The Insignificance of Philosophical Skepticism

(Penultimate Draft - Forthcoming in *Synthese* – Please cite published version)
Jonathan Dixon, dixonj@wfu.edu, Wake Forest University, Winston-Salem, NC, USA

Abstract: The Cartesian arguments for external world skepticism are usually considered to be significant for at least two reasons: they present genuine paradoxes and providing an adequate response to these arguments would reveal something epistemically important about knowledge, justification, and/or our epistemic position to the world. Using only premises and reasoning the skeptic accepts, I will show that the most common Cartesian argument for external world skepticism (i.e. the closure-based skeptical argument) leads to a previously unrecognized self-undermining dilemma: it either leads to a *reductio ad absurdum*, or to avoid this reductio the skeptic must accept that this argument is epistemically *idle* – it does not provide any support for external world skepticism. Either way, this Cartesian argument cannot legitimately threaten or even call into question our beliefs about the external world. And thus, either way, this Cartesian argument for external world skepticism is not epistemically significant – it is not a genuine paradox and adequately responding to it need not reveal anything epistemically important.

Key Words: Cartesian Skepticism, Closure of Knowledge, Brains in Vats, Replies to Skepticism

External world skepticism is the view that we do not or cannot know much, if anything, about the external world. The Cartesian arguments for external world skepticism appeal to global skeptical hypotheses (e.g. I am currently dreaming, a brain-in-a-vat, deceived by an Evil Demon, etc.) and claim that unless one can rule-out or neutralize such hypotheses we do not know much, if anything, about the external world.[1]

Prima facie, external world skepticism and these Cartesian arguments are absurd. It both seems that we know many things about the external world (cf. Lewis 1996: 549), and that appealing to such seemingly ridiculous hypotheses to establish external world skepticism is beyond the pale

[1] Alternatively, external world skepticism is the view that we are not or cannot be *justified* in believing much, if anything, about the external world, and the Cartesian arguments appeal to global skeptical hypotheses to show that we are not justified in believing much, if anything, about the external world. While I will not focus on this justification version of external world skepticism or the Cartesian arguments for it, I believe my arguments in this paper – against the Cartesian arguments for not having *knowledge* of the external world – also apply to and challenge these arguments.

(cf. Frances 2005: 568, Frances 2008: 225, Goldman (2007), Greco 2008: 111).[2] However, despite this position's initial implausibility, philosophers have been perennially occupied with external world skepticism in large part because the Cartesian *arguments* for this view appear to be genuine paradoxes: they have premises that appear to be obviously true, but which logically entail an implausible conclusion (e.g. Greco 2000: 2, Stroud 1984: 14, McGinn 1989: 7, Wright 1991: 89, Pritchard 2002: 217, Pritchard 2005b: 16, and Byrne 2004: 299-300). As such, these arguments are epistemically significant because they, allegedly, provide excellent support for external world skepticism which cannot easily be dismissed out of hand (cf. Williams 2001: 3-4; Vogel (1993), Greco 2000: chap 3).

Additionally, it has been thought that the paradoxical nature of these arguments also, allegedly, makes them epistemically significant in another respect: providing an adequate response to these arguments would reveal something important about knowledge, justification, and/or our epistemic position to the world (e.g. Bernecker and Dretske 2000: 301, Le Morvan 2011: 87-8, and Greco 2000: 2-3). It is largely for this reason that these arguments have held and continue to hold a central place in epistemology. Indeed, many recent advances in epistemology (e.g. reliabilism, sensitivity, safety, and contextualism about knowledge, and externalism about justification, to name a few) have largely resulted from scholars attempting to respond to these skeptical arguments (e.g. Pritchard 2005b: 7, Steup and Neta (2020), Comesaña and Klein (2019), and Brueckner (2014)). Furthermore, current and past theories of knowledge and justification are often categorized and judged to be (in)adequate by how well they respond to these skeptical arguments (e.g. Comesaña and Klein (2019), Pritchard (2002)). In short, however seemingly

[2] Bourget and Chalmers (2014) found that the vast majority of academic philosophers believe external world skeptics is false and Leiter (2020) found that external world skepticism was voted as the "most preposterous" philosophical thesis (see https://leiterreports.typepad.com/blog/2020/01/preposterous-philosophical-views-the-poll-results.html).

implausible external world skepticism may be, tilling the skeptical fields is justified because it has born much epistemic fruit.

However, despite these advances, it has also been thought that these, and other, theories of knowledge and justification ultimately fail to provide a satisfying response to the Cartesian arguments for external world skepticism (e.g. Stroud (1984), Williamson 2000: 27). To many, the aforementioned theories seem only to *beg the question* and thus unsatisfactorily respond to the skeptical paradox because they employ premises or reasoning that the skeptic would not accept (Pryor 2000: 517-520). But ideally, we non-skeptics would respond to the Cartesian arguments for external world skepticism by using *only premises and reasoning that the skeptic would accept* (Stroud 1984: 108, Fumerton 1990: 64, Kornblith 2021: 10-13). Such a response is ideal because it would both definitively show that the Cartesian arguments for external world skepticism are not genuine paradoxes and would do so without relying on controversial philosophical theses which likely beg the question against the skeptic.

Of course, defenders of the above theories of knowledge and justification believe their responses adequately demonstrate that these skeptical arguments are not genuine paradoxes but, instead, rest on false presuppositions. But many other epistemologists also believe that it is likely impossible (e.g. Stroud 1984: 1, 20-1, 33, Williamson 2000: 27, Nagel 1986: 73-74, and Pryor 2000: 517-20) and/or unnecessary (e.g. Nozick 1981: 197-8, DeRose 1995: 49, and Byrne 2004: 303) to provide a non-question begging refutation of the skeptic's position and arguments. Whether or not this is true, it seems that until the Cartesian arguments for external world skepticism are refuted on their own terms, they will continue to hold a central place in epistemology.

Using only premises and reasoning the skeptic accepts, I will show that the most common Cartesian argument for external world skepticism (i.e. the closure-based skeptical argument)[3] leads to a previously unrecognized self-undermining dilemma: it either leads to a *reductio ad absurdum,* or to avoid this reductio the skeptic must accept that this argument is epistemically *idle* – it does not provide any support for external world skepticism. Either way, this Cartesian argument cannot legitimately threaten or even call into question our beliefs about the external world. And thus, either way, this Cartesian argument for external world skepticism is not epistemically significant – it is not a genuine paradox and providing an adequate response to it need not reveal anything epistemically important.

1. Overview of Dilemma

Before delving into the nitty-gritty details of the closure-based skeptical argument, for clarity it will be helpful to provide a brief overview of how my self-undermining dilemma proceeds.

For the time being, I take it as a given that the debate over whether we know anything about the external world should follow the rules of a debate. And like any debate, the skeptic and the non-skeptic both have to provide arguments with plausible premises that logically lead to and provide sufficient support for their respective positions (see section 3 for a defense of these claims). More specifically, each side in this, or any, debate has the following *dialectical burden*:

> **Dialectical Burden**: to provide arguments for their respective positions that are sound / cogent and have reasoning / premises that put one in a position to know, or provide adequate justification to believe, the conclusions of these arguments.[4]

[3] "The most popular way of motivating radical scepticism in the contemporary literature is by appeal to so-called 'closure-based' sceptical arguments" (Schonbaumsfeld 2016: 7). I believe this sentiment is shared by most epistemologists who write on skepticism (e.g. Steup and Neta (2020), Comesaña and Klein (2019)).

[4] The dialectical burden to "provide arguments" need not require that anyone put forward or be a proponent of the view that an argument defends. This is important to note because, with the possible exception of Unger (1975), no one is an external world skeptic. The Cartesian arguments allegedly provide excellent support for external world skepticism independently of anyone raising or advocating for them. Thus, the

Furthermore, like any debate, a winner is declared when one side provides more persuasive arguments for their position and/or demonstrates that their opponents' arguments are not cogent via some counter-arguments (cf. Byrne 2004: 300). A tie or draw results if neither side provides a more cogent argument for their position or if neither side can provide an effective counter-argument.

In the debate over our knowledge of the external world, both sides can attempt to meet their dialectical burden in many ways. While a full account of how each side might accomplish this is outside the scope of this paper, the following examples are sufficient to illustrate how each side can meet their dialectical burden.[5] For example, some non-skeptics (i.e. reliabilists) argue that we have lots of knowledge about the external world because the world is roughly how we perceive it through our reliable sensory and cognitive capacities. And, the Cartesian skeptic argues we do not have much, if any, knowledge of the external world because our sensory experiences are not sufficient to rule-out or neutralize the possibility that some global skeptical hypothesis – whereby the world is not at all how we perceive it – is true.

My self-undermining dilemma proceeds in four steps. First, it assumes, for reductio, that the Cartesian skeptic's argument is sound. For example, let's say that the skeptic argues for the conclusion that 'I don't know that I have hands' on the grounds that, because I *could* be a handless brain-in-a-vat being fed a subjectively indistinguishable computer simulation of the world, my sensory experiences are not sufficient for me to know much, if anything, about the external world.

---

dialectical burden does not have to be met by any particular proponent of a side of this debate (i.e. the skeptic vs. the non-skeptic) but rather only needs to be met by the position or arguments for either side of this debate (i.e. external world skepticism vs. the denial of external world skepticism) (cf. Kelly (2005)). However, for ease of exposition, in this paper I will sometimes refer to the Cartesian *skeptic* as a convenient way to refer to external world skepticism and the Cartesian arguments for it.

[5] See Comesaña and Klein (2019), Pritchard (2002), and Greco (2008) for summaries of various other arguments for, and defenses of, our knowledge of the external world.

Second, my argument shows that this same reasoning "cuts both ways": *if* the skeptic is right that the skeptical hypothesis that one *could* be a brain-in-a-vat is enough to block me from knowing I have hands, then *by that same skeptical reasoning* any anti-skeptical hypothesis where I *could* know the world is as it appears to my reliable senses is likewise enough to block me from knowing that 'I don't know that I have hands.' As such, the skeptic's argument leads to a kind of meta-skepticism: I *don't know* that 'I don't know that I have hands.'

Third, my argument further shows how the reasoning the skeptic utilizes in their, allegedly, sound argument can be used to meet the dialectical burden of at least putting one in a position to know its conclusion, i.e. the skeptic's argument allows me to *know* that 'I don't know that I have hands.' Steps two and three together show that *if* the skeptic's argument meets its dialectical burden, then we reach the following contradiction: I *both* know and don't know that 'I don't know that I have hands.' Fourth, I show that the only way the skeptic can avoid this contradiction is by conceding that their argument does not meet its dialectical burden. Specifically, to avoid this reductio the skeptic must accept that the premises and conclusion of her argument cannot be known or adequately justified. But if the skeptic were to do this, she would be admitting that her argument does not provide adequate support for external world skepticism. Thus, either the skeptic's argument is refuted via *reductio ad absurdum* or *idle* because it does not provide any support for external world skepticism.

2. The Closure Argument

I will now delve into the nitty gritty details of the closure-based skeptical argument and show how it leads to my self-undermining dilemma.[6] This argument reasons that if I know

[6] I should note that "the" closure argument is a misnomer since there are many different ways of spelling out the details of this argument (e.g. see Pritchard (2015)). I do not have space to canvas all these formulations here and will only show how my self-undermining dilemma applies to what I take to be the

something about the external world, e.g. I have hands, then I also know that I am not a handless brain-in-a-vat being fed a subjectively indistinguishable computer simulation of the world (henceforth, BIV); but since I do not know I am not a BIV, it follows that I don't know that I have hands. This argument can be formalized as:

C1. If I know that I have hands, then I know that I am not a BIV. (Closure)
C2. I don't know that I am not a BIV. (Premise)
C3. Thus, I don't know that I have hands. (C1, C2 MT)

And because the skeptic can replace 'I have hands' with almost any proposition about the external world and 'I' could refer to any epistemic agent, this skeptical argument generalizes to impugn most, if not all, of our knowledge of the external world. And while the skeptic could also replace the 'BIV' hypothesis with other global skeptical hypotheses to reach her skeptical conclusion (e.g. I am currently dreaming, hallucinating, being deceived by an evil demon, etc.), I will follow most other commentators and use the BIV hypothesis to explicate this argument.

Step 1: Assume the Closure Argument is Sound

The closure argument is thought to provide excellent support for external world skepticism because it is logically valid and each of its premises seem to be true. Premise C1 is supported by the often-discussed closure principle which concerns the requirements for using competent deduction to extend one's knowledge of a proposition to knowing another proposition. While there are many different versions of this principle (see Kvanvig (2008) and Luper (2018)), this paper will utilize the following formulation:

> **Closure Principle**: For all agents S, and all propositions P, and all propositions Q, if S knows that P, and S knows that P entails Q, and S competently deduces Q from P, and S believes Q on the basis of this competent deduction from P, then S knows, or is in a position to know, that Q.

---

most common way of formulating this argument. But regardless of how this argument is formulated, the following sections should make it clear how my self-undermining dilemma can be applied to any version of this argument.

While the precise nature of the closure principle does not ultimately matter for my arguments below, it is advantageous to the skeptic to utilize this stringent version of the closure principle because it is thought by many that such versions of the closure principle are obviously true (e.g. Fumerton 2005: fn. 5, Pritchard 2015: 14). And while some deny (various versions of) the closure principle (e.g. most famously Dretske (1970) and Nozick (1981), see Vogel (1990b) and Luper (2018) for a discussion), I will assume for the sake of the skeptic's argument that both C1 and this version, or some formulation, of the closure principle is true.

Premise C2 also seems to be obviously true because it plausibly assumes that one's knowledge of the external world is predominately, if not exclusively, based on their sensory experiences. But given the description of the BIV hypothesis, if you were a BIV being fed a flawless and subjectively indistinguishable simulation of the world, then all your sensory experiences about the world would be, *by hypothesis*, misleading. As such, the skeptic reasons that one's sensory experiences do not provide adequate evidence to support my beliefs about the world over the BIV hypothesis. Thus, the skeptic holds that sensory experiences do not allow me, or anyone, to know that they are not a BIV, per C2. And while there is also some controversy about whether this premise is true or is adequately justified by the skeptic (e.g. see Greco 2008: 111, Pryor (2000), Vogel (1990a)), I will also assume for the sake of the skeptic's argument that C2 is true and grant that the above reasoning which supports C2 is sound.[7]

[7] There is also controversy about whether C2 can be motivated independently of the underdetermination-based skeptical argument. If this is the case then, the closure-based and underdetermination-based skeptical arguments are not independent of one another (e.g. see Brueckner (1994), Cohen (1998), Pritchard (2005a), McCain (2013)). In this paper, I will assume in the skeptic's favor that both skeptical arguments are independent of one another (see Prichard (2015, 2018) for a recent defense of this claim).

Granting C1, C2, and the reasoning which supports them is part of my self-undermining dilemma below: e*ven if* we grant the skeptic all this, the closure argument ultimately fails to support external world skepticism.

Step 2: Show How the Closure Argument "Cuts Both Ways"

With these assumptions, the reductio horn of my dilemma continues by showing that the closure argument "cuts both ways" – i.e. this argument is epistemically self-defeating. An argument is epistemically self-defeating iff at least one of this argument's components (i.e. premises, conclusions, or reasoning) epistemically defeats at least one such component.[8] For example, any argument that leads to "the premises of this very argument cannot be known and/or are unjustified" would be epistemically self-defeating. Additionally, any argument that leads to "the conclusion of this very argument cannot be known and/or is unjustified" would be epistemically self-defeating. These are the kinds of self-defeat that will be at issue in this paper (cf. Decker (2014)).[9]

Regarding the closure argument, I argue that the skeptic's own reasoning can be turned around and used to epistemically defeat its own *conclusion,* C3. In other words, the skeptic's own

[8] While Moore (1959) also argues that the dream skeptic's argument "cuts both ways," he means something different by this terminology. He argues that a rival non-skeptical argument is just as well motivated as the dream skeptic's argument. On one interpretation of Moore, this attempts to demonstrate a kind of stalemate with the skeptic. My argument is importantly different since it attempts to show that the skeptic's argument is self-incriminating and defeating – i.e. it can be turned on itself to epistemically defeat its own conclusion.

[9] One might object that an argument leading to the conclusion 'the premises of this very argument cannot be *known*' is both different and not as serious a defect as an argument leading to the conclusion 'the premises of this very argument cannot be *justified*.' To link these kinds of self-defeat in my definition of 'epistemic self-undermining arguments' seems to presuppose something like the following controversial epistemic view of inference: S can use a proposition p as a premise in inferential reasoning only if S knows that p (see Williamson 2000: 184-207; Moon 2013: 2728). A similar worry seems to befall linking arguments that lead to 'the *conclusion* of this very argument cannot be known vs. cannot be justified.' To link these kinds of self-defeat seems to presuppose that suspending belief in p is the proper reaction to discovering that one does not know p (cf. Decker 2014: 1102). In Step 4 of Section 2 (and fn. 21), I argue that my self-undermining dilemma does not make either of these assumptions. Thanks to an anonymous reviewer for raising this concern.

reasoning can be used to establish a kind of meta-skepticism: 'I don't know C3.' To see why, it is important to first note one feature about the closure argument: the closure principle which underlies C1 places *no restrictions* on the kinds of propositions that can be placed into it (i.e. "For all agents S, and *all* propositions P, and *all* propositions Q…"). This open-endedness is necessary for the closure argument to support external world skepticism because it is this feature that allows this argument to generalize to all, or nearly all, propositions about the external world. This open-endedness also allows the skeptic to appeal to any of the common skeptical hypotheses (e.g. BIV, dreaming, Evil Demon, etc.) to challenge our knowledge of the external world (cf. Pritchard 2018: 76). As such, the skeptic's own closure principle allows for their conclusion, C3, to be placed into it.

This leads to the question: *if* one knows C3 (i.e. I know that 'I don't know that I have hands'), then what other propositions can one extend their knowledge to via the closure principle? For my purposes, the salient proposition that we can extend our knowledge to via the closure principle is 'I don't satisfy the conditions for knowing that I have hands':

> If I know C3, and I know C3 entails that 'I don't satisfy the conditions for knowing that I have hands,' and I competently deduce this entailment and believe that 'I don't satisfy the conditions for knowing that I have hands' on the basis of knowing C3, then I know, or am in a position to know, that 'I don't satisfy the conditions for knowing that I have hands.'

This allows for the following epistemic self-defeating challenge to the conclusion of the skeptic's argument, C3:

> C4. If I know C3, then I know that I don't satisfy the conditions for knowing that I have hands. (Closure)
> C5. I don't know that I don't satisfy the conditions for knowing that I have hands. (Premise)
> C6. Thus, I don't know C3. (C4, C5 MT)

For the rest of this subsection, I will further explicate and defend this argument from objections.

The skeptic cannot deny C4 because, as was just explained, the closure principle places no restrictions on the kind of propositions that can be placed into it, and so the skeptic cannot object to placing C3 into the closure principle. However, one might object that C4 does not satisfy the clause in the closure principle that 'S knows P entails Q' because C3 does not entail 'I don't satisfy the conditions for knowing that I have hands.' This objection holds that it is possible for me to not know that I have hands (C3), while it being the case that I *do* satisfy the conditions for knowing that I have hands. This objection can be blocked if we plausibly assume the following biconditional: S knows X iff S satisfies the conditions for knowing X, e.g. I know that I have hands just in case I satisfy the conditions for knowing that I have hands. As such, this biconditional makes C3 (I don't know that I have hands) entail 'I don't satisfy the conditions for knowing that I have hands.'[10] Importantly, this biconditional is not committed to any particular account of knowledge and, like the skeptic, I am not committed to any particular account of knowledge. Rather, in this paper I am following the skeptic in holding that whatever the correct account of knowledge may be, the skeptical closure argument C1-C3 shows that the conditions, whatever they may be, for me to know that I have hands are not satisfied – and thus, I don't know that I have hands.

Furthermore, the skeptic cannot deny C5 because the same considerations that the skeptic provides to support C2 (I don't know that I am not a BIV) explained above also supports C5. Specifically, C5 is established by parity of reasoning with the skeptic's argument for C2, i.e. the argument for C5 is similar enough in structure and content to the argument for C2 that both arguments are logically on a par with one another. As such, the skeptic cannot object to C5 without also objecting to the reasoning they provide for their own premise, C2.

[10] I should note that blocking this objection only requires the claim that: If S knows X then S satisfies the conditions for knowing X. I assume this biconditional because I believe it is true and do not think that taking on this additional commitment will likely be challenged since this is widely accepted claim.

To see this, it is important to recall how the skeptic argues for C2. Everyone, including the skeptic, will grant that it *appears* to me that I have hands on the basis of my sensory experiences. Indeed, the skeptic plausibly assumes that one's knowledge of the external world is predominantly, if not exclusively, based on their sensory experiences. But the skeptic argues that because the world would appear exactly the same to me whether or not I was a BIV, my sensory experiences are unable to provide adequate evidence to support my belief that I have hands over the BIV hypothesis. From this, the skeptic reasons that 'I don't know that I am not a BIV,' i.e. C2.

There are two important things to note about the skeptic's argument for C2. First, this argument is essentially negative. It does not provide any positive evidence for or positive reason to believe the BIV hypothesis is true (e.g. the ticker tape at the bottom of my visual field says "you are a BIV, deal with it"). Instead, it relies on our senses being incapable of differentiating between the real world and skeptical counter-*possibilities* like the BIV hypothesis. And the skeptic need not provide any such evidence for or even believe that the BIV hypothesis is true. In true skeptical fashion, the skeptic need not have any views on the matter. After all, their overall argument C1-C3 just aims to show that we don't have knowledge of the external world, not to show what the world is really like.

Second, because the skeptic's argument for C2 is negative in these ways, the skeptic grants that the world *could* be as it appears to us. For the skeptic, neither the BIV hypothesis nor the hypothesis that the world is roughly how we perceive it through our reliable senses (henceforth, 'the real-world hypothesis') are adequately supported by one's sensory evidence. Since the world would appear exactly the same to me if either hypothesis were true, both hypotheses *could* be the case and our sensory experiences do not favor one over the other. In other words, the skeptic holds that these hypotheses are epistemically *on a par*. Indeed, for all the skeptic has shown about the

external world, the BIV hypothesis *could* be false, the world *could* be roughly how we perceive it, and our senses *could* be a reliable guide to knowing the world. As such, the real-world hypothesis *could* be true, and I *could have hands*. Additionally, it seems that the skeptic also *must* grant the possibility that the real-world hypothesis is true. If they were to deny the possibility that the world is actually as it appears and that I have hands in their support for C2, they would *beg the question* against the non-skeptic, i.e. since knowledge is factive, denying the possibility that I have hands to support C2 just assumes that the conclusion C3 (i.e. I don't know that I have hands) is true. Thus, for the skeptic the BIV hypothesis and the real-world hypothesis are epistemically *on a par*.

But, if the skeptic's negative reasoning for C2 is permissive enough to allow for both the BIV and real-world hypotheses, the skeptic's own negative reasoning is also presumably permissive enough to allow for *any* hypothesis about the world that is likewise consistent with how it appears to us. However, not only does it appear to me that I have hands, it also *appears* to me that 'I satisfy the conditions for knowing that I have hands.' That is, it appears to me that my sense experiences do provide sufficient support to know that I have hands. Indeed, everything would look exactly the same if I did or did not satisfy the conditions for knowing that I have hands. As such, it seems that the skeptic's negative reasoning for C2 also supports the following possibility: I *could* satisfy the conditions for knowing that I have hands. Of course, the skeptic's negative reasoning supports the possibility that 'I could not satisfy the conditions for knowing that I have hands.' Indeed, the skeptic's negative reasoning for C2 above allows the skeptic to conclude 'I don't know that I satisfy the conditions for knowing that I have hands.' This is because if the BIV hypothesis were true, then I would not satisfy the conditions for knowing I have hands (since I would be handless), and hence I cannot know that I satisfy these conditions. In short, the BIV hypothesis rules-out that I satisfy these conditions.

However, there are hypothesis that do entail that 'I satisfy the conditions for knowing that I have hands.' These hypotheses are consistent with how the world appears to us (and thus epistemically on a par with the real-world and BIV hypotheses) but which, contra the skeptic, rule-out that 'I don't satisfy the conditions for knowing that I have hands.' Call these 'anti-skeptical hypotheses.' While there are many possible anti-skeptical hypotheses, this paper will focus on the following example because it is the BIV hypothesis which concerns the skeptic's closure argument:

> *The Anti-BIV Hypothesis*: For whatever reason, the computer program, which underpins the usually flawless simulation for those who are BIVs, is incompatible with my neural network. This incompatibility is such that if I were a BIV, everything I perceived in the simulated world would indicate that it was a simulated object (e.g. have "BIV Simulated Object" written on it with big red letters). Consequently, this simulation cannot deceive me into believing I was in the real world when I was actually a BIV. In other words, the structure of my brain imbues me with a perfectly reliable ability to discern whether I am a BIV or in the real world.

Now, by the skeptic's standard, in this possibility my sensory experiences do provide adequate evidence to believe that I have hands over the BIV hypothesis precisely because, in this possibility, I have the ability to rule out the BIV hypothesis. That is, if I were an Anti-BIV, then when I look at my hands and they do not have "BIV Simulated Object" written across them, I would come to truly believe that I have hands on the basis of this ability. In short, unlike the BIV hypothesis, if I were an Anti-BIV, all my sensory experiences about the world would, *by hypothesis*, *not* be misleading.

Furthermore, a plausible case can be made that, in this scenario, I would satisfy the conditions for knowing that I have hands: true belief with perfectly reliable justification. I say "plausible case" because I, like the skeptic, do not want to assume any specific account of knowledge. I and the skeptic just hold that one needs to be able to rule out qualitatively indistinguishable counter-possibilities, like the BIV hypothesis, in order to know things about the

world. That said, it seems that we can add details to the Anti-BIV hypothesis above for it to satisfy any reasonable conditions for knowing that I have hands.[11] In short, the above case can be amended to make the following true: if the Anti-BIV hypothesis were true, then I would satisfy the conditions for knowing I have hands. And if this were true then, because knowledge is factive, 'I don't know that I don't satisfy the conditions for knowing that I have hands,' i.e. C5, would also follow.

Fortunately, the details of how the Anti-BIV hypothesis should be spelled out to allow me to satisfy the conditions for knowing that I have hands is of no consequence to my argument for C5. Indeed, my argument for C5 does not rely on the truth or cogency of the Anti-BIV hypothesis. Instead, in the same way the skeptic's negative reasoning for C2 uses the BIV hypothesis to challenge the real-world hypothesis, my parity of reasoning argument for C5 uses the BIV hypothesis to challenge the Anti-BIV hypothesis. This negative reasoning dictates that because the world would look exactly the same if either the Anti-BIV or BIV hypothesis were true, my sensory experiences do not provide good reason to believe the possibility that I am an Anti-BIV and 'I satisfy the conditions for knowing that I have hands' over the possibility that I am a BIV and 'I do not satisfy the conditions for knowing that I have hands.' So even if I am an Anti-BIV, I still don't know that I am not a BIV, per C2.

[11] As the case is currently described, it allows for fallible knowledge in-line with many prominent modal accounts of knowledge. For example, if I were an Anti-BIV, my belief that I have hands is sensitive (if I did not have hands, then I would not believe it) and safe (I would believe that I had hands only if I had hands). But, the case can be spelled-out to even allow for certain kinds of infallible knowledge (see Dutant (2007, 2016)) while remaining compatible with the how the world appears to us, e.g. an Evil Demon or BIV Scientist could imbue me with an infallible power to discern whether I am being deceived by an illusion or am experiencing the real world but, from my first person perspective, it appears to me that I am forming my belief based on my sensory and fallible sense perceptions when I am using my infallible power. If this were the case, then I would presumably satisfy the conditions for infallibly knowing that I have hands (e.g. this power makes it that my belief that I have hands cannot be wrong (cf. Dutant 2007: 73) or epistemically unlike any false belief (cf. Dutant (2016))). And given that the closure principle has no restrictions on the kind of hypotheses that can be plugged into it, the skeptic cannot object to this and other anti-skeptical possibilities.

But crucially, the same reasoning "cuts both ways." Because these hypotheses are epistemically on a par, my sensory experiences also do not provide adequate reason to believe the possibility that I am a BIV and 'I don't satisfy the conditions for knowing that I have hands' over the possibility that I am an Anti-BIV and 'I do satisfy the conditions for knowing that I have hands.' So even if I am a BIV, the claim that 'I don't know that I am not an Anti-BIV' follows, by parity of reasoning, from the skeptic's negative reasoning for C2. And since being an Anti-BIV (however one spells out the details) entails that 'I do satisfy the conditions for knowledge,' the claim that 'I don't know that I am not an Anti-BIV' is equivalent to 'I don't know that I don't satisfy the conditions for knowing that I have hands,' i.e. C5. Thus, C5 is established by parity of reasoning with the skeptic's argument for C2. In sum, my argument strategy for C5 is to show that if the Anti-BIV hypothesis is true, then C5 follows; and if we challenge the Anti-BIV hypothesis with the skeptic's negative reasoning for C2, then C5 follows via parity of this negative reasoning.

One might object that because the skeptic's conclusion C3 establishes that 'I don't know that I have hands' it follows that 'I don't satisfy the conditions for knowing that I have hands' and thus it also follows that 'I don't know that I satisfy the conditions for knowing that I have hands.' Indeed, the skeptic's closure argument seems to establish that it is *impossible* for me to satisfy the conditions for knowing that I have hands. As such, the skeptic's closure argument and conclusion C3 rules out anti-skeptical hypotheses like the Anti-BIV hypothesis, and so my parity of reasoning argument for C5 does not work.

This objection fails to understand the dialectic thus far. While I do grant that the skeptic's closure argument is sound, for reductio, the skeptic's conclusion, C3, is not relevant to the reasoning used to support C2. Indeed, the skeptic's conclusion cannot be used to support one of its premises since this would beg the question against the non-skeptic, i.e. it would be assuming

and using C3 in an argument for C3.[12] Recall that the skeptic must grant the possibility that the real-world hypothesis is true to avoid begging the question against the non-skeptic. Similarly, the skeptic *must* allow for the possibility that I *could* satisfy the conditions for knowing that I have hands. Since knowledge is factive, if they were to deny that I could satisfy these conditions in their support for C2, they would just be assuming that C3 (I don't know that I have hands) is true and thus beg the question against the non-skeptic.

Additionally, what my parity of reasoning argument shows is that the skeptic's closure argument does not and cannot establish that it is *impossible* for me to know that I have hands. To accomplish this, it seems the skeptic would have to show something about the world (e.g. show I don't have hands or show that I am permanently incapable of being an Anti-BIV) which, given the skeptic only provides a negative argument for C2, is something neither the skeptic nor their closure argument attempts to do. At best, the skeptic's argument establishes that, *given the appearance* that I can only use my fallible senses to learn about the world, I do not know that I have hands because my senses cannot rule out the BIV hypothesis. But this result also leaves room for the possibility that I *could* satisfy the conditions for knowing that I have hands via an anti-skeptical hypothesis like the Anti-BIV hypothesis. Thus, the skeptic's negative reasoning allows for the possibility that anti-skeptical hypotheses, like the Anti-BIV hypothesis, are true, and leaves room for my parity of reasoning argument for C5. In short, it is precisely because the skeptic's reasoning for C2 dictates that the BIV and Anti-BIV hypotheses are epistemically on a par that the skeptic's own reasoning cuts both ways.[13]

---

[12] However, as will be shown in next the next section, this objection points to a reductio argument against the closure argument. I argue that *if* the skeptic's reasoning *also* allows me to satisfy the conditions for, and hence, know that I don't know that I have hands (via the above biconditional), then the skeptic's own reasoning leads to a contradiction: I both know and don't know that I don't know that I have hands.

[13] I would like to thank my anonymous reviewers for pressing me to make the arguments in this section clearer.

In sum, this subsection argued that the skeptic cannot deny C4 because it relies on the closure principle which underlies C1 of her own argument and the skeptic cannot deny C5 because it is supported by the same considerations that support C2 of her own argument. Thus, the closure argument itself cuts both ways and epistemically defeats its own conclusion, C3.

Step 3: Show How the Closure Argument Meets its Dialectical Burden

One might object that, *even if* the closure argument cuts both ways, this is not a very worrisome result for the Cartesian skeptic. At best this shows that the skeptic's own argument leads to a kind of meta-skepticism: not only can we not know very much, if anything, about the external world, but we cannot even know that we don't know very much, if anything, about the external world. But it is not clear how challenging the knowability of the closure argument's conclusion, C3, contests the validity or soundness of the original closure argument C1-C3. Indeed, the skeptic herself may even accept this meta-skeptical result as bolstering her own argument and skeptical position because it further shows the true depths of our ignorance.

I agree that this meta-skepticism does not *by itself* challenge the closure argument. However, I argue that unless this meta-skepticism can be blocked, the closure argument is refuted by an additional argument. This additional argument shows how the skeptic's own reasoning allows their closure argument C1-C3 to meet its dialectical burden: this, allegedly, sound argument puts one in a position to know, or provides adequate justification for, its conclusion. In other words, if one learns of this sound argument, is convinced the premises are true, and makes the appropriate valid inference to the conclusion, then one is justified in believing or is in a position to know its conclusion, C3. But this result contradicts the meta-skepticism established in the previous subsection – I don't know C3. As such, this shows that the skeptic's own reasoning leads to a

contradiction: I *both* know and don't know C3. For clarity, we can formally represent how the closure argument meets its dialectical burden as:

> C7. If I know C1 and C2, then I know C3. (Closure)
> C8. I know C1 and C2. (Premise)
> C9. Thus, I know C3. (C7, C8 MP)

For the rest of this subsection, I will further explicate and defend this argument from objections. In the next subsection, I will discuss how the skeptic may respond to this reductio.

While there are many plausible principles that can underlie C7, in this paper I will hold that this premise is supported by the above closure principle. In expanded form, C7 says:

> If I know C1 and C2, and I know C1 and C2 entail C3 and I competently deduce this entailment and believe C3 on the basis of knowing C1 and C2, then I know, or am in a position to know, C3.

This allows me to avoid relying on other principles in my reductio argument that the skeptic may deny. Also, because C1 of the skeptic's closure-based skeptical argument depends on the closure principle, if the closure principle supports C7, then the skeptic cannot deny C7 without also undermining their own argument – and thus fail to meet their dialectical burden. Likewise, the skeptic cannot deny C8 without also failing to meet their dialectical burden of putting one in a position to know, or providing adequate justification for, its conclusion. And while C8 can be challenged in a number of ways (see Step 4 below), I will also grant for the time being that the skeptic has met their dialectical burden and assume that C8 is true.

One might object that, strictly speaking, C7 does not utilize the above closure principle. Since C7's consequent is inferred from *two* propositions (C1 and C2) rather than just *one* proposition (I know that I have hands), C7 does not utilize the above "single-premise" closure principle but instead utilizes a "multi-premise" closure principle. As such, this might provide the

skeptic a way to deny C7 while accepting C1.[14] This objection fails for three reasons. First, by this standard, it is not clear that C1 doesn't also use a multi-premise closure principle. After all, the proposition 'I have hands' is not an atomic proposition but is comprised of several different propositions which only together make 'I have hands' incompatible with 'I am a BIV.' For instance, if 'have hands' necessarily includes having fingers, thumbs, palms, etc., then 'I have hands' assumes the conjunction of 'I have fingers,' 'I have thumbs,' 'I have palms,' etc. As such, the skeptic needs to explain why the conjunction which underlies 'I have hands' doesn't necessitate the use of a multi-premise closure principle while the conjunction of C1 and C2 does necessitate the use of a multi-premise closure principle.

Second and relatedly, this point about individuating propositions also applies in the other direction: by this standard, is not clear why the conjunction of C1 and C2 ('If I know that I have hands, then I know I am not a BIV' and 'I don't know that I am not a BIV') cannot be rewritten as a *single* proposition to be plugged into the above "single-premise" closure principle. That is, if 'I have fingers,' 'I have thumbs,' etc. can legitimately be written as the single proposition 'I have hands,' then what prevents C1 and C2 from also being written as a single proposition? For example, this conjunction might be rewritten as: 'the consequent of C1 is false' – i.e. the consequent of "If I know that I have hands, then I know that I am not a BIV" is false. This reformulation is significant because the *truth* of this single proposition entails C3 (I don't know that I have hands), and so from this single proposition the closure principle can be used to extend

[14] I should note that, while this and similar versions of the closure principle are commonly invoked when explicating the skeptic's closure argument, it is not clear that they really are "single-premise" closure principles. For instance, if we hold that the inferential statements of the form 'P entails Q' are propositions, then the closure principle above requires an agent to know two propositions (to know that 'P' and know that 'P entails Q') in order to deduce another proposition (Q). However, I will grant for the sake argument that the skeptic intends to use a single-premise closure principle and just assume for ease of exposition that inferential statements that instantiate 'P entails Q' are not propositions (cf. Luper 2018).

our knowledge to C3 in the same way that C7 does.[15] Thus, contra the above objection to C7, this premise and argument above need not require a "multi-premise" closure principle. As such, to make their objection stick, the skeptic will also need to explain why C1 and C2 cannot be written as a single proposition and explain why 'I have hands' can legitimately be written as a single proposition.

Lastly, *even if* C7 utilizes a multi-premise closure principle while C1 does not, the skeptic will have to accept something like a multi-premise closure principle to meet their dialectical burden of putting one in a position to know, or providing adequate justification for, their argument's skeptical conclusion. (Or, at the very least, the skeptic can *ad hocly* deny a multi-premise closure principle is generally true but accept that it holds true only for their skeptical argument.) Indeed, only a complete skeptic about our ability to gain knowledge from arguments (i.e. via reasoning and inference) would deny that we are able to gain knowledge in the way C7 suggests (i.e. by coming to know a(n), allegedly, true conclusion by competently following and being convinced by a(n), allegedly, flawless argument). And since the external world skeptic only aims to challenge our knowledge of the external world and not our ability to reason in this way, the skeptic cannot deny C7 without also broadening their skepticism to impugn our general ability to gain knowledge from arguments. And if the skeptic were to do this, it would raise questions about the coherence of providing arguments to prove external world skepticism when their own argument dictates that no argument can accomplish this.

Step 4: The Closure Argument is Refuted or Idle

[15] In case this is not readily apparent, here is the full rendering of this application of the closure principle: If I know that 'the consequent of C1 is false,' and I know that 'the consequent of C1 is false' entails C3 (I don't know that I have hands), and I competently deduce C3 from 'the consequent of C1 is false', and I believe C3 on the basis of this competent deduction from 'the consequent of C1 is false,' then I know, or am in a position to know, that C3.

With steps 1-3, the Cartesian skeptic's closure-based skeptical argument succumbs to the first horn of my self-undermining dilemma: this argument is refuted via a *reductio ad absurdum* since the reasoning this argument utilizes leads to a contradiction between C6 and C9: I *both* know and don't know C3 – i.e. I both know and don't know that I don't know that I have hands. This reductio argument rests on four premises: C4, C5, C7, and C8. As was shown above, the skeptic cannot deny C4, C5, or C7 since these premises are supported by the very same considerations that support the skeptic's own closure argument. This only leaves premise C8 (i.e. 'I know C1 and C2') for the skeptic to deny. If the skeptic does not deny this premise, then at least one of the skeptic's original premises C1, C2, and/or the closure principle is false. Therefore, the skeptic must deny C8 to block this reductio and hold that we don't know at least one of the premises of their argument, C1 or C2.

But the skeptic cannot do this without succumbing to the second horn of my self-undermining dilemma: the closure-based skeptical argument is epistemically *idle*, i.e. this argument does not meet its dialectical burden of putting one in a position to know, or providing adequate justification for, external world skepticism. To see this, recall that the debate over whether we know anything about the external world should follow the rules of a debate. Both the skeptic and the non-skeptic have to provide arguments with plausible premises which logically lead to and provide sufficient support for their respective positions. The skeptic's closure argument fails to meet this dialectical burden if the skeptic denies C8 and admits that her own argument dictates that we cannot know at least one of its premises, C1 or C2. Indeed, by failing to meet its dialectical burden, the skeptic's own reasoning dictates that this argument is not cogent support for external world skepticism. To illustrate, if the skeptic denies C8, then we non-skeptics are entitled to ask: why is C8 false? Why can't we know the premises C1 and C2 of this skeptical

argument? Is the fault with us, or are your premises defective? Any answers to these questions the skeptic provides will be an admission that, by their own reasoning, their closure-based skeptical argument does not provide adequate justification for, or puts one in a position to know, its conclusion.[16]

There are at least three different ways to contest this horn of my self-undermining dilemma. First, the skeptic could *deny* C8 but argue that, contrary to my claims in the previous paragraph, their argument is not epistemically idle because it does meet its dialectical burden. Second, they could *accept* C8 and argue that the resulting contradiction is not a problem for their skepticism. Lastly, the skeptic could altogether deny the need for their argument to meet a dialectical burden. I will elaborate and respond to the first two objections below and elaborate and respond to the third objection in the next section.[17]

Regarding the first objection, the skeptic might claim that their closure argument is not idle since they can deny C8 while at the same time meeting their dialectical burden. They could argue that, as stated, the dialectical burden requirement runs together two epistemic notions that should be kept separate: 'putting one in a position to know'[18] vs. 'providing adequate justification' for external world skepticism. Given that 'position to know' is factive while 'justification' is not, the skeptic can do the latter without doing the former. Indeed, the skeptic only claims to show that

[16] Furthermore, it is an interesting fact that the skeptic is forced to deny C8 *if* C6 and C7 are true because together C6 and C7 entail ~C8. In this and the following section, I defend C8 and the dialectical burden requirement from objections.

[17] Thanks to an anonymous reviewer for raising these objections.

[18] This paper does not assume any specific account of 'position to know' but relies on how it is commonly used and understood in English. That said, I am tempted to say that at a minimum a deductive argument puts S in a position to know its conclusion C only if there is at least one world where the argument is not just valid but sound, and if S were to learn of this argument in this world, S could make the appropriate inference(s) in this argument to come to know C. A defense of this claim is outside the scope of this paper. See Yli-Vakkuri and Hawthorne (2022) and reference therein for a discussion of 'position to know' and the difficulties in precisifying this notion.

external world skepticism is true, not to show that they *know* that external world skepticism is true. To show that p is one thing, to show that we know that p is another. As such, the skeptic need not know the premises of their argument, per C8, or their conclusion to meet their dialectical burden. Instead, their premises can provide adequate justification for, without allowing one to know, their conclusion. Furthermore, connecting these two notions seems to assume the following controversial epistemic view of inference: S can use a proposition p as a premise in inferential reasoning only if S knows that p (see Williamson 2000: 184-207; Moon 2013: 2728). Alternatively, connecting these two notions might assume that suspending belief in p is the proper reaction to discovering that one does not know p (cf. Decker 2014: 1102). The skeptic is not committed to, and thus can deny, these assumptions. Therefore, the skeptic's argument is not necessarily idle if they deny C8, since they can maintain that their argument does provide adequate justification for, just not knowledge of, external world skepticism.

This objection ultimately fails because of the structure of my self-undermining dilemma. To see why, recall that to avoid the reductio horn of my dilemma, the skeptic has to deny C8. In other words, to avoid a contradiction, the skeptic's own reasoning forces them to accept that at least one of their premises can never be known (i.e. deny C8). Additionally, C4-C6 also establishes that the skeptic's own reasoning is epistemically self-defeating, i.e. this reasoning entails that the conclusion of their own argument, C3, cannot be known. But plausibly, if it is impossible to know some proposition, then one can never be in a position to know this proposition (cf. Rosenkranz 2018: 317; Yli-Vakkuri and Hawthorne 2022: 1325). Thus, the premises of their argument can *never* put one in a position to know that external world skepticism is true. But if the skeptic's own reasoning prohibits this, then these premises also do not provide adequate justification for external world skepticism.

In other words, this horn of my self-undermining dilemma shows that the skeptic violates something like the following epistemic principle concerning how 'position to know' and 'adequate justification' are related: In deductive arguments, if S's premises / reasoning provides adequate justification for conclusion C, then there must be a possible world in which these premises / reasoning put S in a position to know C.[19] But given that the skeptic's own reasoning forces them to deny C8 (i.e. I know C1 and C2) in order to avoid a contradiction, and their own reasoning also dictates (via C4-C6) that they cannot know external world skepticism is true, there is no possible world in which the skeptic's own reasoning puts them in a position to know external world skepticism. Indeed, the skeptic's own reasoning prohibits such a world. Thus, via the above principle, the skeptic's own reasoning also does not provide adequate justification to believe external world skepticism.[20] Or, as I put it above, the skeptic's own reasoning forces them to admit that their closure argument is idle because it does not meet its dialectical burden.

As such, my self-undermining dilemma does not presuppose either of the claims mentioned above from Williamson, Moon, or Decker which respectively say that it is impermissible to use p in one's inferential reasoning or justifiably believe p if one does not know p. Rather, my dilemma relies on a different and much weaker logical principle that makes use of a *special case* where 'position to know' and 'adequate justification' are logically related: when one's deductive

[19] 'In deductive arguments' is needed to avoid Gettier-style counterexamples, e.g. there is no possible world in which a stopped clock puts you in a position to know the time (even if it is displaying the correct time), but the stopped clock still might provide adequate justification to believe the time it displays. Limiting the scope of this principle avoids such counterexamples since the premises of a deductive argument logically entail their conclusion while Gettier-cases, like the stopped clock case, only provide inductive support for their conclusions. And given that the skeptic's reasoning concerns the deductive closure argument, this clause does not beg the question against the skeptic.

[20] See Rosenkranz (2018) for a discussion of similar principles that link propositional justification and 'being in a position to know.'

argument can *never* put one in a position to know its conclusion, then this argument does not provide adequate justification for its conclusion.[21]

Regarding the second objection, the Cartesian skeptic could both accept C8 and accept that her closure argument leads to a contradiction but deny that this is a problematic result. For instance, the skeptic could argue that this is a beneficial result for their skepticism since it just reveals some inconsistency within our conceptions of knowledge or truth. In response, while other kinds of skepticism may coherently be able to take this route and accept the above contradiction (see fn. 28), this is not a route the Cartesian skeptic can take. As mentioned above, the Cartesian skeptic is commonly interpreted as only providing an argument which intends to demonstrate that we don't have knowledge of the external world (see Comesaña and Klein (2019)). Indeed, as also mentioned above, the skeptic's closure argument does not assume any particular account of knowledge but argues that – whatever the correct account of knowledge may be – we do not know that we have hands. But, if the skeptic were to accept that their argument reveals this kind of deep inconsistency, this would both considerably broaden their skepticism to impugn our general conceptions of knowledge and truth, and undercut the aim of their argument. This is because an argument cannot demonstrate p (i.e. we lack knowledge of the external world) while at the same time showing an

[21] In response to fn. 9 specifically, this also provides a special case where the self-defeat involved in 'the premises / conclusion of this very argument cannot be known' leads to the kind of self-defeat involved in 'the conclusion of this very argument cannot be justified.' If a deductive argument is self-undermining in a way that makes it impossible to know its conclusion, then this argument also epistemically self-undermines its own justification for its conclusion. As such, this also shows that my account of epistemic self-defeating arguments does not assume the claims from Williamson, Moon, or Decker. However, I should note that I am not opposed to the claims from Williamson, Moon, or Decker. Indeed, I believe there is some plausibility to these claims. That said, I believe my self-undermining dilemma is a more convincing and powerful response to the closure argument if it does not rely on such controversial claims.

inconsistency in this demonstration or an inconsistency in something this demonstration presupposes. Thus, one would no longer be a Cartesian skeptic if they took this route.[22]

3. The Significance and Limits of the Above Dilemma

In this paper, I have shown that the most common Cartesian argument for external world skepticism – the closure-based skeptical argument – leads to a self-undermining dilemma. This is a significant result because this demonstrates that this prominent Cartesian argument for external world skepticism is not significant in either of the ways mentioned in the introduction: because the reasoning which it utilizes is *self*-undermining, it is not a genuine paradox, and it can be refuted without revealing anything epistemically important about knowledge, justification, and/or our epistemic position to the world. Additionally, because my dilemma demonstrates that the above Cartesian argument is *self*-undermining, this meets the ideal standard of showing how the skeptic's argument can be non-question beggingly refuted by relying only on premises that the skeptic accepts. Indeed, it is a feature, and not a bug, of my dilemma that it refutes the closure argument while not appealing to any controversial epistemological theses (e.g. reliabilism, sensitivity, safety, and contextualism about knowledge, and externalism about justification, to name a few).[23]

[22] However, in response to this, the Cartesian skeptic could accept some paraconsistent logic like dialethism. This would allow the Cartesian skeptic to coherently accept that her closure argument leads to a contradiction while also denying that this is a problematic result. Indeed, depending on the details, this may even allow the Cartesian skeptic to show that their argument does demonstrate that we lack knowledge of the external world, despite this argument's reasoning leading to a contradiction. While this is an interesting possibility, I will have to set aside such concerns since they are outside the scope of this paper. However, I should note that I am not aware of anyone who argues that the Cartesian skeptic does or should adopt a paraconsistent logic, let alone dialethism. So, if my self-undermining dilemma leads to the skeptic to adopt an unorthodox logic, then I take this to be a significant result and one that further contributes to the literature on Cartesian skepticism. Thanks to an anonymous reviewer for raising these objections and showing me the need to respond to them.

[23] As Lycan (2019: 53, 77) confidently puts it: "on *any* philosophical topic, the person who propounds an analysis is going to get creamed."

But, unlike other responses to Cartesian skepticism, my self-undermining dilemma does not attempt to deny a specific premise of the skeptic's argument, nor does it provide an explanation as to why the argument seems compelling. Rather, my self-undermining dilemma only demonstrates that the *combination* of the skeptic's premises, and the reasoning which supports them, cannot be used to threaten our knowledge of the external world. Where this skeptical argument fails, my dilemma does not say (cf. Rinard 2018: 242).[24]

And it is for these reasons that some will be dissatisfied with my self-undermining dilemma and this general approach to responding to external world skepticism (e.g. Fumerton 1995: 43-53, Cohen 1998: fn. 26, Greco 2000: 61-64). Some believe that an adequate response to external world skepticism must provide some diagnosis as to what is mistaken with either external world skepticism itself or point to some flaw in the arguments for it (e.g. Williams (1996), Greco 2000: 61, cf. Pritchard 2015: 16, 24, Lycan 2019: 27, 32). This thought should be resisted. If my self-undermining dilemma is sound, it shows that the skeptic's closure argument is not a genuine paradox and that epistemologists have unknowingly been chasing their tails in attempting to respond to it. Furthermore, my dilemma does provide some explanation of why the closure argument seems compelling: it is not obvious that it is self-undermining and epistemologists have thus far failed to see this.[25]

[24] Rinard and others (e.g. Wright (1991), Wilson (2012), and Lai (2020)) have also attempted to show that the Cartesian skeptic's argument is self-undermining. While there are many similarities between our different ways of attempting to accomplish this, a full explanation of this is beyond the scope of this paper.
[25] However, I should note that other epistemologists have gestured at something like my self-undermining dilemma but have not found it to be a worthwhile response to the arguments for external world skepticism. For example, Greco (2000: 62) writes that "this kind of response is dismissive because it does not engage any particular skeptical argument…and cannot be sustained on even a superficial consideration of actual skeptical arguments." On the contrary, my dilemma is not dismissive since it engages directly with the reasoning of the closure argument and poses a significant challenge to this prominent skeptical argument.

To be clear, with these remarks I am not disparaging the many other responses to Cartesian skepticism or the epistemic fruit they have born. I only mean to express that all this ingenuity has been philosophically unnecessary since my self-undermining dilemma demonstrates that the closure argument doesn't genuinely threaten our knowledge of the external world. Indeed, I believe that it is likely advantageous to continue to provide non-self-undermining responses to the closure argument since this might provide more insights into the nature of knowledge, justification, and our epistemic position to the world. There are cases when genuine insights can be attained by knowingly operating under false assumptions. Just like how falsely attributing intentional mental states to artifacts can be advantageous in predicting their behavior,[26] more epistemic fruit might be born from continuing to till the skeptical fields.[27]

However, one might also not be satisfied with my self-undermining dilemma because it relies on the notion of meeting a "dialectical burden." For example, the skeptic might altogether deny this requirement and hold that appealing to this notion to undermine certain arguments for external world skepticism misunderstands the role these arguments serve. The skeptic might claim that the purpose of the closure argument is not to convince us or put us in a position to know its skeptical conclusion, but to demonstrate that some of our baseline epistemic commitments are problematic and need revision (cf. Pritchard 2015: 16, 24). Indeed, on this interpretation, my self-undermining dilemma might support this skeptical position since it demonstrates, via the reductio

[26] See Dennett (2009) for a discussion and a nuanced defense of the view that attributing intentional mental states to (many) artifacts is not just instrumentally useful but a theory of attributing real or genuine mental states to artifacts.

[27] This strategy of granting something one believes to be false to gain more insight into a problem has been advantageous in other areas of philosophy too, e.g., Thomson (1971) grants what she takes to be the false assumption that the fetus is a person from the moment of conception to argue that, *even if* this is true, abortion is still permissible in many cases.

horn of my dilemma, that some of our baseline epistemic commitments are inconsistent and in need of revision.

These thoughts should also be resisted. In general, the dialectical burden enshrines something like the following meta-philosophical commitment: a philosophical view should be taken seriously to the extent to which there are some, hopefully compelling, arguments or considerations in its favor. While a full defense of this claim is outside the scope of this paper, I take this commitment to be obviously true and motivated by the following thought: There are many possible but ludicrous philosophical positions and we should ignore these views until there are compelling reasons to take them seriously. Indeed, to deny the dialectical burden requirement amounts to the claim that, while doing philosophy, we need not provide good reasons or arguments for our philosophical beliefs. Here I am in full agreement with Weintraub (2008: 129) when she says of Hume's problem of induction that: "Like any skepticism worth its salt, it has to be based on an argument rather than merely issue a challenge." The same is true for external world skepticism.[28]

However, I agree that *if* the purpose of the closure argument is only to show that we need to revise what seem to be our "baseline" epistemic commitments, then the closure argument accomplishes this – provided that the reductio horn of my dilemma is sound. After all, by deriving a contradiction, this reductio shows that at least one of the premises of the closure argument is

[28] This defense of the dialectical burden might go some way in addressing and responding to a worry posed by Pyrrhonian skepticism – which on some interpretations (see Lammenranta (2008)) would deny the need to meet a dialectical burden since they do not believe in or assent to any epistemological propositions. The worry is that requiring a dialectical burden might illegitimately rule-out this (and other) more ambitious kind(s) of skepticism. While I believe that Pyrrhonianism is implausible, this is outside the scope of this paper. Instead, I focus only on Cartesian skepticism since it is commonly thought that this latter kind of skeptic aims to put us in a position to know its skepticism (see Comesaña and Klein (2019)). However, if appealing to Pyrrhonian skepticism is the only way for the Cartesian skeptic's closure argument to avoid my dilemma, then this, I claim, is still a significant result.

false. As mentioned above, I agree that many epistemic insights have been, and will likely continue to be, gained by attempting to respond to this skeptical argument. But if the closure argument is intended to *threaten* our knowledge of the external world, then, as my dilemma demonstrates, it fails to accomplish this. To reiterate, my dilemma refutes the closure argument by using only reasoning and premises that the Cartesian skeptic accepts. As such, my argument does not beg the question against the skeptic and intends to dissuade them of external world skepticism. This is significant because convincing the skeptic is an ideal that many epistemologists have strived to achieve in responding to the skeptic. Indeed, my self-undermining dilemma meets its own dialectical burden by demonstrating that our underlying epistemic commitments that the closure argument utilizes cannot be made into an argument that threatens our knowledge of the external world. So, speaking metaphorically, we non-skeptics and skeptics should revise our epistemic commitments *together*.

4. Conclusion

To be clear, this paper does not attempt to provide a cogent argument that external world skepticism is false or a cogent argument that proves the external world exists as (roughly) how we perceive it. My aim is only the modest goal of refuting one Cartesian argument for external world skepticism: the closure argument. However, the external world skeptic might retort that *even if* my dilemma non-question beggingly refutes the closure argument, this does not defeat external world skepticism since there are other Cartesian arguments for this view (e.g. the underdetermination argument). In other words, while I may have removed one argument from the Cartesian skeptic's arsenal, they have other resources at their disposal (cf. Weintraub 2008: 130). At best, I have won a battle in this debate, not the war.

In response, while I believe that a version of my self-undermining dilemma can refute other Cartesian arguments for external world skepticism, I do not have space to demonstrate this here. However, a potentially important implication of my self-undermining dilemma is that it might provide a blueprint for responding to Cartesian (and potentially other) skeptical arguments. Specifically, its four steps (i.e. assume an argument is sound for reductio, show how the argument "cuts both ways," show how the argument meets is dialectical burden, and show that it is refuted or idle) provide a powerful anti-skeptical strategy for refuting skeptical arguments. To speculate, if this strategy is widely successful, then we non-skeptics can potentially win the war against the Cartesian skeptic. For, if they cannot provide a cogent argument using global skeptical hypotheses to support external world skepticism, they will be forced to surrender due to lack of resources (cf. Byrne 2004: 301). Of course, we non-skeptics have to meet our dialectical burden of defending and arguing for the real-world hypothesis. But this is made considerably easier if we do not have to contend with arguments involving global skeptical hypotheses.[29]

Work Cited

Bernecker, Sven and Dretske, Fred, 2000, "Scepticism: Introduction", in *Knowledge: Readings in Contemporary Epistemology*, Sven Bernecker and Fred Dretske (eds.), Oxford: Oxford University Press, 301-306.

Bourget, David, and Chalmers, David, 2014, "What Do Philosophers Believe?" *Philosophical Studies*, 170(3): 465-500.

[29] I would like to thank the audiences at the following conferences for providing helpful feedback on this paper: 2022 North Carolina Philosophical Society and the South Carolina Society for Philosophy Joint Conference, the 2021 Wake Forest University Philosophy Colloquium Series, and the 2021 Varieties of Anti-Skepticism (LVI Reuniones Filosóficas) Conference. I would also like to thank Plínio Smith and Michael Williams for their insightful commentaries for the 2021 Pacific APA. A special thanks to my two anonymous reviewers whose comments made this paper significantly better than it otherwise would have been. Thanks to Hilary Kornblith for encouraging this project and discussing skepticism with me over the years. Lastly, thanks to Elizabeth Hupfer for reading and discussing several different versions of this paper and for making this project possible.

Brueckner, Anthony, 1994, “The Structure of the Skeptical Argument”, *Philosophy and Phenomenological Research,* 54(4): 827-835.

Brueckner, Anthony, 2014, “Skeptical Mystery Tour”, in *Current Issues in Epistemology,* Ram Neta (ed.), New York: Routledge, 119-129.

Byrne, Alex, 2004, “How Hard Are the Skeptical Paradoxes”, *Nous,* 38(2): 299-325.

Cohen, Stewart, 1998, “Two Kinds of Skeptical Argument”, *Philosophy and Phenomenological Research,* 58(1): 143-159.

Comesaña, Juan and Klein, Peter, 2019, "Skepticism", *Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.).

Decker, Jason, 2014, “Conciliationism and Self-Incrimination”, *Erkenntnis,* 79:1099–1134.

Dennett, Daniel, 2009, “Intentional Systems Theory”, in B. McLaughlin, A. Beckermann & S. Walter (Eds.), *The Oxford Handbook of Philosophy of Mind,* Oxford University Press

DeRose, Keith, 1995, “Solving the Skeptical Problem”, *Philosophical Review,* 104(1): 1-52.

Dretske, Fred, 1970, “Epistemic Operators”, *Journal of Philosophy,* 67(24): 1007-1023.

Dutant, Julien, 2007, “The Case for Infallibilism,” in C. Penco, M. Vignolo, V. Ottonelli & C. Amoretty (eds.), *Proceedings of the 4th Latin Meeting in Analytic Philosophy,* Genoa: University of Genoa, pp. 59-84.

Dutant, Julien, 2016, “How to be an Infallibilist,” *Philosophical Perspectives,* 26(1): 148-171.

Fumerton, Richard, 1990, “Metaepistemology and Skepticism”, in *Doubting: Contemporary Perspectives on Skepticism,* ed. M.D. Roth and G. Ross. Dordrecht, Holland: Kluwer.

Fumerton, Richard, 1995, *Metaepistemology and Skepticism,* Roman & Littlefield Publishers.

Fumerton, Richard, 2005, “The Challenge of Refuting Skepticism”, in *Contemporary Debates in Epistemology,* (Eds.) Steup, Matthias, and Ernest Sosa, Blackwell Publishing, pp. 85-97.

Goldman, Alan H., 2007, “The underdetermination argument for brain-in-a-vat scepticism”, *Analysis,* 67(1): 32-36.

Greco, John, 2000, *Putting Skeptics in Their Place,* Cambridge: Cambridge University Press.

Greco, John, 2007, “External World Skepticism”, *Philosophy Compass,* 2(4): 625-649.

Greco, John, 2008, “Skepticism about the External World”, in *The Oxford Handbook to Skepticism,* (Eds.) Greco, John, Oxford University Press, pp. 108-128.

Kelly, Thomas, 2005, “The Epistemic Significance of Disagreement,” in T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, vol. 1. Oxford: Oxford University Press, pp. 167-196.

Kornblith, Hilary, 2021, *Scientific Epistemology: An Introduction*, Oxford University Press.

Kvanvig, Jonathan L., 2008, “Closure and Alternative Possibilities”, in *The Oxford Handbook to Skepticism,* (Eds.) Greco, John, Oxford University Press, pp. 456-483.

Lai, Changsheng, 2020, “The Self-Hollowing Problem of the Radical Sceptical Paradox”, *Erkenntnis,* 85(5): 1269-1288.

Lammenranta, Marcus, 2008, “The Pyrrhonian Problematic”, in *The Oxford Handbook to Skepticism,* (Eds.) Greco, John, Oxford University Press, pp. 9-33.

Le Morvan, Pierre, 2011, “Healthy Skepticism and Practical Wisdom”, *Logos and Episteme,* 2 (1):87-102

Luper, Steven, 2018, "Epistemic Closure", *Stanford Encyclopedia of Philosophy* (Winter 2018 Edition), Edward N. Zalta (ed.).

Lycan, William, 2019, *On Evidence in Philosophy*, Oxford University Press.

McCain, Kevin, 2013, “Two skeptical arguments or only one?”, *Philosophical Studies,* 164(2): 289-300.

McGinn, Marie, 1989, *Sense and Certainty: A Dissolution of Skepticism,* Blackwell: New York.

Moon, Andrew, 2013, “Remembering entails knowing”, *Synthese,* 190(14): 2717-2729.

Moore, G.E., 1959, “Certainty”, in *Philosophical Papers* 226-251. Reprinted in *G. E. Moore: Selected Writings* 171-96.

Nagel, Thomas, 1986, *The View from Nowhere,* Oxford University Press.

Nozick, Robert, 1981, *Philosophical Explanations,* Cambridge, MA: Harvard University Press.

Pritchard, Duncan, 2002, “Recent Work on Radical Skepticism”, *American Philosophical Quarterly,* 39(3): 215-257.

Pritchard, Duncan, 2005a, “The Structure of Sceptical Arguments”, *The Philosophical Quarterly*, 55(218): 37-52.

Pritchard, Duncan, 2005b, *Epistemic Luck,* Oxford: Clarendon Press.

Pritchard, Duncan, 2015, *Epistemic Angst,* Princeton University Press.

Pritchard, Duncan, 2018, “Epistemic Angst”, *Philosophy and Phenomenological Research,* 93(3): 70-90.

Pryor, James, 2000, “The Skeptic and the Dogmatist”, *Nous,* 34(4): 517-549.

Rosenkranz, Sven, 2018, “The Structure of Justification”, *Mind,* 127(506): 309-338.

Rinard, Susanna, 2018, “Reasoning One’s Way Out of Skepticism”, in K. McCain & T. Poston (Eds.) *The Mystery of Skepticism: New Explorations,* Brill Studies in Skepticism, Vol 2, pp. 240-264.

Schonbaumsfeld, Genia, (2016), *The Illusion of Doubt,* Oxford University Press.

Steup, Matthias and Ram Neta, 2020, "Epistemology", *Stanford Encyclopedia of Philosophy* (Fall 2020 Edition), Edward N. Zalta (ed.).

Stroud, Barry, 1984, *The Significance of Philosophical Skepticism,* Oxford: Clarendon Press.

Thomson, Judith Jarvis, 1971, “A Defense of Abortion”, *Philosophy and Public Affairs,* 1(1): 47-66.

Unger, Peter, 1975, *Ignorance: A Case for Skepticism,* Oxford: Clarendon Press.

Yli-Vakkuri, Juhanni and Hawthorne, John, 2022, “Being in a Position to Know”, *Philosophical Studies,* 179: 1323-1339.

Vogel, Jonathan, 1990a, “Cartesian Skepticism and Inference to the Best Explanation”, *Journal of Philosophy,* 87(11): 658-666.

Vogel, Jonathan, 1990b, “Are there counterexamples to the closure principle?”, in M. Roth & G. Ross (Eds.), *Doubting: Contemporary Perspectives on Skepticism,* Dordrecht: Kluwer Academic Publishers.

Vogel, Jonathan, 1993, “Dismissing Skeptical Possibilities”, *Philosophical Studies,* 70(3): 235-250.

Vogel, Jonathan, 2004, “Skeptical Arguments”, *Philosophical Issues,* 14(1): 426-455.

Weintraub, Ruth, 2008, “Skepticism about Induction”, in *The Oxford Handbook of Skepticism,* (Eds.) Greco, John, Oxford University Press, pp. 129- 148

Williams, Michael, 1996, *Unnatural Doubts: Epistemological Realism and the Basis of Skepticism,* Princeton, NJ: Princeton University Press.

Williams, Michael, 2001, *Problems of Knowledge: A Critical Introduction to Epistemology,* Oxford University Press.

Williamson, Timothy, 2000, *Knowledge and its Limits,* Oxford University Press.

Wilson, Jessica M., 2012, "The regress argument against Cartesian skepticism", *Analysis,* 72(4): 668-673.

Wright, Crispin, 1991, "Scepticism and dreaming: Imploding the Demon", *Mind,* 100(1): 87-116.