaway of our discussion thus far is that there are reasons to think that the second horn of the Argument from High Indirect Priors will work only for a smallish number of people—the “default Mooreans,” as we might call them. But if Pinillos is right about the psychology underlying (most of) our skeptical inclinations, these folks are precisely the ones who are unlikely to experience the global skeptical inclination in the first place. As he explains (2023: 86), if your principal basis for believing that you aren’t a BIV is an external fact like that of you

having two hands, then your principal basis *will be* sensitive to the content of your belief, since in the nearest possible worlds in which you are a mere BIV, you won't have two hands. As a result, your PBS mechanism won't deliver the skeptical inclination.

In brief, then, our worry with the second horn of the Argument from Indirect High Priors is that it seems to work best, if not only, on those who need it least—the default Mooreans. (These folks are also the ones least likely to accept premise 4 in Anti-Skeptical Argument 1, for what it is worth.) This leaves us with a pressing question: how does the second horn of the Argument from Indirect Priors play out for the rest of us?

### 3.2. *First Horn*

Pinillos articulates the first part of the Argument from Indirect Priors as follows:

For (b) [of Norm] to be true, it would have to be the case that at time  $i$ , prior to learning  $E$ , my rational credence in  $p$  was below the ignorance marker (recall that any credence below the ignorance marker does not yield knowledge). But to determine this I need to know what  $E$  is supposed to be. If it is the totality of experiences or beliefs I've ever had, then  $i$  must be the time when I was barely sentient. In that case, I doubt there could be a probability function  $C_i$  defined over complex propositions like  $p$  that captures my rational doxastic states, if any, at that time. If there isn't such function, then (b) could not be true. Norm would then not predict SK (I don't know I am not a BIV). (2023: 152)

The idea being put forward is that if the evidence that matters vis-à-vis Norm and your belief that you aren't a BIV is the totality of your experiences or beliefs up to and including the time you encountered the BIV scenario, then assessing condition (b) on Norm will involve determining what

your rational confidence in “I am not a BIV” was before you acquired this totality. According to Pinillos, this involves rewinding the clock back to when you were barely sentient. But now, Pinillos observes, we run up against a strange question: what was your rational confidence in the proposition that you weren’t a BIV when you were an infant or a toddler? The best response may be that there is no answer to this question. As Pinillos puts it, it is doubtful that a probability function exists for assigning rational credences in complex propositions for barely sentient children. If so, then condition (b) won’t be met when the relevant evidence consists of your totality of experiences and beliefs. And thus, Norm won’t deliver the verdict that you don’t know you aren’t a BIV, which gets us close to premise 3 in Anti-Global Skepticism 1.

We aren’t certain that the preceding line of argumentation is quite right, though there may be another route to the conclusion available to Pinillos. One of the key moves in Pinillos’ argument is that if the pertinent evidence for assessing the applicability of Norm to your non-BIV belief is the totality of your experiences and beliefs, then the relevant time slice for assessing your prior confidence in the non-BIV proposition is when you were barely sentient. Do we need to go that far back? Consider the totality of your experiences and beliefs up to and including the time that you considered the BIV scenario. Let’s say that this happens at time  $f$ . Then  $E$  will be the totality of your experiences and beliefs at  $f$ . Call this totality  $T$ . Importantly, just moments before  $f$ , you didn’t have  $T$  quite yet. Instead, you had a set of experiences and beliefs that looks a lot like  $T$ . But this set—call it  $T^*$ —will be missing some experiences or beliefs present in  $T$ . So, on this way of approaching the matter, we shouldn’t ask what your rational credence was in “I am not a BIV” when you were barely sentient; we should ask what it was just moments before you were considered the BIV error possibility.

If the argument we just sketched works, then we need to ask: what was your rational confidence in the proposition that you aren't a BIV just prior to being presented with the error possibility? We don't know the answer to this question. But, since the beliefs found in your totality of beliefs and experiences at the time includes many beliefs about the external world like the belief that you have two hands, perhaps a Pinillos-friendly case can be made along the lines of the second horn of the Argument from Indirect High Priors. We leave this matter to Bayesians to work out.

#### **4. Anti-Global Skepticism 2**

What we would like to do next is to draw the reader's attention to a feature of Norm that has yet to make much of an appearance in our discussion: condition (c). A case could be made that this condition may make it a bit too easy to argue that Norm doesn't generate the verdict that you don't know you're not a BIV. But, as we'll suggest, this consideration, in concert with claims suggested in the previous section, might be spun in Pinillos' favor, helping to generate a second anti-skeptical argument.

It turns out that not any old bit of evidence is a candidate for Norm. As we noted in Section 2, *E* needs to be the strongest proposition learned during the timeframe. Here, *strongest proposition learned* is defined as the proposition that was acquired during the timeframe that entails all other propositions learned during the period (2023: 136). In addition, Pinillos writes: "We assume that learning a proposition is assigning it probability 1..." (2023: 136). Those are some mighty high hurdles to clear, and it is not obvious to us that the evidence to which Pinillos appeals in laying out, for example, the Argument from Indirect High Priors manages to clear them. Consider having two hands. If this is the evidence that is relevant to your belief that you aren't a BIV, then it is hard to see how condition (c) of Norm could be satisfied. The proposition would have to (i) have been

acquired during the timeframe, (ii) entail all other propositions learned during the period, and (iii) be assigned a probability 1. But for (i) to hold of “I have two hands,” evidently, the timeframe would have to go very far back, perhaps to when you were barely sentient. That would mean that (ii) is unlikely to hold: since the proposition doesn’t entail all that much, it will often cease to be the strongest proposition once anything else is learned. Moreover, the proposition seems ill-suited for an assignment of probability 1. This value is usually reserved for the truths of logic and mathematics.

Indeed, thinking of a candidate that meets conditions (c) of Norm is not as easy as it might seem. A mathematical truth learned last week may be given the proper probability level, but it won’t entail everything learned over the week if you are a bit too free in their probability assignments (e.g., assigning probability 1 to “George Washington crossed the Delaware”). Maybe what would work is a mathematical truth learned just prior to considering the BIV scenario. It could be given probability 1, and it would trivially entail all the propositions learned in the relevant timeframe if it was the only thing learned. But a mathematical truth would make for a weird bit of evidence on which to base your belief that you aren’t a BIV. Can defeating Global Skepticism be as easy as pointing out that any natural candidates for satisfying condition (c) of Norm either won’t entail everything learned in the relevant timeframe or won’t be assigned maximal confidence?

But perhaps all this unclarity about how Norm plays out with respect to your belief that you aren’t a BIV can be spun in Pinillos’ favor. Recall that his principle of defeat includes a parenthetical: “If you believe heuristic  $H$  is supposed to approximate norm  $N$  and that  $H$  predicts  $p$ , but you also believe that  $N$  does not predict  $p$  (or *withhold belief that  $N$  predicts  $p$  as a result of considering the issue*), then you are not justified in believing  $p$  if this belief arose through  $H$ ” (2023: 149; emphasis added). Maybe the main takeaway is that it is simply unclear what Norm has

to say about “I am not a BIV.” If so, then you might withhold belief as to whether Norm predicts that you don’t know that you aren’t a BIV, which would make for a slightly different argument against Global Skepticism:

### *Anti-Global Skepticism 2*

1. You believe that PBS is supposed to approximate Norm (thanks to the arguments of Chapter 5).
2. You also believe that PBS predicts that you know don’t know that you aren’t a BIV (in view of the work presented in Chapter 2).
3. In addition, you withhold belief on whether Norm predicts that you don’t know you aren’t a BIV (because of the discussion found in Chapter 6).
4. Your belief that you don’t know that you aren’t a BIV arose through PBS.
5. Defeat.
6. Therefore, you are not justified in believing that you don’t know that you aren’t a BIV.

Notice that this argument is nearly identical to Anti-Global Skepticism 1. Its only difference is in premise 3. Specifically, instead of claiming that you believe that Norm doesn’t predict that you don’t know you aren’t a BIV, this version asserts that you merely withhold belief on the matter.

## **5. Developing an Effective Response to Global Skepticism**

### *5.1. Pinillos’ Audience*

One question that may occur to some is: who is the audience for Pinillos’ two anti-global skeptical arguments? It is natural to think that his intended audience includes the global skeptic. But if so,

then one might wonder about the dialectical effectiveness of his counterarguments. As Pinillos appears to recognize (2023: 156–157), many global skeptics are likely to claim that his arguments beg the question against their position. Do they? In part, what makes assessing this charge so difficult is that the conditions on begging the question are less clear than textbooks in informal logic often suggest. There are entire monographs written on the fallacy, after all (e.g., Walton 1991). With that said, a case could be made that Anti-Global Skepticism 1 and 2 are guilty of the fallacy. Specifically, to the extent that begging the question involves relying on premises that you have reason to think that your interlocutor doesn't accept (Hazlett 2006), it is hard to see how Pinillos could avoid the charge. If global skeptics don't accept propositions about being two-handed, it is unlikely that they will allow the claims of advanced cognitive scientific research, including research of the sort to which Pinillos appeals in motivating his cognitive account of our skeptical inclinations.

Then again, it is likely that Pinillos' primary audience doesn't include the global skeptic. Notice that his principle of defeat uses the second-person pronoun 'you' (hence our use of the word in our formulation of Anti-Global Skepticism 1 and 2). This suggests that Pinillos' arguments are directed at his readers, most of whom are probably not global skeptics since they make for a rare breed. Also, after noting that the global skeptic might not be persuaded by the arguments that he gives in response, Pinillos asks us to "imagine an agent who is not sure that she is justified in believing the skeptic's premise" (2023: 157). It appears, then, that Pinillos' primary audience is a neutral bystander, someone who is on the fence as to whether she is justified in believing that she doesn't know she is not a BIV. One advantage of targeting the neutral bystander is that worries about begging the question largely dissipate. We have no reason to think that the neutral bystander objects to empirical considerations. Instead, we might imagine this person taking on a role akin to



a judge. She is sitting at an imaginary bench. The global skeptic has just had their turn in giving their argument, and she is now ready to hear Pinillos' counterarguments, and, if those arguments include empirical results, so be it.

With that said, there are challenges in turning to the neutral bystander. What qualifies a person as such? And how do we work out her reactions to the two sets of arguments? For example, consider us, the authors of this symposium contribution. Would we count as neutral bystanders? If so, then it might be worth noting that the anti-global skeptical arguments didn't work on us. To be sure, we believe Pinillos' arguments are strong; some of the best that we've seen. It's just that, despite the impressive range of considerations that Pinillos leverages on behalf of the first and second premises of Anti-Global Skepticism 1 and 2 (viz., "You believe that PBS is supposed to approximate Norm" and "You believe that PBS predicts that you know don't know that you aren't a BIV"), our reaction was, "Well, no, not quite." Belief just felt like too far of a bridge for us. Instead, the belief contents struck us as two eminently plausible hypotheses, well worth further exploration. Then again, perhaps we aren't neutral (or rational!) enough to count as neutral bystanders. We are philosophers, after all. So, let us set aside the possibility that we are among the neutral bystanders that Pinillos has in mind and consider how some such person might or should respond to the argumentative exchange.

### *5.2. Taking a Cognitive Turn*

A complaint commonly directed toward skeptical refutations is that they do not produce conviction, often failing to persuade others of the anti-skeptical positions. Indeed, a shared feature of many anti-skeptical arguments is that they are far more complicated and far less intuitive than the skeptics' own. This observation, first explored in detail by Stroud (1984), has played an

important role in shaping contemporary responses to global skepticism. Consider Moore's famous demonstration—first waving one hand in the air and then the other. It is a virtue of his response that it is as compact and direct as the skeptic's own argument. But at best it results in a stalemate. At least arguably, it falls short of a fully satisfactory response because it does not address the original lure of our skeptical intuitions. If our knowledge of the external world is basic and beyond question, then why is the skeptic's position so attractive and powerful?

Pinillos is clear that an important virtue of adopting a cognitive approach like his is that it can provide an explanation of our attraction to skepticism as an extension of quotidian forms of doubt. In this way, we can view Pinillos as positioning his cognitive account in line with a tradition of recent approaches to skeptical problems, perhaps best exemplified by linguistic contextualists like Stewart Cohen and Keith DeRose, who aim not merely to respond to skepticism, but also to explain its appeal. Pinillos writes: "...it is not enough to resolve this tension by explaining which of our commitments or attractions lead to truth and which lead to falsity. A satisfactory explanation needs to say why we have these commitments and feel the attraction to skepticism in the first place. We can think of this type of account as a bit of therapy to calm our intellectual anxieties" (2023: 48).

As we will contend, by taking a cognitive turn, Pinillos has contributed to a promising line of inquiry in the epistemology of skepticism. A more fine-grained description of the processes, systems, or mechanisms that give rise to our skeptical inclinations has the potential to offer a more thorough diagnosis of the conditions under which we oscillate between unquestioned belief and skeptical doubt, and thereby provide a more complete kind of therapy. And yet, as we will argue, making the most of the cognitive turn in replying to skepticism, even when it comes to the neutral bystander, will require—perhaps counterintuitively—a much more thorough commitment to

gathering and working from empirical data than has so far been attempted by anyone, including Pinillos.

The cognitive turn that Pinillos takes can be understood as part of a larger and comparatively recent tradition in philosophy that aims to provide a response to the global skeptic while also giving an error theory to explain the appeal of their argument. Michael Williams dubs the proponents of this approach “The New Sceptics” (1996: xvi). Many New Sceptics assume that “there is something in the concept of knowledge that includes a standing invitation to scepticism,” to borrow a phrase from Bernard Williams (2015: 47). More broadly, the New Sceptics start with the idea that global skepticism is “intuitive” or “natural” in some way; or, as Pinillos frames it, “the attraction to skepticism is a typical human response” (2023: 37). If the inspiration for this approach arises from Hume’s observation that skeptical inclinations are the inevitable conclusion of *rational reflection*—or, as he puts it, “sceptical doubts arise naturally from profound and intense reflection” (1739/1975: 218)—modern applications have tended to propose that our skeptical inclinations are a byproduct of such psychological machinery as the *availability heuristic* (e.g., Williamson 2005), *focal bias* (e.g., Gerken 2013), and *epistemic egocentrism* (e.g., Nagel 2010).

The New Sceptical approach can be understood as taking many forms, with linguistic contextualism being perhaps the most prominent contemporary example of this tradition. A cognitive approach of the sort proposed by Pinillos is an exciting advance on linguistic approaches to anti-skeptical replies because it offers the possibility of a more complete description of the conditions in which skeptical inclinations arise—conditions that seem to vary more widely than those encompassed by the semantics of ‘knows’. Whatever those semantics may be, they can be employed to convey attitudes that result from a rich variety of underlying information processing (Williams 1996: 26–31). One thing that the work in the psychology of heuristics and biases

tradition show, as exemplified by such well-known phenomena as confirmation bias and disconfirmation bias, is that there can be a great deal of underlying structure that is relevant to understanding our assertions and commitments that is left unexplained by the semantics of the terms that we use to convey those commitments. Thus, understanding the underlying processes, mechanisms, or systems that give rise to our skeptical inclinations and by which they are transformed into skeptical arguments and positions is surely essential to any plausible diagnosis of the skeptic's pull.

But for all this, and for all the promise of the diagnostic approach of the New Sceptics, it is by no means universally conceded that skeptical inclinations are *in fact* all that “natural” or “intuitive,” or that they depend on nothing more than our ordinary epistemic practices. The *neopragmatist tradition*, one that Williams sees running through Wittgenstein and expressed in the ordinary language philosophy of Austin, contends that global skepticism and many of its other forms only becomes plausible from a distinct theoretical perspective, one where every claim can be given and owes justificatory grounding that is open to question. As Clarke (1972: 754) puts it: “What is the sceptic examining: our most fundamental convictions, or the product of a large piece of philosophizing about empirical knowledge done before he comes on stage?” A great deal hangs or falls depending on how we answer this question, depending on which perspective—new sceptic or neopragmatist—we adopt. If global skepticism is “intuitive” and “belongs to the human condition” to use Williams’ apposite turn of phrase (1996: 12), then at some level skepticism demands a response, maybe a concessive one. But if skepticism is “a problem internal to a set of theoretical ideas that we are not bound to accept,” then a satisfactory escape from the problem may be possible by showing that dependence. A fundamental insight of Williams (1996) is that the answer given to this question determines what kind of response to the skeptic is possible. Much,

then, hangs on the question, are skeptical doubts “intuitive” or “ordinary”? But how then to answer this question? Pinillos turns to recent results in experimental philosophy to provide evidence that the skeptical pull is indeed intuitive. But it is by no means clear that the empirical literature supports this conclusion, especially if what is wanted are strongly felt skeptical inclinations.

### *5.3. Salience Effects and the Need for a More Powerful Model*

The literature on salience (aka, salient alternative, skeptical pressure, or error possibility) effects plays an important role in Pinillos’ case that skepticism is intuitive (or is an outgrowth of widespread inclinations), as well as in motivating his PBS mechanism.<sup>1</sup> A *salience effect* arises when merely mentioning an unrealized error possibility reduces or even forestalls our willingness to ascribe knowledge to an agent when we would have otherwise ascribed knowledge if the alternative had remained unmentioned. Salience effects have been influential in the literature on linguistic contextualism and relevant alternative epistemology. They’ve also been the subject of many empirical studies (for reviews, see Pinillos 2016, Gerken 2017, and Dinges 2019). Salience effects are often appealed to as a way of explaining how we take our first step in global skeptical arguments: if we tend to judge that we don’t know  $p$  when an unrealized but uneliminated error possibility is made salient, then when the skeptic raises the possibility that we may be mere BIVs,

---

<sup>1</sup> Pinillos also reviews the empirical literature on folk judgments of lottery-style arguments to make his case that skeptical intuitions are typical of human responses. From our perspective, he is right to gather as much and as broad an empirical case as he can in defending his position on the intuitive nature of skeptical doubts. We focus on salience effects because of their familiarity and because we believe there is an important nuance in their interpretation that is relevant to deciding this question.

we are inclined to judge that we don't know we're not BIVs. But once we give into this inclination, the global skeptical conclusion follows quickly and almost inevitably.

In reviewing the results of these studies, Pinillos concludes that they show that salience effects, and thereby skeptical inclinations, are a typical human response, found across cultures and throughout ordinary epistemic practice, in ways that are not theoretically laden (2023: 38). But a closer look at the empirical results adds important nuance to the conclusions we should draw. Consider that the first empirical studies of salience effects did not find that folk inclinations mirrored the intuitions of professional philosophers (Buckwalter 2010, May et al. 2010). Later studies, using different scenarios, did find that mentioning uneliminated error possibilities can reduce ordinary individuals' willingness to attribute knowledge, and to sometimes withdraw it (e.g., Schaffer & Knobe 2010, Nagel, San Juan & Mar 2013, Buckwalter 2014). But while there are many possible explanations for this difference (for overviews, see Gerken 2017 and Dinges 2019), it does suggest that the intuitions are unsteady, if not delicate, depending on the materials and conditions of presentation. Overlooking this is likely a mistake. Moreover, while several empirical studies show that salience effects do exist, and that they appear to persist across potential moderators such as need-for-cognition (Alexander, Gonnerman & Waterman 2014) and participant perspectives (Gerken et al. 2020), they also show that their strength can vary, and that the mention of an uneliminated possibility of error does not always cause a subject to deny knowledge where they would once attribute it, but rather to attribute it with less confidence. This fact deserves attention, too.

It is worth stressing that philosophical skepticism, whether in local or global forms, is not merely a statistically detectable diminution of confidence in the proposition that we know; it is to admit that knowledge of something we would ordinarily assert with great confidence is, in fact,

not had—and is impossible to have. The same studies Pinillos appeals to as showing that salience effects appear across cultures also suggests that different scenarios engender reduced inclinations to ascribe knowledge of different forces across individuals, and that the size of the effect also varies across cultures (Waterman et al. 2018). Something similar may be true of scholars from different academic backgrounds (Starmans & Friedman 2020).

Global skepticism is a dark room we only rarely find ourselves lost within. Developing an effective response to the global skeptic, if it is to calm our anxieties, must do more than merely describe how the skeptic's argument gains a grip on us; it must explain when it doesn't. To appreciate the point, consider again the intended audience: the neutral bystander. She has just heard the global skeptical argument (or was just reminded of it), but she is unsure whether she is justified in believing that she doesn't know she's not a BIV. The skeptic's argument is short and sweet, and, presumably, exerts some pull on the bystander. To effectively counter an argument of this sort, the cognitivist is going to need an especially compelling account of what gives rise to our skeptical inclinations. And the persuasiveness of this account is going to be a direct function of its predictive power. For, as Musgrave (1974: 2) observes, "in assessing the confirmation or evidential support of a hypothesis, we must take into account especially (and perhaps even exclusively) the success or failure of its *predictions*." As such, merely putting forward and defending a blueprint of the architecture of the underlying psychology is unlikely to win over the bystander, or at least it shouldn't. What is preferable, if not wanted, is a full-blown model that predicts and explains *variability in skeptical inclinations*: for example, why some materials work and don't work; why we observe a reduced tendency to ascribe knowledge here but a flat-out denial of knowledge there;

and why skepticism lands differently in this population than this other population.<sup>2</sup> Having a model of this sort will help to instill in the bystander justified confidence that the anti-skeptical cognitivist has correctly and exactly described the path that so many of us follow *en route* to the global skeptical inclination. Thus, perhaps somewhat ironically, rather than retreating to the *a priori* and inward-looking stance, it may be that developing an effective response to the global skeptic will require a great deal more empirical data than we have so far.

Without question, Pinillos has presented a promising framework for developing a comprehensive model of our skeptical inclinations with robust predictive power. And it is worth noting that he has taken some of the most substantial steps in the literature at developing a cognitive account of relevant and irrelevant alternatives (2023: 219). Despite these advancements, we believe that his PBS and metacognitive framework fall short of fully meeting the challenge of

---

<sup>2</sup> In motivating his cognitive approach, Pinillos argues that the gold standard for any kind of causal theory is data produced by experimental methods (2023: 14). Yet he also argues that the experimental method has its drawbacks, and that empirical studies inevitably rely on the responses of individuals who are not used to making subtle distinctions and may not be attending closely to the scenarios they're examining. To overcome this problem, Pinillos is explicit in his willingness to "draw on professional philosophers' judgments" to help guide our interpretation of cases (e.g., 2023: 14). But, if putting forward a compelling counterargument to Global Skepticism requires, as we have been suggesting, giving a cognitive account that captures *when skepticism is intuitive, how strong the intuitions are, in what circumstances, and to whom*, then supplementing or supplanting empirical data with refined philosophical judgments runs the risk of overlooking variations in our skeptical inclinations. Moreover, the very fact that philosophical traditions, pragmatist and New Sceptical, respond to these scenarios differently prompts questions about the wisdom of "drawing on philosopher's judgements." Whose judgements count? [As an aside: all these variations and disagreements are actually grist for the skeptic's mill—at least for a skeptic of a Pyrrhonian stripe.]



replying to the global skeptic. The simplicity of the skeptic's argument demands a correspondingly complex and detailed anti-skeptical reply with substantial predictive power, if the whole structure is to provide the kind of therapy that the New Sceptics point to as a desideratum of any anti-skeptical project. As Sections 2–4 of this paper suggest, this project is going to be enormously complex, and will almost certainly involve many theoretically contested propositions. Furthermore, as highlighted in Section 5, there is still enough variation in our skeptical inclinations that is left unexplained to feel confident that the phenomena of skepticism have been fully described, even by as thorough and inventive a reconstruction as has been developed by Pinillos here.

## **6. Conclusion**

In concluding our discussion of Pinillos' *Why We Doubt*, we want to stress that the book is a significant contribution to philosophical discussions of skepticism, particularly in its innovative and comprehensive approach to addressing global skepticism. While we have focused predominantly on the cognitive and Bayesian aspects of Pinillos' response, it is evident that the richness of his arguments invites further inquiry and development. His integration of the cognitive sciences with epistemology not only enriches our understanding of skeptical inclinations but also challenges us to consider more nuanced models and gather more empirical data to robustly tackle skepticism, even its more global forms. As we advance, it is crucial to build on Pinillos' insights, exploring the depth and breadth of skeptical challenges and responses, ensuring a dynamic and evolving conversation within the philosophical community and beyond. This engagement not only refines our theoretical perspectives but also enhances our practical understanding of knowledge, belief, and justification in the face of skepticism.

## References

- Alexander, J., Gonnerman, C., & Waterman, J. 2014. "Salience and Epistemic Egocentrism: An Empirical Study." In J. Beebe (ed.), *Advances in Experimental Epistemology*, 97–118. New York: Continuum.
- Buckwalter, W. 2010. "Knowledge Isn't Closed on Saturday: A Study in Ordinary Language," *Review of Philosophy and Psychology* 1 (3): 395–406.
- Buckwalter, W. 2014. "The Mystery of Stakes and Error in Ascriber Intuitions." In J. Beebe (ed.), *Advances in Experimental Epistemology*, 145–173. New York: Bloomsbury Academic.
- Clarke, T. 1972. "The Legacy of Skepticism," *The Journal of Philosophy* 69 (20): 754–769.
- Descartes, R. 1988 [1641]. *Meditations on First Philosophy*. In *Descartes: Selected philosophical writings*, 73–122. Translated by J. Cottingham, R. Stoothoff, & D. Murdoch. Cambridge: Cambridge University Press.
- Dinges, A. 2019. *Knowledge, Stakes, and Error: A Psychological Account*. Frankfurt am Main: Klostermann.
- Gerken, M. 2013. "Epistemic Focal Bias," *Australasian Journal of Philosophy* 91 (1): 41–61.
- Gerken, M. 2017. *On Folk Epistemology: How We Think and Talk about Knowledge*. Oxford: Oxford University Press.
- Gerken, M., Gonnerman, C., Alexander, J., & Waterman, J. P. 2020. "Salient Alternatives in Perspective," *Australasian Journal of Philosophy* 98 (4): 792–810.
- Hazlett, A. 2006. "Epistemic Conceptions of Begging the Question," *Erkenntnis* 65 (3): 343–363.
- Hume, D. 1975 [1739]. *A Treatise of Human Nature*. Edited by L. A. Selby-Bigge, revised by P. H. Nidditch. Oxford: Clarendon Press.

- May, J., Sinnott-Armstrong, W., Hull, J. G., & Zimmerman, A. 2010. "Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study," *Review of Philosophy and Psychology* 1 (2): 265–273.
- Moore, G. E. 1939. "Proof an External World," *Proceedings of the British Academy* 25 (5): 273–300.
- Musgrave, A. 1974. "Logical Versus Historical Theories of Confirmation," *The British Journal for the Philosophy of Science* 25 (1): 1–23.
- Nagel, J. 2010. "Knowledge Ascriptions and the Psychological Consequences of Thinking about Error," *Philosophical Quarterly* 60 (239): 286–306.
- Nagel, J., San Juan, V., & Mar, R. A. 2013. "Lay Denial of Knowledge for Justified True Beliefs," *Cognition* 129 (3): 652–661.
- Nozick, R. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Pinillos, Á. 2016. "Experiments on Contextualism and Interest Relative Invariantism." In J. Sytsma & W. Buckwalter (eds.), *A Companion to Experimental Philosophy*, 349–358. Malden, MA: Wiley Blackwell.
- Pinillos, N. Á. 2023. *Why We Doubt: A Cognitive Account of Our Skeptical Inclinations*. Oxford: Oxford University Press.
- Reid, T. 1983 [1785]. *Essays on the Intellectual Powers of Man*. In R. E. Beanblossom & K. Lehrer (eds.), *Reid: Inquiry and Essays*. Indianapolis, IN: Hackett.
- Russell, B. 1912. *Problems of Philosophy*. London: William & Norgate.
- Schaffer, J., & Knobe, J. 2010. "Contrastive Knowledge Surveyed," *Noûs* 46 (4): 675–708.
- Starmans, C., & Friedman, O. 2020. "Expert or Esoteric? Philosophers Attribute Knowledge Differently than All Other Academics," *Cognitive Science* 44 (7): e12850.

- Walton, D. N. 1991. *Begging the Question: Circular Reasoning as a Tactic of Argumentation*.  
New York: Greenwood.
- Waterman, J., Gonnerman, C., Yan, K., & Alexander, J. 2018. “Knowledge, Subjective Certainty,  
and Philosophical Skepticism: A Cross-Cultural Study.” In M. Mizumoto, S. Stich, & E.  
McCready (eds.), *Epistemology for the Rest of the World*, 187–214. Oxford: Oxford  
University Press.
- Williams, B. 2015. *Descartes: The Project of Pure Inquiry*. New York: Routledge.
- Williams, M. 1991. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*.  
Cambridge, MA: Blackwell.
- Williamson, T. 2005. “Contextualism, Subject-Sensitive Invariantism and Knowledge of  
Knowledge,” *The Philosophical Quarterly* 55 (219): 213–235.