Inference to anti-skepticism

Can we be *a priori* justified in rejecting skepticism?

Inferens til antiskeptisisme

Er vi berettiget til å avvise skeptisisme *a priori*?

Luca Caiti

Institutt for filosofi og førstesemesterstudier

Universitetet i Bergen

Masteroppgave i FIL0350

Vår 2023

Veileder: Prof. Sorin Bangu

Abstract

In the following thesis, I will discuss the topic of radical skepticism and deal with abductivism as a possible anti-skeptical strategy. I will argue that abduction—or, Inference to the Best Explanation (IBE)—can be utilized as an epistemically valuable strategy against the skeptic. I shall focus on simplicity and will offer my own account of *a priori* ontological simplicity, based on the view called "rationalist abductivism". I will defend the position that the "brain in a vat" is less simple than the commonsense hypothesis. Along the way, I will argue that the abductive principle of simplicity, together with a number of other considerations, provide an epistemic justification for believing in the commonsense hypothesis; moreover, I will note that the BIV skeptical hypothesis actually assumes, even if just implicitly, the existence of an external world. First of all, I will present the skeptical argument (in chapter 1), and then discuss some of the major anti-skeptical strategies (in chapter 2). Then, in chapter 3 I will discuss abductivism and Vogel's approach. Next, I will face several objections to abduction and abductivism (chapter 4). Finally, in chapter 5, after the discussion of some relevant suggestions and BonJour's approach, I will present my own account, followed by my reply to some possible objections to it.

Sammendrag

I den påfølgende oppgaven vil jeg diskutere temaet radikal skepsis og ta for meg abduktivisme som en mulig antiskeptisk strategi. Jeg vil argumentere for at bortføring – eller slutning til den beste forklaringen (IBE) – kan brukes som en epistemisk verdifull strategi mot skeptikeren. Jeg skal fokusere på enkelhet og vil gi min egen redegjørelse for a priori ontologisk enkelhet, basert på synet kalt «rasjonalistisk abduktivisme». Jeg vil forsvare synet om at "common sense"-hypotesen er en enklere forklaring enn "hjerne-på-tank"-hypotesen. Underveis vil jeg argumentere for at det abduktive prinsippet om enkelhet, sammen med en rekke andre hensyn, gir en epistemisk begrunnelse for å tro på "common sense"-hypotesen. Jeg vil i tillegg argumentere for at "hjerne-på-tank"-hypotesen implisitt antar eksistensen av en ytre verden. Jeg vil først presentere det skeptiske argumentet (i kapittel 1), og deretter diskutere noen viktige antiskeptiske strategier (i kapittel 2). I kapittel 3 vil jeg diskutere abduktivisme og Vogels tilnærming. Deretter vil jeg møte flere innvendinger mot bortføring og abduktivisme (kapittel 4). Til slutt, i kapittel 5, etter diskusjonen av noen relevante forslag og BonJours tilnærming, vil jeg presentere mitt eget syn, etterfulgt av mine tilsvar mot noen mulige innvendinger.

Acknowledgments

I would like to thank, first of all, Sorin Bangu, for his interesting lectures and his commendable and patient work as a supervisor. To the fifth floor, for the time spent together in the trenches. A special thanks to Oda Ystaas and Sebastian Frost Bø, for their help with the translation in Norwegian.

Contents

Introduction

Throughout the centuries, the philosophical debate around skepticism has grown rich and developed widely. In general, there are two different types of skepticism; here I shall call them *practical* or *folk* skepticism, and *theoretical* or *philosophical* skepticism. The first one refers to a sort of practical mindset, or inquisitive attitude according to which we should suspend our judgment and question the information that reaches us. In other words, we should cultivate a critical spirit towards what we hear or see. This is, usually, what we mean by "skepticism" in our everyday life. The second type is a theoretical kind of skepticism, and this is the subject of the philosophical discussion.

This second kind can be of several different varieties and can be either about *knowledge* or about *justified belief*. In the following thesis, I shall discuss a specific kind of philosophical skepticism, called radical or Cartesian skepticism, which traditionally questions the possibility to have knowledge—or even any justified belief—about the external world. This type of skepticism is central in the contemporary debates, in particular in the form of a well-known scenario: the situation in which we are no longer normal human beings experiencing reality, but are "brains in a vat" (BIVs) connected to a supercomputer that simulates reality for us.

After the presentation of the radical skeptical argument, I shall discuss critically some of the best known strategies to overcome it. However, my focus will be the anti-skeptical approach called "abductivism", and the idea of one of its important proponents, Jonathan Vogel. In addition, I shall argue in favor of this approach and endorse a rationalist *a priori* kind of abductivism. In the last chapter, I shall offer my own version of this type of abductivism, which does provide us, together with some additional considerations, a strong epistemic justification to reject the radical skeptical hypothesis.

Here is the plan for my discussion. In the first chapter, I will present the radical skeptical argument in its two most relevant forms: the one relying on the Closure Principle, and the one based on the Underdetermination Principle. I will explain these principles and argue that, in order to deal with the strongest version of skepticism, we should focus on the second form of the argument.

In the second chapter, I will analyze some of the best-known anti-skeptical strategies advanced in epistemology. I will start with G.E. Moore's response to the skeptic, followed by Wittgenstein's reply. I will then discuss two accounts which are broadly speaking Wittgensteinean, that is, Michael Williams' and Crispin Wright's. After that, I will take a critical look at James Pryor's approach. Finally, I will outline Hilary Putnam's ideas on how to offset the skeptic. During the presentation of these authors, I will offer, whenever relevant, my own estimations of the force of these anti-skeptical arguments.

In the third chapter, I will deal with the particular anti-skeptical strategy of abductivism. This strategy is based on the inferential reasoning called abduction, also referred to as Inference to the Best Explanation (IBE). The basic idea is that one tries to find, from the observation of a certain phenomenon, the most likely—*or, best*, all things considered—explanation of such a phenomenon, through an appeal to explanatory principles. I will first outline abduction and abductivism in general; then, I will discuss Vogel's own abductive approach, as spelled out in his influential paper *Cartesian Skepticism and Inference to the Best Explanation* (1990). Despite my sympathy for this strategy, I shall raise several objections against Vogel's position. Some of the considerations I will provide here will be utilized also in the last chapter.

In the fourth chapter, I will examine some of the most important criticisms of the abductive approach, and I will try to rebut them. I will start with Richard Fumerton's article *Skepticism and Reasoning to the Best Explanation* (1992), followed by Duncan Pritchard's objections elaborated in his recent book, *Epistemic Angst: Radical Skepticism and the Groundlessness of our Believing* (2015). Next, I will turn to a discussion of James Beebe's paper *The Abductivist Reply to Skepticism* (2009). I will also present some of Van Fraassen's famous arguments against abduction, based on his 1989 book *Laws and Symmetry*. Finally, I will formulate an additional objection to abductivism, and also sketch a possible way to answer it.

In the fifth chapter, I will consider the particular kind of abductivism called "rationalist abductivism". I discuss it since, I argue, this is the most promising anti-skeptical strategy available. I will first address the suggestions made by some authors (in particular Fumerton and Beebe), and I will critically present Laurence BonJour's position. Then, I will advance my own *a priori* account, focusing on

simplicity as the key notion and central criterion. This account, together with some additional considerations initially raised in the third chapter, will be sufficient, I believe, to give us epistemic justification to reject the skeptic. Finally, I will analyze some possible objections to these ideas.

1. Skepticism and the brains in a vat

Since the publication of his *Meditations* in 1641, Descartes' dream argument, along with his Evil Genius argument, have been of great influence in the modern epistemological debate; in fact, it would not be wrong to claim that modern philosophical theorizing about knowledge and justification begins with these arguments. In his First Meditation—"What can be called into doubt"—Descartes writes:

> Whatever I have up till now accepted as most true I have acquired either from the senses or through the senses. But from time to time I have found that the sense deceive, and it is prudent never to trust completely those who have deceived us even once. Yet although the senses occasionally deceive us with respect to objects which are very small or in the distance, there are many other beliefs about which doubt is quite impossible, even though they are derived from the senses – for example, that I am here, sitting by the fire, wearing a winter dressing-gown, holding this piece of paper in my hands, and so on. […] How often, asleep at night, am I convinced of just such familiar events – that I am here in my dressing-gown, sitting by the fire – when in fact I am lying undressed in bed! Yet at the moment my eyes are certainly wide awake when I look at this piece of paper; I shake my head and it is not asleep; as I stretch out and feel my hand I do so deliberately, and I know what I am doing. Indeed! As if I did not remember other occasions when I have been tricked by exactly similar thoughts while asleep! As I think about this more carefully, I see plainly that there are never any sure signs by means of which being awake can be distinguished from being asleep. The result is that I begin to feel dazed, and this very feeling only reinforces the notion that I may be asleep. Suppose then that I am dreaming, and that these particulars – that my eyes are open, that I am moving my head and stretching out my hands – are not true. Perhaps, indeed, I do not even have such hands or such a body at all. (Descartes 1986, 12–3)

And he goes on:

> I will suppose therefore that […] some malicious demon of the utmost power and cunning has employed all his energies in order to deceive me. I shall think that the sky, the air, the earth, colours, shapes, sounds and all external things are merely the delusions of dreams which he has devised to ensnare my judgment. I shall consider

> myself as not having hands or eyes, or flesh, or blood or senses, but as falsely believing that I have all these things. (Descartes 1986, 15)

These arguments have led to what is nowadays called “Cartesian” or “radical” skepticism, which in essence questions our knowledge of the existence of the external world. Generally, this type of skepticism states that we don’t know, and (worse) cannot know all the things we commonly believe to know about the external world. In particular, the skeptic says, we can’t acquire any sort of justification for our perceptual knowledge about the external world. As opposed to what we commonly take for granted, the skeptic shows us that there might be alternative explanations for why we have our sensory experiences. For instance, instead of perceiving the external world through our senses, as we do commonly believe to be the case, we might be under some sort of sensory deception.

The classic Cartesian arguments have been reformulated by Hilary Putnam in his *Reason, Truth and History* (1981) into what is now the standard scenario discussed in the contemporary debates, i.e., the case of the brains in a vat (BIVs). Putnam presents it as follows:

> [I]magine that a human being (you can imagine this to be yourself) has been subjected to an operation by an evil scientist. The person’s brain (your brain) has been removed from the body and placed in a vat of nutrients which keeps the brain alive. The nerve endings have been connected to a super-scientific computer which causes the person whose brain it is to have the illusion that everything is perfectly normal. There seem to be people, objects, the sky, etc; but really all the person (you) is experiencing is the result of electronic impulses travelling from the computer to the nerve endings. The computer is so clever that if the person tries to raise his hand, the feedback from the computer will cause him to ‘see’ and ‘feel’ the hand being raised. Moreover, by varying the program, the evil scientist can cause the victim to ‘experience’ (or hallucinate) any situation or environment the evil scientist wishes. He can also obliterate the memory of the brain operation, so that the victim will seem to himself to have always been in this environment. It can even seem to the victim that he is sitting and reading these very words about the amusing but quite absurd supposition that there is an evil scientist who removes people’s brains from their bodies and places them in a vat of nutrients which keep the brains alive. The nerve endings are supposed to be connected to a super-scientific computer which causes the person whose brain it is to have the illusion that... (Putnam 1981, 5–6)

Imagine, Putnam continues, that all human beings or even all sentient beings are brains—or just nervous systems—in a vat. What the supercomputer would cause to us is, in fact, a collective hallucination, so that everyone would live this sort of virtual (simulated) life, interacting with each other only through this massive simulation.

Now, there are two main ways of reconstructing the skeptical argument. The most classic one consists of a skeptical hypothesis (*SH*)—that is, any scenario depicting a massive deceiving—and the so-called Closure Principle (CP). This principle can be formulated in the two following ways:[1]

(a) If *S* knows that *p*, and *S* knows that *p* entails *q*, then *S* knows that *q*.[2]
(b) If *S* knows that *p*, and *S* competently deduces from *p* that *q*, thereby forming a belief that *q* on this basis while retaining their knowledge that *p*, then *S* knows that *q*.

The difference between these two formulations is that the second one makes the inferential structure of the CP more explicit, highlighting all the requirements for the inference to go through.

Now, by assuming CP (to derive (ii) below), the generic skeptical argument can be presented in the form of a *modus ponens*:

(i) *S* does not know that ¬*SH*.
(ii) If *S* does not know that ¬*SH*, then *S* does not know that *p*.
(iii) Therefore, *S* does not know that *p*.

Or, also, in the form of a *modus tollens*:

(i) *S* does not know that ¬*SH*.
(ii) If *S* knows that *p*, then *S* knows that ¬*SH*.
(iii) Therefore, *S* does not know that *p*.

---

[1] The CP can be formulated in many other different ways. For my purposes I have presented two of them, the most common one and a more sophisticated one. For an in-depth discussion of CP, see Pritchard (2015, Ch.1).

[2] "*p*" and "*q*" stand for generic propositions about the external world, e.g., "The sky is blue", or "I have hands".

Importantly, the skeptical argument can be also rendered as the following paradox (Pritchard 2015, 15):

(i) One is unable to know the denials of radical skeptical hypotheses.
(ii) The closure principle.
(iii) One has widespread everyday knowledge.

Or, utilizing the contemporary example of the BIVs (Pritchard 2015, 15), we can write:

(i) One cannot know that one is not a BIV.
(ii) If one cannot know that one is not a BIV, then one cannot know that *p.*
(iii) One knows that *p*.

The other way of conceiving the skeptical argument consists of, again, a skeptical hypothesis (*SH*) and another principle called Underdetermination Principle (UP). This principle has been defined with different nuances by several authors. The two definitions I find more interesting are the following:[3]

(a) If *q* is a competitor to *p*, then one can know *p* only if one can non-arbitrarily reject *q*, where non-arbitrarily means that one rejects *q* on the base of its less epistemic merit.
(b) If *S* knows that *p* and *q* describe incompatible scenarios, and yet *S* lacks a rational basis that favors *p* over *q*, then *S* lacks knowledge that *p*.

Accordingly, *S*'s belief that *p* is underdetermined, in the sense that *S* does not and cannot know whether they are subject to a skeptical scenario or not, so both the skeptical and the ordinary scenario have the same epistemic strength and none is more credible than the other. In other words, «the rational support provided by our perceptual experiences does not seem to epistemically favor our ordinary perceptual beliefs over the kind of scenarios depicted by radical skeptical hypotheses» (Pritchard 2015, 29).

[3] The first one is Vogel's (2005, 73) and the second one is Pritchard's (2015, 30).

So, the skeptical argument can then be reformulated as follows:

(i) *S*'s belief that *p* is underdetermined by *S*'s belief that *SH*.
(ii) In order for *S* to know that *p*, *S*'s belief that *p* must not be underdetermined by *S*'s belief that *SH*.
(iii) Therefore, *S* does not know that *p*.

Suppose you see a tree (Vogel 2004, 427). Usually, you would take the tree to be the source of your sensory experience. However, if there is another equally grounded reason to believe that something else is the source of your experience, your common belief that you are actually seeing a tree becomes an arbitrary belief and does not amount to knowledge. The aim of a skeptical argument, in its underdeterministic sense, is to show us that each and every one of our perceptual beliefs about the external world suffers from this issue. Here our commonsense explanation underdeterministically collides with the skeptical explanation, making it impossible to discern which of the possible scenarios is true. Therefore, we are not able to know anything about the external world and we are forced to suspend the judgment even on our most basic beliefs about it.

The skeptical paradox seen above can be also conceived with the UP instead of the CP, as follows (Pritchard 2015, 32):

(i) One cannot have rational support that favors one's belief in an everyday proposition over an incompatible radical skeptical hypothesis.
(ii) The underdetermination principle.
(iii) One has widespread everyday knowledge.

Or, utilizing the contemporary example of the BIVs:

(i) One cannot have rational support that favors one's belief that *p* over the BIV hypothesis.
(ii) If one cannot have rational support that favors one's belief that *p* over the BIV hypothesis, then one does not know that *p*.
(iii) One knows that *p*.

Traditionally, the various strategies against the skeptical argument (or paradox) have been focusing on the first formulation, i.e., the one with the Closure Principle in play.

First of all, one could tackle skepticism by attempting to provide a refutation of the very Closure Principle.[4] The line of argument is divided into three different categories: the first one consists of advancing some counterexamples where the principle in question fails;[5] the second one argues that the Closure Principle leads to some unacceptable epistemological consequences;[6] the third one maintains that such a principle is incompatible with some epistemological theories.[7]

The problem with such views is that it is hard to find a promising and effective argument against CP (especially in its elaborated, (b) form above). First, those counterexamples have been criticized and, thus far, none sounds really convincing.[8] Second, the alleged unacceptable epistemological consequences are the result of the problematic relation between the Closure Principle and other principles (e.g., Ampliativity or Entailment). But this relation and the very soundness of these other principles have been questioned.[9] Third, the incompatibility of CP with some epistemological theories (e.g., Nozick's reliabilism) depends upon the validity of those theories, and they, in turn, have been contested.[10]

Moreover, when we consider CP's second and more sophisticated formulation—provided above as (b)—things may become even worse. For, as noted by Pritchard (2015, 14), it seems rather unclear how we could draw an appropriate deductive conclusion from our own knowledge without coming to know, by means of that very same deduction, its conclusion. In fact, DeRose (1995) has argued that to deny the Closure Principle is to endorse some possible "abominable conjunctions", for

---

[4] It is important to note that these authors have not necessarily rejected the Closure Principle for anti-skeptical proposes, but just for criticizing the principle *per se*.

[5] See, among others, Dretske (1970, 1015–16) and Audi (1988, 77).

[6] See, among others, Huemer (2001), Sharon & Spectre (2017) and Comesaña (2017). For a general overview of this position, see Comesaña's and Klein's entry "Skepticism" in *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), URL=<https://plato.stanford.edu/archives/win2019/entries/skepticism/>.

[7] For example, Nozick's account of knowledge (1981, 172–87). For a more detailed overview, see the entry cited in fn. 6.

[8] See, among others, DeRose and Warfield (1999).

[9] For example, Pryor (2014) has argued against Entailment.

[10] Nozick's "reliabilist" account of knowledge, for instance, has been highly criticized throughout the years. For a review of these ideas, see Goldman's and Beddor's entry "Reliabilist Epistemology" in *The Stanford Encyclopedia of Philosophy* (Summer 2021 Edition), URL=<https://plato.stanford.edu/archives/sum2021/entries/reliabilism/>.

instance, that I know I am seated at my own desk right now and, at the same time, I don't know whether I am a brain in a vat or not.[11]

The second kind of anti-skeptical strategies tries to tackle the skeptic's thought that we don't know—or, we are not justified in accepting—the denials of the radical skeptical hypotheses. According to these views, we are somehow able to acquire knowledge that we are not subject to a skeptical scenario and thus that we do possess knowledge about the external world after all. There are many different lines of argument in favor of this strategy, and I will discuss some of the best known ones in the following chapter.

At this point, one may wonder whether the second formulation of the skeptical argument—the one with the Underdetermination Principle—could give us any enlightenment. Indeed, I find it surprisingly strange that the skeptical argument has always been presented, with relevant exceptions, utilizing the Closure Principle. First of all, the UP seems much more convincing, and thus even *more undeniable* than the CP; in fact, as seen above, many objections have been raised towards the latter, but the same cannot be said about the former. Furthermore, and more importantly, I believe the Underdetermination Principle grasps the essence of the skeptical position.

The very reason why the skeptic has the epistemic power they have is precisely because we can't discern whether we are truly deceived or not. Also, the UP seems implicitly presupposed in the CP-formulation, in particular in point (i): we don't know the denials of the radical skeptical hypotheses—that $\neg SH$—exactly because of the underdetermined situation the skeptic forces us to face. After all, isn't it in the definition of a skeptical hypothesis that, if it were true, we wouldn't be able to know not only "that *p*", but also that we were in a skeptical scenario? If our beliefs weren't underdetermined, we wouldn't be unable to know the denials of the radical skeptical hypotheses. So, the first point of the skeptical argument (and thus the whole argument) formulated with the Closure Principle does seem to depend upon the one formulated on the basis of the Underdetermination Principle.

Therefore, in light of both the possible weaknesses of the CP and this additional thought, I believe it will be not only wise but also fruitful to consider the skeptical position formulated as an underdetermination argument. In fact, the anti-skeptical

[11] Obviously, knowing that I am seated at my own desk right now entails the denial of not knowing whether I am a brain in a vat or not.

strategies that try to prove that we do know the denials of the skeptical scenarios will try to also prove as well that we have rational support that favors one's belief in an everyday proposition over an incompatible radical skeptical hypothesis. Clearly, they wouldn't be able to prove that we know, or are justified to believe, that we are not under a skeptical scenario, if they weren't first and foremost able to prove that we have such a rational support.

Now, before we move on, some final reflections. One may legitimately wonder what would be the point of discussing skepticism after all. Isn't the skeptical position, one could ask, irrelevant to our lives? However, what we should ask instead is not whether skepticism is relevant with regard to our lives, but if it is relevant with regard to our epistemological discussions. What's the point then, one continues, even on epistemological grounds? Isn't that just a sort of airy-fairy mumbo jumbo over nothing?

Let me explain why I think skepticism is indeed relevant for epistemology, and even—why not!—for our lives. First of all, as stated by Pritchard (2005, 192), the skeptic is not a dialectical opponent we have to face, but rather, our intellect, which has discovered the inconsistency of our beliefs and which is struggling with it. Therefore, skepticism is relevant for our philosophical research, because it makes us face problems that, if we want to truly call us *philosophers*, we can't ignore. Secondly, I believe the role of skepticism is to take to the extreme logical consequences the epistemic principles and considerations on which we usually debate. Skepticism is like the testbed of epistemology; like the "final tribunal" of our epistemological quest.

Finally, one may still argue that the rejection of the skeptical hypothesis is actually quite easy and straightforward. The skeptical position—or better, the conclusion of the skeptical argument—is absurd; therefore, so they say, we can simply reject it. But this is, in fact, a wrong conception. If it's our intellect what we are dealing with, it's not the skeptic who tells us absurd things, but *us*. Then, simply claiming that the skeptical conclusion is absurd is clearly not enough. The following pages were hence born to answer the skeptical challenge and provide, hopefully, an epistemic justification to reject the skeptical position.

2. Anti-skeptical strategies

As we saw in the previous chapter, a possible way to confront the skeptical argument is to reject that (i) *S* does not know that ¬*SH*, or, in terms of the paradox, that (i) One is unable to know the denials of radical skeptical hypotheses. I also argued that to reject this point is to reject, first and foremost, that (i) *S*'s belief that *p* is underdetermined by *S*'s belief that *SH*, or, utilizing the paradox, that (i) One cannot have rational support that favors one's belief in an everyday proposition over an incompatible radical skeptical hypothesis.

In this second chapter, I will present the views of some of the most important authors who have tried to answer the skeptic by taking this path. First, I will discuss G. E. Moore's argument and then Wittgenstein's reply to both Moore and the skeptic. Moore has famously argued that we can actually infer from merely seeing our hands that there is an external world. In reply, Wittgenstein maintains that not only is Moore completely off track, but that so is the skeptic as well. Central to his objection is the concept of "hinge proposition". Second, I will analyze two authors who have built their own strategies on Wittgenstein's remarks; these are Michael Williams and Crispin Wright. They both draw on the idea of the hinge propositions to conceive their own anti-skeptical strategies. Whilst Williams advances his inferential contextualism, Wright defends his own non-evidential notion of warrant. Third, I will analyze an author who has replied in the opposite direction of Wright, i.e., James Pryor. He argues in favor of Moore's argument and provides some argumentative support for why it works after all. Finally, I will discuss Hilary Putnam's confutation, based on his semantic externalism. He argues that, even if the skeptical hypothesis is physically feasible, it can't possibly be true.

2.1 Moore and Wittgenstein

A famous anti-skeptical response is George E. Moore's. He presented it in *A Defence of Common Sense* (1925) and in *Proof of an External World* (1939). He argues that we can show that the skeptical hypotheses are false on the basis of the "commonsense view of the world" (Moore 1925). His well-known proof of the existence of the external world can be rendered as follows (Moore 1939, 165–6):

(1) Here is one hand

(2) Here is another hand

Hence,

(3) Two human hands exist

(4) Human hands are external objects

Therefore,

(5) External objects exist.

According to Moore, this is «a perfectly rigorous proof» (1939, 166), for proposition (3), which is known to be true (since it follows from (1) and (2)), is different from the conclusion (5), which also follows from the premises.

In contrast, Wittgenstein tackles the problem of skepticism in *On Certainty* (1969), writing his remarks as a reply to Moore's own anti-skeptical argument. First of all, when Moore claims that he knows a proposition like (5), "There are external objects", he should be able to offer compelling reasons supporting such a claim. When someone says "I know", they are referring to a possibility of demonstrating the truth of that claim.[12] Secondly, a compelling reason that supports a knowledge claim has to be more certain than the proposition we claim to know. However, «my having two hands is, in normal circumstances, as certain as anything that I could produce in evidence for it» (*OC* §250).

Furthermore (Child 2011, 193), in order for Moore to have a *proof*, the argument presented has to increase Moore's—and our—knowledge, by giving some grounds to believe something that he didn't previously know. Nevertheless, the proof Moore suggests is unable to do so. In fact, Moore could acquire a justification to believe that his hands exist only through the acquisition of a justification to believe that there are external objects. In other words, it is possible to infer the conclusion from the premises only if we assume the conclusion. Clearly, this is a circular argument.

But there is a complication: according to Wittgenstein, even if Moore cannot claim to know that there are external objects, neither can the skeptic raise doubts about it. Wittgenstein maintains that every epistemic inquiry presupposes taking something for granted, since «the *questions* that we raise and our *doubts* depend on the fact

[12] See *OC* §243: «One says "I know" when one is ready to give compelling grounds. "I know" relates to a possibility of demonstrating the truth. Whether someone knows something can come to light, assuming that he is convinced of it. […]».

that some propositions are exempt from doubt, are as it were hinges on which those turn» (*OC* §341). It is in «the logic of our scientific investigations that certain things are *in deed* not doubted» (*OC* §342), as «part of our *method* of doubt and enquiry» (*OC* §151). We just can't inquire into everything at once: for certain things all we have is assumptions, because «if I want the door to turn, the hinges must stay put» (*OC* §343).

So, for Wittgenstein, certain propositions—called "hinge propositions"—function like pillars above which our epistemic enquiries grow.[13] These propositions are of two kinds: first, they are generic empirical propositions like "I have two hands", "For months I have lived at address A" (*OC* §70), "All human beings have parents" (*OC* §240), "I have never been on the moon" (*OC* §269); second, they are basic mathematical propositions like "Twelve times twelve equals one hundred and forty-four" (*OC* §§447-8). Since we presuppose these propositions every time we think and speak, we cannot claim to know them, but they cannot also be called into doubt.

Moreover, the never-ending doubt that the skeptic wants us to face is not a legitimate epistemic enterprise at all, for «if you tried to doubt everything you would not get as far as doubting anything. The game of doubting itself presupposes certainty» (*OC* §115). As an example, Wittgenstein talks about a teacher who is trying to educate a pupil who «will not let anything be explained to him, for he continually interrupts with doubts, for instance as to the existence of things, the meaning of words, etc. The teacher says "Stop interrupting me and do as I tell you. So far your doubts don't make sense at all"» (*OC* §310).

According to the so-called "Therapeutic reading" of *On Certainty*, the skeptical doubt is not simply wrong, but rather nonsense.[14] The skeptic fails to understand the way our language works; their words lack meaning. Since «what your words say depends upon what they are *doing*—how they are at work—in a context of use», it is

---

[13] In the contemporary debate, some authors have re-elaborate the idea of the hinge propositions in different ways. Besides Michael Williams and Crispin Wright, with whom I will deal in the next section, an example can be found in Annalisa Coliva's work (2015). Other authors who have written in this direction are Duncan Pritchard (e.g., 2012, 2021, 2022) and Danièle Moyal-Sharrock (e.g., 2016 and Sandis & Moyal-Sharrock (2022)).

[14] There are different possible readings of Wittgenstein and *On Certainty*. For example, Danièle Moyal-Sharrock and William H. Brenner (2005) distinguish between four different readings: the Framework reading, the Transcendental reading, the Epistemic reading, and the Therapeutic reading. For both my purposes and conciseness I have just outlined the last one, which I find to be the more interesting.

far from clear what the skeptic is «doing with their words—i.e., what the context of use is supposed to be—and hence what it is that they are saying» (Conant 1998, 242). The skeptic faces a dilemma (Conant 1998, 250): on the one hand, they can stay within our language-games but then their words do not express that never-ending doubt they want to advance; on the other hand, they can get out of our language-games, but then, as we saw, it is no longer clear what they mean by their words.

Significantly, Wittgenstein's purpose is not to offer a refutation of the skeptic, that is, to prove the skeptic wrong. Rather, he questions the very meaning of their claim. He describes all the different interpretations of what the skeptic might mean with their words to demonstrate the inconsistency of their desire, as they want to have multiple alternatives simultaneously without committing to any one in particular.

Granted the Therapeutic reading, we may legitimately ask whether Wittgenstein's remarks succeed as an anti-skeptical strategy. On the one hand, the considerations offered in *On Certainty* seem to be of a more fundamental nature than the canonical epistemic response. In fact, it wouldn't even make sense to answer the skeptic trying to find an epistemic solution, since the skeptic is misplaced in the first place. On the other hand, one could argue that Wittgenstein's response is not epistemically valid. If Wittgenstein does truly not want to refute the skeptical argument, but only wants to show the linguistic issue faced by the skeptic, then the approach he offers, even if still interesting and compelling, is not purely epistemic.

In order to find a *purely epistemic* response, we need to provide a proof that the skeptic is wrong; or, better put, we need to find an *epistemic justification* to believe that the skeptic is wrong. Wittgenstein does not seem to offer this. Still, he states that the skeptic is not epistemically legitimized because their doubt presupposes certainty, and there can't be such a thing as the never-ending doubt they propose. Nevertheless, so I believe, it is not so clear whether the skeptic does necessarily propose that never-ending doubt, even just implicitly. When the skeptic presents us the alternative scenario of us being brains in vats and being controlled by a supercomputer, he depicts a very straightforward scenario. It is true that, following the logic of the skeptic's argument, there could be an infinite chain of supercomputers controlling one another so that the skeptic would find themselves in the position of holding a very strange position. It is also true that the skeptic's

argument itself presupposes some kind of certainty. However, the skeptic's point still remains: that we could be brains in vats and that we don't know whether we truly are or not.

Moreover, I believe that, if we followed Wittgenstein, we would not still be epistemically justified—at least fully—in believing that we are not BIVs. To explain my point, let me introduce the following distinction.[15] I think that, in our specific context, there are two different types of justification that we could (try to) provide: let me them "negative justification" and "positive justification". According to the first type, someone—us—is justified in believing ¬A because the person who has advanced A—the skeptic—is misled or incoherent (or whatever would undermine their legitimacy as a proponent of A). According to the second type, someone—us—is justified in believing ¬A because they have a justification independent from the person—the skeptic—who has advanced A to believe ¬A. I maintain that the negative justification, when present, is surely relevant and might even be epistemically valuable, but only the *positive* justification will completely satisfy completely the epistemic demands.

Therefore, what Wittgenstein offers is, at best, a negative justification for believing that we are not BIVs. But the skeptic could legitimately ask for a positive justification, that is, for a higher level of justification. In other words, if we want to answer the skeptic in a fully-fledged epistemic way, Wittgenstein's response does not seem to appeal as completely satisfying.

2.2 Williams and Wright

Two authors who have taken inspiration from Wittgenstein in their anti-skeptical approaches are Michael Williams and Crispin Wright. The first one (Williams 1991) has advanced an account called "inferential contextualism".[16] Williams' fundamental idea is that in every context there is a set of fixed beliefs—called, after Wittgenstein, "hinge beliefs"—that cannot be subject to epistemic evaluation within that same context. The validity of other beliefs can then be evaluated based on this established

[15] I have conceived this distinction on my own. Since it is not such an incredible finding, I do presume that someone has talked about it in the past. Still, I couldn't find anything in the literature to refer to.

[16] Throughout the years, different types of contextualism have been advanced, even though William's remains the most famous one. For example, Keith DeRose (e.g., 1992) is a strenuous defender of this position in the contemporary debate.

set of hinge beliefs. In every specific context certain presuppositions are logically required, for «in particular contexts of inquiry, certain propositions stand fast as a matter of *methodological necessity*» (Williams 1991, 124, *my italics*).

Furthermore, this hinge role can be played by different beliefs in different contexts, with the inferential structure changing accordingly. Thus, the very same beliefs that play a hinge role in one context will not have this special feature in another context. So that if in one context they are exempt from epistemic evaluation, in another they could be tested. Thus, testing a hinge is, in a certain sense, impossible, for, as we try to do it, we find ourselves in another context where the hinge in question becomes a "normal" belief subject to epistemic evaluation.

Let us consider the following well-known example, due originally to Fred Dretske (1970, 1015–16). Suppose to enter a zoo and see some zebras. In a normal situation, no one would reasonably start questioning if the zebras are really zebras and not cleverly painted mules. Thus, in this context, the proposition "There are some zebras" is not object of doubt at all and functions as a hinge belief. However, if we found ourselves in another context where we had good reasons to dispute whether the animals we see were truly zebras and not cleverly painted mules, the proposition "There are some zebras" wouldn't be a hinge anymore, but a "normal" empirical proposition object of enquiry.

Williams maintains that no supporting ground is needed in order to claim the hinge beliefs of a specific context. In addition, he argues that no hierarchy of contexts occurs; that is, no context is epistemically superior or more demanding than another—e.g., the BIV context as opposed to an ordinary real-world context. So, if it is true that we do not have knowledge about our ordinary beliefs in the skeptical context, we don't necessarily lack that same knowledge in other everyday contexts, which possess the same epistemic status as the skeptical one. In ordinary contexts, we can claim to know, in a tacit way, the presupposed hinge beliefs, even if we are unable to offer an evidential ground for such knowledge.

Moreover, Williams maintains that those propositions which stand fast and which we tacitly know are not needed to be necessarily true. Indeed, he writes:

> The general moral here is that questions about a proposition's epistemic status must always be separated from questions about its truth. If epistemic status is fixed by the direction of inquiry, epistemic status is context-sensitive. Truth however is not. A

> proposition is either true or not. But, […], we cannot say, in a similarly unqualified way, that a proposition is either open to doubt or not. Sometimes it will be and sometimes it won't. Generally speaking, a proposition is neither true because it stands fast not stands fast because it is true. (Williams 1991, 124)

Now, several objections can be raised to Williams' contextualist account. Here I want to focus on one particular criticism which I find significantly compelling.[17] First of all, if we want to truly know some propositions, e.g., the denials of skeptical hypothesis, and we want to truly claim that knowledge, then inferential contextualism puts us in a rather precarious situation. If we say we know P and at the same time we admit that we have knowledge of P only in certain contexts—that is, not *always*—can we say that we really know P? It is reasonable to point out that knowledge bound to circumstance is *not* knowledge, especially if we are aware of such binding. Moreover, if the skeptical hypothesis is not (shown to be) false in all contexts, the anti-skeptical strength of such an approach is flawed, since our commonsense beliefs about the external world will not be completely epistemically safe. So, Williams's approach seems to leave us quite unsatisfied as well.

The second author—Crispin Wright (1985, 2004, 2014)—has proposed an account based on the so-called "conservatism about perceptual justification".[18] He argues that for certain propositions, to believe them is a requirement in any epistemic evaluation and, hence, that no supporting ground can be offered for them. Indeed, «one *cannot but* take certain such things for granted» (Wright 2004, 189), for we could even investigate the presuppositions involved in a particular case but, in doing so, we would be forced to make further presuppositions likewise general. Whenever I claim to have justification for a proposition, it is only by means of some presuppositions which lack any evidence at all. This is, Wright maintains, a necessary truth (Wright 2004, 189).

Nevertheless, these beliefs are epistemically supported, even if such support cannot be earned. It will be a sort of unearned epistemic support—an "unearned

---

[17] I have formulated this objection by adapting one found in Pritchard (2005, 213–5).

[18] The introduction of this term is due to James Pryor (2004). In general, this view maintains that, in order for a subject S to have a prima facie justification to believe P, S needs an antecedent and independent justification to believe a hypothesis H which counts as a precondition of S's perceptual experience. For example, S's justification for believing that S is looking at star X during the night with a telescope depends upon the justification that the telescope is functioning correctly.

warrant", as he calls it—but still positive and valuable. The point is that in any sort of epistemic enquiry I have to take risks, to bet on certain things—e.g., the reliability of my senses—just as I bet and take a risk, even if minimal, on the solidity of a chair every time I sit down on it. However, Wright nuances, we are not entitled to believe those kind of propositions in a fully fledged way, but only to *accept* them; that is, to take them for granted. Whenever we are warranted in assuming that P without knowing whether P is likely true, we are entitled to *accept* that P, but not to (fully) believe that P.

The propositions one is entitled to take for granted and through which one pursues every epistemic enquiry are called by Wright "cornerstone propositions". These propositions are nothing but our anti-skeptical commitments, e.g., that we are not BIVs, and that there is an external world. Moreover, he advances the idea of "entitlement of cognitive project" (2004, 191–2), as a non-evidential notion of warrant:[19]

A subject S with a given cognitive project is entitled to trust a proposition P if

(i) P is a *presupposition* of the project, i.e., if to doubt P (in advance)—or weaker: being open-minded about P[20]—would rationally commit one to doubting (or being open-minded about) the significance of the project;
(ii) there is no sufficient reason to believe that P is untrue; and
(iii) the attempt to justify P would involve further presuppositions in turn of no more secure a prior standing... and so on without limit; so that someone pursuing the relevant enquiry who accepted that there is nevertheless an onus to justify P would implicitly undertake a commitment to an infinite regress of justificatory projects, each concerned to vindicate the presuppositions of its predecessors.

Moreover, Wright notes (2004, 192) that whenever we are engaged in a project where we have to make further presuppositions in order to support the already accepted project's presuppositions, and when the new presuppositions are no more

[19] This specific definition, which I find particularly clear, is Pedersen's (2009, 445).

[20] «Doubt is a stronger attitude than open-mindedness», Pedersen writes, for «to doubt that P amounts to holding a positive attitude towards the negation of P whereas open-mindedness involves a positive attitude towards neither P nor its negation» (2009, 445).

secure (or we are not going to lose anything and we may gain something by doing so), we are entitled to just take the first ones on trust. He writes:

> This line of reply concedes that the best sceptical arguments have something to teach us—that the limits of justification they bring out are genuine and essential—but then replies that, just for that reason, cognitive achievement must be reckoned to take place *within such limits*. The attempt to surpass them would result not in an increase in rigour or solidity but merely in cognitive paralysis. (Wright 2004, 191)

Admittedly, according to Wright (2004, 206), the best we can offer is a *skeptical solution* to the skeptical argument. Such a solution might not be able to block skepticism entirely; however, Wright argues that it is capable of limiting the damage. This strategy concedes to the skeptic that we cannot indeed advance any claim to know[21] any of our presuppositions (or cornerstones), but it then tries to make a compromise, adding that we are nonetheless warranted, even if in a different non-standard way, to take them for granted.

Several criticisms have been raised towards Wright's unearned warrant. The first line of objection is that such an entitlement provides a merely pragmatic/practical, rather than epistemic, justification.[22] Wright acknowledges that, if we conceded the skeptical conclusion, it would lead us to an intellectual paralysis and a self-defeat, and thus we must reject it and recognize the validity of our ultimately groundless beliefs. Nevertheless, our need for such hinges does not represent an epistemic reason to believe in the truth of those hinges but, rather, a pragmatic strategy to support our acting *as if* they were true (Pritchard 2005, 207). Significantly, the skeptic is questioning the *epistemic* status of our common beliefs, and not so much the practical validity or usefulness of taking such beliefs for granted. Therefore, Wright does not seem to provide *any* solution, not even skeptical, to skepticism. In other words, Wright does nothing to "nip Cartesian skeptical argument in the bud" (Jenkins 2007, 27).[23]

---

[21] «Note: have no *claim* to know, not: do *not* know» (Wright 2004, 206, ft. 23).

[22] Among others, Pritchard (2005) and Jenkins (2007).

[23] See Pedersen (2009) for a defense of Wright's approach against this criticism.

A second line of objection follows the so-called "leaching problem".[24] This criticism stems from the fact that the cornerstones are justified and justifiable without providing any evidence. If what we have at the base of our beliefs is a mere acceptance of those cornerstones, also our general epistemic enquiries—cognitive projects, as Wright would call them—will be likewise only taken for granted, without any superior epistemic status. For example, if one is justified in believing that there are two hands in front of them only through an unearned warrant or entitlement—that is, a mere unevidenced acceptance, even if rational—to believe that there is an external world, then one is likewise entitled in the same unevidenced base to believe in the presence of their hands. Hence, they will not know they have two hands, nor will they be really justified in believing they have them.

Even though Wright has replied to the most relevant objections in a recent paper (Wright 2014), including the two I have presented here, I still do believe that the problems highlighted above remain rather compelling. Unfortunately, I don't have the space here to elaborate on this debate. Nevertheless, if such criticisms remain in play, Wright's account will also leave us rather unsatisfied.

2.3 Pryor

An author who has advanced an opposed anti-skeptical position to Wright's, based on the so-called "liberalism about perceptual justification",[25] is James Pryor (2000, 2004). According to Pryor, Moore does have a justification to believe that his hands exist and that the external world exists. In fact, Moore «*can* acquire justification to believe there's an external world by having experiences of hands and reasoning in the way he does» (Pryor 2004, 351).

The basic idea is that when our experience represents us two hands, this *per se* justifies us prima facie in believing that there are two hands. Such a justification does not involve any antecedent presupposition about our experience and remains

[24] Wright himself is aware of the problem and he replies that «there is leaching, but that it is at one remove and can be lived with» (Wright 2004, 207). His explanation (2004, 207–9), which I can't report here, has not seemed very convincing. In fact, Aidan McGlynn (2014) has recently advanced an objection.

[25] The term refers to the view opposed to conservatism about perceptual justification (see Pryor 2004 and fn. 18 above). This view denies that S is required to have an antecedent justification to believe H in order to have a prima facie justification to believe P. Still, it values ¬H—i.e., the evidence against H—as undermining S's prima facie justification for believing P.

valid as long as we are merely experiencing the two hands in front of us. Of course, there could be some hypotheses—e.g., the BIV-skeptical scenario—that may undermine our justification. However, in order for us to acquire that justification, we do not need to first have a justification to believe that those hypotheses are false. In other words, the claim is that we do have some basic perceptual beliefs, formed directly from our experience, which we are prima facie justified to have.

To clarify: what constitutes a prima facie justification to believe P is not simply that at first glance it seems like we are justified in believing P. We have a prima facie justification when, even if that same justification can be undermined or defeated by further evidence, «*in the absence of* such further evidence—and when the further evidence itself gets defeated—your prima facie justification to believe P must constitute some degree of *all things considered* justification to believe P» (Pryor 2004, 353). In a footnote (2004, 353, fn.7), Pryor clarifies that what we are dealing with here is not a justification to *fully* believe P, but just to *partially* believe P. As he states, we can have a *degree* of justification to believe P even if such a degree does not justify us to fully believe P.

In Pryor's opinion, what makes our perceptual experiences have this epistemic feature is precisely what he calls the "phenomenology of perception":

> I think there's a distinctive phenomenology: the feeling of *seeming to ascertain* that a given proposition is true. This is present when the way a mental episode represents its content makes it feel as though, *by* enjoying that episode, you can *thereby just tell* that that content obtains. We find this phenomenology in perception and in memory. When you have a perceptual experience of your hands, that experience makes it feel as though you *can just see* that hands are present. […] My view is that our perceptual justification comes from that phenomenology. Having the phenomenology of seeming to ascertain P is *what makes us* have prima facie justification to believe P. (Pryor 2004, 357)

It could happen, Pryor continues, that our perceptual experience leads us to believe more than what our phenomenology presents us. For instance, we can see a man with a blue uniform, therefore enjoying the "phenomenology of seeming to ascertain" that such a man is present, and believe the police are here. Nevertheless, in this case, our perceptual experience—a man with a blue uniform—is not enough

per se to justify us prima facie in believing the police are here. We don't enjoy the phenomenology of seeming to ascertain that the police are here. Our perceptual experience justifies us in believing only the propositions it gives us the phenomenology of seeming to ascertain.

Moreover, Pryor (2004, 366–70) also provides an interesting analysis of the meaning of doubting and the skeptical argument. In his view, doubt is not necessarily some degree of disbelief, for sometimes it is more just a matter of agnosticism or suspended judgment. Pryor calls doubt in this second sense "hypothetical doubt". In the case of the skeptical hypotheses, we could regard them as equally likely to be true or false. Significantly though, that does not mean we actually believe them. When the skeptic tells us about the alternative scenarios there could be, we don't stop believing our commonsense view of the world, but we just entertain the hypothesis that such believing could be false, with all the consequences that would follow. Those hypothetical doubts, however, are not enough *by themselves* to undermine our perceptual justification.

Imagine, says Pryor, you think to possess a prima facie justification to believe P. However, a trustworthy philosophical authority provides you a negative higher-order justification: that you don't—or even can't, I would add—actually have any justification, not even prima facie, to believe P. As a consequence, your first-order prima facie justification will be, at least to some degree, undermined. In addition, even if what you acquire is just a negative higher-order *belief*, your prima facie justification will be undermined as well, for such a belief will rationally obstruct you from believing P. Keeping this in mind, let's consider the following example:

> You have a visual experience of your hands. That experience gives you some prima facie justification to believe you do have hands. Initially, you're *inclined to believe* you have hands, on the basis of your experiences. Then along comes a skeptic. He starts presenting various undermining hypotheses U_1, U_2,… He argues that since you have no antecedent justification *against* those hypotheses, your experiences don't really give you justification—even prima facie justification—to believe you have hands. He argues that if *you did* have justification to believe there's an external world, you ought to be able to rationally persuade him that there is, but you can't. As it happens, the skeptic's arguments are flawed. Before you encountered him, you *did* have justification to believe you have hands and so on, contrary to what he's claiming. But

> the skeptic is a smooth dialectician. His arguments sound pretty compelling to you. You don't discern their flaws. (Pryor 2004, 367)

Since the skeptic's argument, Pryor goes on, sounds compelling and you can't see its flaws, you are apparently justified in believing the conclusion. So, having listened to the skeptic makes you justified in believing that your perceptual experience does not provide any justification, even if you really have that (prima facie) justification.

Otherwise, we could deny that we are justified to believe the conclusion of a flawed, even if compelling, argument. What the skeptic's argument does is persuading us: it persuades us that your perceptual experience doesn't provide any justification whatsoever. Such an argument does not really make us believe that, for example, we are actually brains in a vat, but it makes us believe the negative higher-order claim that we don't have that prima facie justification. In this second case, the skeptic's argument blocks us from believing in the existence of our hands based on our perceptual experience, thereby making us no longer justified in believing we have hands.

Nevertheless, no matter which of the two outcomes we choose, for neither of them entails that the skeptic is right. Pryor writes:

> The skeptic makes claims about *all subjects*, even subjects who haven't heard his argument. On the story I just told, those claims are false. But subjects who *do* hear the skeptic's arguments, and are partly taken in by them, do really end up with some of the epistemic difficulties the skeptic says we all suffer from. Skepticism isn't the truth about all of us, then. It's just a disease that some of us catch. The way to cure the disease is to realize that skepticism *isn't* the truth about all of us: the skeptic's arguments are flawed. So our negative higher-order beliefs are false. When we give those negative higher-order beliefs up, then the prima facie justification our perceptual experiences gave us all along will be undefeated and unobstructed. (Pryor 2004, 368)

According to Pryor, the problem with Moore's argument is of a dialectical kind: it is not dialectically effective against the skeptic. The skeptic doubts that our perceptual experience provides any justification to us, and precisely such a doubt impedes Moore from persuading him. Nothing is wrong with Moore's argument and its justificatory structure: what is wrong «is that the skeptic has doubts he ought not to have» (Pryor 2004, 369).

However, it is still true, Pryor maintains, that Moore does not give a very philosophically satisfying response. His own argument does indeed provide us some reasoning for being justified to believe in the existence of the external world, but we would need to articulate much argumentative support, which Moore does not provide. In fact, a philosophically satisfying response to the skeptic will be formed exactly by such an argumentative ground and not by the mere argument.

As I understand it, Pryor's own account is nothing but an attempt to articulate that lacking support. He accepts Moore's argument and tries to give some explanation of why it actually works. But, let me ask, *does* it? I believe a first objection could be the following. Pryor's strategy concerns only *perceptual* justification, which by itself does not say anything regarding the truth or falsity of skeptical hypotheses. Although Pryor's analysis of the skeptical doubt is definitely worth considering, still, if we truly are BIVs our prima facie perceptual justification doesn't help us very much after all.

Let us imagine to actually be brains in vats, which thing is, as it stands now, underdetermined. Pryor maintains that if we didn't have the negative higher-order beliefs the skeptic makes us face, the (prima facie) justification we would acquire through our perceptual experience would count as compelling and we would be justified in believing we are not BIVs. The problem is that Pryor's approach doesn't offer any rational support to favor the commonsense hypothesis over the skeptical scenario, and thus to "give up" those negative higher-order beliefs.

For, how could our sensory experience provide us a relevant justification if we actually were brains in a vat—which is, once again, underdetermined? Indeed, the argument Pryor advances seems to fall in a kind of circularity: for such a strategy to work—that is, for our sensory experience to provide a relevant justification—we must not be brains in a vat, but that is precisely what the strategy has to *epistemically* confute. In other words, Pryor seems to fall into the fallacy called *petitio principii*, implicitly assuming the conclusion in the premises.

Moreover, how is stating that skepticism is a problem only for those who engage with it and isn't the truth about all of us, but only, *at best*, of those who hear its argument, supposed to give the skeptic a decisive counterargument? It might be consistent with Pryor's view, but that doesn't lead us any further.

In addition, a second objection may be raised by the very skeptic Pryor confronts: if Pryor himself admits that the (prima facie) justification we apparently possess is not a full-degree justification, but only a justification of some degree, then the question is: how could such view succeed in its anti-skeptical aim? The skeptic may legitimately ask for a more powerful epistemic justification, that is, a "full-degree" justification. Not being able to provide it, the skeptic continues, just does skepticism a favor, proving precisely its point. So, it seems like Pryor's approach is not entirely satisfying when it comes to finding an epistemic justification for rejecting skepticism.

2.4 Putnam

In *Reason, Truth and History* (1981, in particular 5–17) Putnam offers his own confutation of the skeptical argument, based on a position called "semantic externalism". Roughly put, in the field of philosophy of language, semantic externalism is the view according to which the meaning of a word or a proposition is determined, or at least influenced, by external factors, i.e., factors external to the individual speaker or thinker. Examples of these external factors are the social and cultural practices of the society in which the subject lives, and, in general, the environment of the subject. Specifically, some semantic properties of a word depend, for their individuation, on one's external environment. The typical case is the one of two identical individuals whose words and thoughts, even if apparently the same, might have different meanings just because they live in different environments.

In opposition to the skeptic, Putnam argues that the hypothesis that we are brains in a vat is physically feasible but can't possibly be true, for the BIV scenario is somehow self-refuting. The claim is that if we actually were brains in a vat, we wouldn't even be able to think, let alone say, that we were. Although there is a possible world in which every human is a brain in a vat—and has our own same experiences and thinks and says every word we think and say—those humans can't refer to what *we* refer to; in particular, they can't think or say that they are brains in a vat.

Importantly, the envatted brains are certainly conscious, functioning, and intelligent brains, but their words cannot refer to what *our* words refer to. Specifically, their words cannot refer to external objects, or, more precisely, to the

external objects our words refer to. For example, their word “tree”, even if arranged in an identical sequence to *our* word “tree”, does not refer to our—i.e., an *actual*—tree. Putnam writes:

> For there is no connection between the *word* “tree” as used by these brains and actual trees. They would still use the word “tree” just as they do, think just the thoughts they do, have just the images they have, even if there were no actual trees. Their images, words, etc., are qualitatively identical with images, words, etc., which do represent trees in *our* world; but […] the brains in a vat are not thinking about real trees when they think “There is a tree in front of me” because there is nothing by virtue of which their thought “tree” represents actual trees. (Putnam 1981, 12–3)

The brains in a vat think about trees and their words refer to trees just because the supercomputer’s program connects the system of language to non-verbal inputs and outputs. Taking inspiration from Turing’s Test, Putnam states that not even the supercomputer’s program refers to actual trees, for it is just a device programmed to give a certain input (“tree”) to the envatted brains, and this without really knowing what trees are or having a non-verbal interaction with them.

Moreover, to add more complications to this picture, it can be argued that when the brain in a vat thinks “There is a tree in front of me” what happens is that it is referring to trees “in the image” (i.e., in the virtual simulation), or to the supercomputer’s input that causes the experience of the tree. What an envatted brain means when it says “There is a tree in front of me” is that there is a tree in the image which is in front of the “me” in question. Alternatively, they can mean that the electronic signal which causes that representation is coming right at that time from the supercomputer. Similarly, the word “vat”, when thought or said by such a brain, does not and cannot refer to a real vat, but only to a vat *in the image* (or to the related signal). Indeed, in that case, the word “vat” doesn’t have any causal connection to a real vat.

Therefore, if we really were brains in a vat, what we would mean by the proposition “We are brains in a vat” would be, if anything, that were brains in a vat *in the image*. Nevertheless, Putnam continues, the hypothesis that we are brains in a vat contains also the fact that we are not brains in a vat in the image; that is, that our virtual experience is not that we are brains in a vat. Hence, finally, if it’s true that we

are brains in a vat, then the proposition "We are brains in a vat" says something false, or better, *it is* necessarily false. Consequently, it is self-refuting.

Throughout the decades, several criticisms of Putnam's anti-skeptical strategy have been advanced. In general, they focus on two different possible claims: that Putnam's argument is question-begging, or worse, that it is self-refuting too. An example of the first objection has been advanced by Anthony Brueckner (1986), while Crispin Wright (1992) articulated an example of the second. I will outline only the latter.[26]

First of all, let us start with Wright's own reconstruction of Putnam's argument. The argument starts from two fundamental premises (Wright 1992, 73–4):

(i) My language is disquotational;[27]
(ii) In the BIV's language, "BIV" does not refer to BIVs.[28]
From these two premises everything else follows:
(iii) In my language, "BIV" is a meaningful expression;[29]
From (i) and (iii):
(iv) In my language, "BIV" refers to BIVs;
Hence, from (ii) and (iv):
(v) My language is not the BIV's language;
But, from the definition of vat-English,
(vi) If I am a BIV, my language is the BIV's language;
So,
(vii) I am not a BIV.

According to Wright, there is a problem with the second premise—(ii) In BIV's language, "BIV" does not refer to BIVs. The reason is that, together with (iv)—In my language, "BIV" refers to BIVs—it is inferred that (v) My language is not the BIV's language. Now, the speaker can conclude that their language is not the BIV's

[26] I have used Pritchard (2016) as a guide.

[27] *My* language is the language we are using for reflecting on the matter, that is, common English. Moreover, it is "disquotational" in the sense that it «permits correct, homophonic characterisation of the references, satisfaction-conditions and truth-conditions of meaningful expressions within it» (Wright 1992, 73).

[28] This premise follows from semantic externalism.

[29] That "BIV" is meaningful is a prima facie plausible consideration, which shouldn't be contentious at all.

language, what they cannot conclude is that *no one* speaks the BIV's language. Significantly though, if the argument is valid, it can be equally phrased in the first or third person. The subject can be "I", "Crispin Wright", or anyone else, still there is no difference in the legitimacy of the argument (Wright 1992, 76–7).

Nevertheless, it is not possible, Wright maintains, to prove *a priori* that any other person other than me is not a brain in a vat. He writes:

> Without supplementary information, you cannot validly infer anything from (i) and (iii) about how to *specify* what is the reference of 'brain-in-a-vat', as used in [another person]'s language. All you can infer is that a specification in CW's language would be homophonic. That is the same thing as (iv) only if it is presupposed that your language —the language in which the argument is presented—is [the other person]'s. (Wright 1992, 77)

That even if the language of two different people sounds the same it doesn't mean that they mean the same thing—this is precisely the point of semantic externalism.

We could nonetheless save the argument, Wright thinks, by adding another initial premise (Wright 1992, 78):[30]

(P) The following argument is formulated in *X*'s language, which coincides with the language of the assessor of the argument in the meaning assigned to all the expressions involved, including "BIV" and "BIV"s language".

But how can I know whether my use of "BIV" means for me, the assessor of the argument, what it means for Putnam, the proponent of the argument? It seems then that «the kind of knowledge of content that externalism seems to forbid (or at least problematize) is the same kind of knowledge that is needed to make Putnam's anti-skeptical argument work» (Pritchard 2016, 82).

Setting aside Wright's objection, things can get even worse for Putnam's argument. According to Pritchard (2016, 75–7), Putnam does not mean to claim that we *know* that we are not brains in a vat, but rather to claim that we cannot truly think that we are brains in a vat. Whereas the former is an epistemological claim, the latter is a *semantic* claim. What Putnam does is trying to challenge the widespread idea

[30] I have revised it, due to presentation purposes.

that brains in a vat and the non-envatted counterparts do think the exact same propositions, just with a different truth value: whereas the brains in a vat's thoughts are massively false, the non-envatted brains' thoughts are mostly true, at least when it comes to skepticism. However, these two different claims must be distinguished and kept separate. Besides, an epistemological claim requires an epistemological justification.

Moreover, Pritchard argues (as many others) that Putnam's argument builds on the hypothesis that such brains in a vat have *always* been envatted and controlled by a supercomputer.[31] Therefore, Putnam's argument is valid at best only for those scenarios where this premise is true, and not for all possible forms of BIV skepticism. In reply to this criticism, one could argue that if we got envatted later in our lives, we would have presumably received a sort of brainwashing, so that we wouldn't be able to remember anything at all. If so—and this applies to the other scenarios where we haven't always been BIVs—Putnam's argument still holds. But this is not the case, for there could be a scenario where we have been envatted later in our lives *without* receiving any brainwashing; for, if the virtual simulation matches perfectly the real world, we are not even able to tell the difference. Hence, there is at least one scenario where we haven't always been envatted, and we can refer to real things in the world and even BIVs, even if we *are* BIVs.

To conclude, Putnam's semantic externalism seems rather unhelpful in tackling radical skepticism.

[31] To make it as explicit as possible, if we were envatted later in our lives, we would have experienced the world and so we would be able to refer to actual trees and even BIVs.

3. IBE, Abductivism, and Vogel

Another anti-skeptical strategy is represented by abductivism, which relies on abduction, also called Inference to the Best Explanation (IBE). According to this inferential process, one tries to find, after the observation of a certain phenomenon, its most likely—*best*—explanation through an appeal to explanatory principles, or criteria. The most cited and discussed of these criteria is, undoubtedly, simplicity. Following the abductive reasoning, we are justified in taking a non-skeptical hypothesis as true—usually, our commonsense scenario—because it is the best explanation compared to the radical skeptical hypothesis. So, if such a strategy is successful, we will be able to acquire the rational support needed to favor our belief in an everyday proposition over an incompatible radical skeptical hypothesis, and thus reject the skeptic's claim.

An attempt on this line has been made by Jonathan Vogel. Whereas he argues against some detrimental abductivist approaches, he advances the thought that the commonsense explanation is indeed simpler than the skeptical scenario with respect to their explanatory structures.

3.1 IBE and Abductivism

Abduction is standardly conceived as a third alternative mode of reasoning between induction and deduction (while clearly closer to the former). In valid deductive reasoning, the truth of the premises necessitates the truth of the conclusion. As it is usually put, the truth of the premises "guarantees" the truth of the conclusion. In inductive reasoning, the truth of the conclusion does not necessarily follow from the truth of the premises. Induction is thus an ampliative way of inference, as the conclusion takes us "further" than what its premises by themselves say. The basis of inductive inferences is heterogeneous, but it is usually made of considerations over statistical data.

Abduction is likewise ampliative, but here the conclusion is inferred through an appeal to explanatory considerations, or principles, and not by invoking merely statistical data. While reasoning abductively, we start with collecting some observations—our premises—and then we conceive a hypothesis—our conclusion

—which best explains those observations. This type of inferential process is very important and often utilized in science and scientific methodology.

In fact, IBE is very common in the history of science.[32] An example is Charles Darwin and his theory of evolution. In the book *The Origin of Species* (1859) he mentions a large variety of facts that can be explained by the theory of evolution, but not by the common view of that time, according to which species were created separately from one another by God. He writes:

> It can hardly be supposed that a false theory would explain, in so satisfactory a manner as does the theory of natural selection, the several large classes of facts above specified. It has recently been objected that this is an unsafe method of arguing; but it is a method used in judging of the common events of life, and has often been used by the greatest natural philosophers. (Darwin 1962, 476)

Another example is Antoine Lavoisier and his oxygen theory of combustion. He writes:

> I have deduced all the explanations from a simple principle, that pure or vital air is composed of a principle particular to it, which forms its base, and which I have named the oxygen principle, combined with the matter of fire and heat. Once this principle was admitted, the main difficulties of chemistry appeared to dissipate and vanish, and all the phenomena were explained with an astonishing simplicity. (Lavoisier 1862, 623, translated by Thagard (1978, 77–8))

Lavoisier's theory opposed the then-accepted theory of phlogiston. Roughly, The argument Lavoisier advances is based on the consideration that his theory is able to explain the fact that bodies undergoing combustion increase in weight instead of decrease. Significantly, in order for the defenders of the competing theory to explain the same fact, they have to make extremely odd assumptions—e.g., that the phlogiston has got negative weight. Since the oxygen theory explains the phenomenon in question without making such odd assumptions, Lavoisier infers that it is the best explanation.

In general, IBE can be rendered as follows (Fumerton 1992, 161):

---

[32] For a more detailed analysis see Thagard (1978, 76–8).

1) O (a description of some observed phenomena).
2) Of the set of available and competing explanations E_1, E_2,…, E_n capable of explaining O, E_1 is the best according to criteria C_1, C_2,…, C_n.

Therefore,

3) E_1.

Common explanatory criteria are:[33]

(i) "Ontological Simplicity": Other things being equal, a theory that postulates the existence of fewer (kinds of) entities should be preferred over a theory that postulates more.[34]
(ii) "Explanatory (Consequent) Simplicity": Other things being equal, a theory that raises fewer further explanatory questions should be preferred over a theory that raises more.[35]
(iii) "Explanatory (Primitive) Simplicity": Other things being equal, a theory that posits fewer primitive explanatory notions should be preferred over a theory that posits more.[36]
(iv) "Explanatory Elegance/Straightforwardness": Other things being equal, a theory whose structure is more elegant or straightforward should be preferred over a theory that is less elegant or straightforward.[37]
(v) "Explanatory Power": Other things being equal, a theory that is capable of explaining more things should be preferred over a theory that is not capable of explaining as many things.
(vi) "Conservatism": Other things being equal, a theory that is more in accordance with our previously accepted beliefs should be preferred over a theory that is less in accordance with those beliefs.

---

[33] See Beebe (2009, 609–611) for a more exhaustive list. Here I have taken inspiration from it, revising some criteria.

[34] This is probably the most common abductive criterion cited in the literature. It is often referred to just as "simplicity".

[35] See Harman (1992, 202).

[36] See Lycan (2002, 415). The distinction between consequent and primitive is mine and I have added it in order to distinguish the two different types of explanatory simplicity.

[37] Beebe (2009, 609) classifies this principle as another type of explanatory simplicity. Nevertheless, I believe the elegance or straightforwardness of a theory is not necessarily linked to its simplicity, so I have classified it as a different type.

Ideally, we should always have a set of criteria which need to be balanced against one another. So, in assessing if one hypothesis is a better explanation than other hypotheses we should always utilize several different criteria. The reason is simple: a hypothesis may be much better according to one criterion, but at the same time much worse according to all the other criteria. In this case, we oughtn't to take it as the best explanation, preferring another hypothesis which satisfies better all the criteria considered in harmony.

There are different ways of describing the epistemic strength required for abduction. Usually, what we need is a hypothesis that best explains the phenomenon in question, and also some "good reasons" for considering it as the best. Such a hypothesis will be thus taken as the most probable to be true. However, some put the matter in a stronger way. For example, Gilbert Harman writes:

> In general, there will be several hypotheses which might explain the evidence, so one must be able to reject all such alternative hypotheses before one is warranted in making the inference. Thus one infers, from the premise that a given hypothesis would provide a "better" explanation for the evidence than would any other hypothesis, to the conclusion that the given hypothesis is true. (Harman 1965, 89)

So, not only do we have to provide "good reasons" for selecting one of the given hypotheses, but we are also under the obligation to rule out all the other hypotheses in play. The most straightforward way to interpret the ruling out—or rejecting—of the other hypotheses is that we have to provide grounds for their falsity. This point is not so obvious, since it could well be that, even if one hypothesis is better, being better *per se* does not warrant us to rule out all the others, let alone proving them false. Indeed, some of them might remain in play as reasonable alternatives despite us taking only one.

Moreover, the selected hypothesis must not be *the most probable* to be true, but, simply, true. So, according to this stronger conception of abduction,[38] we have to prove both the chosen hypothesis as true and all the alternatives as false. I believe that this stronger conception might be useful in a scientific context for assessing the

[38] Note that this is not a full account of Harman's views on abduction. I am just using his words for extrapolating a different description of abduction.

relevance of the alternative possible explanations, but does not represent a problem from a mere epistemic point of view. When I have grounds to believe that one hypothesis is a better explanation given certain criteria, thereby being the most probable to be true, I am therefore epistemically justified in accepting that hypothesis, and not the others.

One could object that the truth of the selected hypothesis and the falsity of the alternatives must not be just in some degree probable, but they must be somehow certain. Otherwise, one continues, we wouldn't be really justified to believe the chosen hypothesis. But I object to this. First of all, this would lead us far away from abduction: if we already knew the truth of the hypothesis and the falsity of all the others, what would we even use abduction for? It would be simply superfluous. Secondly, the truth and falsity of a phenomenical proposition are never certain; asking for such a requirement is not epistemically fair nor relevant.

In the philosophical debate, one important role abduction plays is as the ground for a possible anti-skeptical strategy. Following abductive reasoning, one is able to acquire an epistemic justification for embracing a non-skeptical hypothesis, which is explanatory better when compared to the skeptical one(s). In the case of skepticism, we will have the following structure:

1) *O* (a description of our sensory experience).
2) Of the set of available and competing explanations, that is, our ordinary commonsense hypothesis and several skeptical hypotheses, capable of explaining *O*, the ordinary commonsense hypothesis is the best according to criteria C_1, C_2,…, C_n.

Therefore,

3) Our ordinary commonsense hypothesis is, in all probability, true.

According to Beebe (2009, 611), abductivists about radical skepticism must commit to the three following theses:

(1) Belief in the existence of the external world can be epistemically justified via IBE.
(2) Belief in the existence of the external world can *only* be epistemically justified via IBE.

(3) Showing that belief in the existence of the external world is epistemically justified can only be achieved via IBE.

Now, whilst I agree with the first statement, which follows from the definition of abductivism, I prefer to be more cautious with the second, and I find the third one rather unnecessary or, at least, redundant. The reason for this cautiousness is the following. The fact that IBE is a way or can be a way to justify our commonsense beliefs does not exclude that this justification may be achieved in other ways. An abductivist can claim such a justification, without claiming abduction to be the *only* way. Besides, it doesn't really seem necessary for abduction. How could another anti-skeptical strategy, just for being another anti-skeptical strategy, undermine *a priori* abductivism? It seems to be rather odd indeed.

Moreover, the third thesis follows from the second one. If I can epistemically justify the belief in the existence of the external world only by means of IBE, it follows that I can show such an epistemic justification only by means of IBE. I believe that what Beebe means by "showing" is something close to "claiming", to put it in epistemological terms. However, in this circumstance, nothing really changes, and the third thesis still follows from the second.

3.2 Vogel

A prominent proponent of abductivism is Jonathan Vogel. In his well-known paper, *Cartesian Skepticism and Inference to the Best Explanation* (1990), he distinguishes between a "real-world hypothesis" (RWH) and a "computer skeptical hypothesis" (CSH).[39] Whilst the former represents what we ordinarily believe about the external world, the latter represents the possible skeptical scenario of us being brains in vats. Vogel writes:

> The skeptic's position will be empty unless he can provide us with reason to think that a satisfactory competitor exists (in particular, a sufficiently rich competitor that is not unduly burdened with ad hoc explanatory posits). The lesson here is that the skeptic needs to frame an alternative that matches the RWH very closely. If a skeptical hypothesis can be made sufficiently similar in relevant respects to the RWH, then, one might expect, that the skeptical hypothesis will match the RWH in explanatory

[39] The CSH is also called "isomorphic skeptical hypothesis" (ISH) in Vogel (2005).

> adequacy. To the extent that explanatory virtues like coherence, depth, and simplicity are matters of theoretical structure, a skeptical hypothesis that is isomorphic to the RWH will explain things just as well as the RWH does. (Vogel 1990, 660)

In order for the CSH to be a significant alternative to the RWH, Vogel continues, the former hypothesis has to satisfy two constraints: first of all, it must have items corresponding to the elements of the RWH; secondly, it must posit similar properties and relations of these items to those of the RWH. In particular, the CSH relationships among causes and effects must correspond to those in the RWH.

Let us consider the following example. Right now I am working on this master's thesis: I see my hands typing these words on my laptop's keyboard, the piles of books and papers on my desk, and the beautiful landscape outside the window. I see trees, mountains, and houses, all covered in snow and surrounded by a surprisingly sunny bright blue sky. At least, that is what *I think* I see. The skeptic wants to persuade us to believe that all this could be just a simulation and that we are nothing but brains in vats controlled by a supercomputer. However, in order for the CSH skeptical scenario to be compelling it must mimic perfectly the RWH. In particular, the explanatory structure must mirror the RWH's: «there are *some* entities bearing *some* properties that are related in ways exactly analogous to those specified by the RWH» (Vogel 1990, 661).

What makes the difference is that, according to the skeptic, the entities and their properties in question are *different* from the ones in the RWH. Most importantly, what causes my tactile experience of my fingers typing on the keyboard, of my eyes seeing all the things on my desk and the Norwegian nature is nothing but a supercomputer which induces these specific sensations. Also, the sensory experiences are programmed in such a way that the relations among them are the same as in the RWH. For example, the cause of my typing on the keyboard is the cause of my seeing the words I'm typing on the laptop's screen.

We are instinctively tempted to reject such a scenario and, practically speaking, we would never take it into serious consideration. Conceptually, though, things are more complicated; in fact, it is very easy to misunderstand what is wrong with the CSH (if anything). One could argue that we are entitled to set the CSH aside because it has inevitably more unexplained explainers than the RWH (Vogel 1990, 661–2). For example, the CSH will not be able to answer fundamental questions

about where the computer that controls us comes from or why it does so. Nevertheless, the RWH does have unexplained explainers as well—e.g., the RWH will not be able to answer questions about where the universe comes from or why it has the laws it has—so that in the end «both the CSH and the RWH invoke ultimate regularities that are not themselves explained» (Vogel 1990, 662). Hence it is far from clear, Vogel maintains, how the RWH could be better than the CSH.[40]

Moreover (Vogel 1990, 662–3), one could argue that we are entitled to reject the CSH because the RWH is ontologically simpler, and therefore preferable, to the CSH. The latter seems to postulate—at least tacitly—the existence of more entities than the former, e.g., the computer and its central processing unit. However, according to Vogel, ontological simplicity—so understood—can be invoked to also argue the exact opposite, i.e., that the CSH is simpler than the RWH. Whilst the first one is committed to the existence of just the supercomputer and the brains in a vat, the second one assumes many more things to exist: from the trees outside my window to my hands writing this piece of work right now. Secondly, that positing fewer entities is a theoretical and epistemological value is not really straightforward, and could quite easily be opposed. Finally, the skeptic could revise their scenario in such a way that the central processing unit would have no role anymore, in such a way that «the elements of the computer memory act directly on each other, and on the seat of consciousness, in causal patterns that mirror those of the RWH» (Vogel 1990, 662).

Despite all this, Vogel believes that simplicity may still help in tackling the skeptical argument. Whereas the RWH posits different objects with specific spatial properties, so does the CSH, but in a different way: in the CSH objects are just portions of a computer disk, which are then transmitted to our brains so that we "see" them in a simulated world. For example, the trees I see out the window are just portions of a computer disk that are transmitted to my brain. Clearly, in the CSH there is an additional step: what in the RWH is a tree, in the CSH is, first of all, a part of a disk and, then, a (fake) tree.

---

[40] Vogel does not explicitly provide a definition of what he means by "unexplained explainers", nevertheless, it can be easily inferred from the text. To put it explicitly, an unexplained explainer is any ultimate regularity whose purpose is to explain some features of the relevant scenario and which stands on a higher level to the occurring phenomena. This regularity is "unexplained" because it is somehow taken for granted or accepted as it is.

Utilizing one of Vogel's examples, in the RWH something's being spherical explains by itself why that something behaves like a sphere; however, in the CSH what behaves like a sphere is not a sphere but a portion of a computer disk, which does not explain by itself why a piece of metal connected to some wires should behave like a sphere. Therefore, the skeptic must account for this additional feature and thus provide a more extended explanation. In doing so, the skeptic exposes themselves to the risk of postulating a more complicated—less simple—explanatory structure.

Specifically, for each object in the RWH there has to be a corresponding object in the CSH, i.e., a portion of the computer disk. Furthermore, if the RWH assigns a certain propriety to an object, the CSH has to assign a corresponding property to the portion of the computer disk which represents that same object. If a tree in the RWH has the property "location", so that it is placed in one specific point in space and time, the CSH has to posit, for the counterpart object-tree, a property "pseudo location", so that this counterpart object-tree is also placed at that same specific point. Generally, «what the RWH explains by reference to genuine locations, the CSH will explain in terms of these pseudo locations» (Vogel 1990, 664).

Now, the problem for the skeptic is that there is something in the RWH that does not need any explanation, which does need to be explained in the CSH: that, at a specific time, two different objects have two different genuine locations. In fact, that two objects cannot have the same (genuine) location at the same time and place is a necessary truth of our physical world. To put this in more formal terms (Vogel 2005, 77), if *X* has the property of being located at a genuine location *L*, and *Y* is distinct from *X*, then *Y* does not have the property of being located at *L*. But the same cannot be said for the CSH: on a computer disk, we could surely write that two objects have the same property "pseudo location".

In particular, if we conceive this property as consisting of coordinates (x, y, z), which are written in an object's file, we can surely nominally ascribe the *same* coordinates to two *different* objects in their files. Nevertheless, this is not possible since the CSH has to mimic the RWH. Therefore, there has to be a rule—some sort of exclusion principle—which regulates what we can write on the computer disk such that no two objects are assigned the same coordinates in their files, and thus have the same pseudo locations. However, in doing so «the CSH has to add an

extra empirical regularity, to which no regularity in the RWH corresponds. Such an addition will make the CSH inferior to the RWH on simplicity grounds» (Vogel 1990, 665).

Finally, the skeptic could advance the idea that the reason why different objects in the CSH have different pseudo locations is some other necessary truth. The property "pseudo location" would be encoded by a physical property *P*, which two physical objects cannot have at the same time. However, Vogel maintains (1990, 665), there are no such physical properties. Indeed, he maintains, if there is a physical object at one place, there could be an identical object at another place with the same properties except for location.

3.3 Objections to Vogel

I believe several objections can be raised to Vogel's approach. In the next pages, I will address each of them in turn, following the order in which Vogel has stated the relevant thoughts.

(1) To start off, we recall that Vogel argued against the thought that the CSH has more unexplained explainers than the RWH, and the reason was that the latter hypothesis also has some unexplained explainers, so it is not clear how it could be better. But I think that he is too hasty in rejecting such a view. Indeed, the RWH may actually be better—i.e., it may invoke fewer unexplained explainers—for the following reason.

When the skeptic cunningly makes us face their alternative explanation, they are not saying just what they want to say: their statements might have some unintended consequences. In particular, when the skeptic says we may be BIVs and controlled by a supercomputer, they are *inevitably* positing the existence of a physical-external world, even if not the one we—as BIVs—imagine. In other words, they don't commit to just the existence of the supercomputer and the brains in vats. They also commit —logically—to the existence of the world, even if only tacitly. The BIV skeptical scenario has to posit *by definition* the existence of an external world, for otherwise *where* would the supercomputer and those brains *be*?[41]

[41] I will utilize this first consideration in the fifth chapter, where it will be a fundamental step of my own anti-skeptical argument.

Now, the relevance of this point is that one could make use of such a consideration to argue that the RWH is better than the CSH. Basically, in the CSH the physical world might well be different from the world we commonly believe to know (the RWH one), but the differences—if any—are likely to be on the lower level of phenomena and not on the higher level of regularities. This is a move, so they argue, we are entitled to make because the skeptic has committed themselves to it when they have first depicted their scenario. However, the line of reasoning continues, if the CSH has the same regularities as the RWH and then we add the regularity governing the supercomputer, then the CSH turns out to have *more* unexplained explainers than the RWH.

But here an objection may be raised: that we don't know how the real world is in the CSH, so it may have some different regularities compared to the virtual world we have always experienced. Significantly, one can't say to know—nor to be justified in believing—that the relevant differences are on the lower level of phenomena and not on the higher level of regularities. Therefore, such a reply is doomed.

While I recognize the cogency of such an objection, I think the possible approach I have displayed does successfully show that Vogel rejects a relevant possibility—that the RWH has fewer unexplained explainers than the CSH—in a rather superficial way. In addition, I believe that the point I made earlier stands, after all: that the CSH does not only postulate the existence of a supercomputer and the BIVs but also the existence of a physical world.

(2) Moreover, Vogel argues against the view that the RWH is ontologically simpler than the CSH, providing three reasons, as seen above. To his third objection—that the skeptic could revise their scenario in such a way that the central processing unit would have no role anymore—I reply that, even if the central processing unit would have no role anymore, there would still be a supercomputer and the brains in vats too, so it is rather unclear how this revision could be relevant. To his first and second objections—that ontological simplicity can count in favor of the CSH being simpler and that it is not necessarily such a relevant principle—I reply that I disagree with both. Let me explain why.

The starting point, which remains relevant as it is a logical consequence of the skeptic's argument, is that the CSH posits the existence of an entire world—the real one, where the supercomputer and the brains are located. Therefore, it seems

reasonable to say that the CSH posits everything present in the RWH and, additionally, the supercomputer and its machinery. Consequently, the CSH cannot in any way be simpler than the RWH. One could reply, again, that I simply don't know whether the possible world where this supercomputer exists contains precisely everything existing in the RWH. In particular, the world as it exists in the RWH—the one we are experiencing right now—might have many more entities than the world where the supercomputer is (according to the CSH).

Nevertheless, this objection leads us precisely to my next point. It doesn't really matter which scenario has the largest number of entities. It could even be that the CSH is the one with fewer entities, but the RWH would still remain simpler. In order to understand how this can be so, we have to ask ourselves a question: when Vogel talks about ontological simplicity, what is he referring to?

In general, as seen earlier in this chapter, ontological simplicity can refer to basically two things: either the total number of individual ("token") entities or the total number of *kinds* ("types") of entities. To explain the difference, we can apply the concept of ontological simplicity to the theory of elementary particles in physics, comparing two different hypotheses. Whereas on the first meaning of ontological simplicity, we would choose the hypothesis which posits the existence of fewer particles in general, on the second meaning we would choose the hypothesis which posits the existence of fewer kinds of particles, no matter their total number.

I believe no one in the scientific community would choose a hypothesis basing their decision on the first meaning—the total number of particles—but rather, they might do it following the second meaning—the total number of kinds of particles. In fact, a physical hypothesis postulating the existence of a smaller total number of particles does not seem to be necessarily preferable as a matter of simplicity, but one that states the existence of fewer kinds of particles—e.g., two: bosons and fermions—might have much higher chances of being preferable on simplicity grounds.

Now, from a mere epistemic point of view, ontological simplicity understood as the total number of entities seems much weaker than ontological simplicity understood as the total number of *kinds* of entities. In fact, if we refer to the former meaning, the CSH is not necessarily simpler than the RWH nor vice versa, and, anyhow, it wouldn't be of great help in tackling the skeptical argument.

Nevertheless, if we look at the total number of *kinds* of entities, the RWH might actually be simpler than the CSH.

Specifically, we can think of the RWH and the CSH as positing the existence of *n* kinds of entities, depending on what we count as possible kinds. For example, we can think of (physical) objects, minds, mental states, and causation. Alternatively, we could count mental states as falling under the category "minds". Anyway, what would make the difference between the two scenarios in question is the category "(physical) objects". The problem arises when we have a closer look at what constitutes a physical object in the two different hypotheses. In particular, since the CSH posits the existence of also an external physical world, it postulates *necessarily* the existence of two different kinds of (physical) objects, whilst the RWH posits the existence of only one kind.[42]

The explanation of this relevant feature is the following. In the RWH, the only kind of (physical) object is the one made up of all the things that we perceptually experience in the real world. Instead, in the CSH there are two different kinds: the virtual objects—corresponding or not to the RWH physical objects—and the real physical objects outside the virtual simulation—at the very least, the supercomputer, the brains in vats, and some sort of world where these are located. It seems reasonable to maintain that these two kinds of objects are intrinsically different and have different properties. Hence, they ought to be conceived as distinct.

Let me explain it in more detail. The are three ways to conceive the peculiar objects in the CSH simulation. First, from our perspective of BIVs, they are fully and plainly physical, but are at the same time virtual (by construction). Second, from the CSH physical world, these objects are virtual as parts of the simulation, but at the same time physical, as they correspond to physical counterparts (parts of the supercomputer's disk and nervous stimuli). Third, consequently, they are also at the same time physical in two different senses: fully and plainly physical for the BIVs (as we do, in fact, commonsensically think of them) and physical as parts of the supercomputer's disk and electrical stimuli.

I believe this construction leaves room to think that those objects are conceivable as a very peculiar and intrinsic kind. But the CSH, as stated earlier, posits also, by definition, the existence of fully and plainly physical objects—the supercomputer, the

[42] This second consideration, which follows from the first, will be utilized as well in the fifth chapter as an additional fundamental step of my own anti-skeptical argument.

BIVs, and the world where these are located. The problem is, the RWH only posits the existence of fully and plainly physical objects. Therefore, the RWH posits fewer kinds of entities.

One could argue that, if the objects in the simulation are virtual, then they are not physical, and thus they are not entities at all. But, first of all, not being physical does not mean not being *real*. Those objects are, indeed, real, in the sense that they are truly happening—they are *really* located—somehow in our brains (by definition). Secondly, they are, in a certain sense, physical as both parts of the supercomputer's disk and nervous stimuli in our brains.

Significantly though, those objects are at the same time virtual: to those parts of the supercomputer's disk and to those nervous stimuli correspond some virtual objects which are different from the metallic chips and electric impulses per se, and yet they correspond to those metallic chips and electric impulses. Also, they do so in a very specific way: only one particular part of the supercomputer's disk and one particular nervous stimulus will be one particular virtual object.

In addition, these objects are physical for the BIVs in another way as well: they are conceived and perceived as fully and plainly physical (as we do commonsensically think). Therefore, these objects, which are distinct from the fully and plainly CSH physical objects, so I believe, have an intrinsic property which makes them a "new" different kind of objects, and hence an additional kind of entities.

One could legitimately argue against the cogency and relevance of such a distinction. Admittedly, this point is rather debatable. Nevertheless, the very fact that we are facing this problem shows that the skeptic has left some room for interpretation and that their argument is quite unstable. Indeed, I maintain, the skeptic's assertions have some "side effects" the skeptic is not fully aware of. Significantly, the very construction of the skeptical hypothesis leads to this sort of "ontological impasse"; and we are entitled to make such a move because of the way the very skeptic has constructed their alternative scenario.

Finally, ontological simplicity in its second sense—as a matter of how many kinds of entities there are—might be a theoretical and epistemological value that we should take into consideration. In the fourth and fifth chapter, I will deal with this issue.

(3) Now, let us consider Vogel's positive account. Firstly, Vogel might be wrong in pointing out that there is no principle, or regularity, in the RWH such that two objects do not have the same genuine location. What he should have said here is that there is no *explicit* principle or regularity of this kind, but the fact that in the RWH such a principle is implicit—i.e., taken for granted—does not mean it is not present. Indeed, Vogel seems to contradict himself: he states that it is a necessary truth that two objects cannot have the same genuine location at the same time and place, and then he argues that there is no regularity in the RWH for which two objects cannot have the same genuine location at the same time and place. That is why the skeptic would have to include an additional explanation, making the CSH less simple.

Nevertheless, the RWH's necessary truth that two objects cannot have the same genuine location at the same time and place is exactly the regularity that the CSH has to mimic. It is, after all, a principle of the RWH, even if an implicit one. In addition, the fact that we commonly take it for granted does not allow us to say that it does not need any explanation whatsoever. Significantly, if there is such a regularity in the RWH, then the CSH does not have to provide a more complex explanation apparatus and therefore the CSH is not less simple than the RWH. If we want to appeal to simplicity, we ought to find another way, as, apparently, Vogel's approach has failed.

(4) Moreover, there could also be another possible objection. If it's true that the virtual simulation in the CSH mirrors the world in the RWH, the same cannot be said with certainty for the world where the supercomputer and the brains in vats are. Indeed, we don't know anything about it. All we can reasonably presume, as I did, is that the objects have the same properties and relations and that the regularities are the same, but we can't say to know it. Therefore, when Vogel refers to the CSH mimicking the RWH he actually should have referred to the CSH *virtual simulation* mimicking the RWH.

(5) In addition, the statement that there could be two objects which are identical except for their location is not as obvious as Vogel thinks. First of all, two objects are never in actuality *perfectly* identical even without considering their locations. Second, some things might be different simply *by virtue of* their different locations. For example, two homozygous twins who are perfectly similar biologically but who

have lived their first few years (or even just months) of life in two different places will likely develop differences in their behaviour and character.

(6) Finally, Vogel's statement that two objects cannot have the same location at the same time and place seems to be undermined by modern physics. As current research shows, there are some kinds of elementary particles called "bosons" which have the particular feature of being able to share the same location at the same time and place with other bosons.[43] Even if Vogel refers to physical objects of a bigger scale, his general statement seems to be at least problematic.

[43] Bosons are one of the two fundamental kinds of elementary particles (the other being fermions). According to the so-called Standard Model of particle physics, there are five types of bosons: one (the Higgs boson) explains how other particles have mass, and four others (Photons, Gluons, Neutral weak bosons, Charged weak bosons) act as force carriers. For a more detailed description see the webpage "The Standard Model" on the official CERN website (URL=<https://home.cern/science/physics/standard-model>). CERN is the European Council for Nuclear Research (in French *Conseil Européen pour la Recherche Nucléaire*), a leading research institute in the field of particle physics.

4. Abductivism and its critics

In recent years, many authors have raised objections towards abduction, arguing that it does not represent an epistemically rational way of reasoning. In particular, some of them have also argued that abductivism is not an effective anti-skeptical strategy. An important epistemologist arguing along these lines is Richard Fumerton. He has maintained that (1) abduction works only under the debatable premise that the phenomenon we want to explain has actually a possible explanation; that (2) an ambiguity is present in the very notion of best explanation; that (3) even when considering the number of kinds of entities, simplicity leads nonetheless to the problem that some skeptical hypotheses might be simpler than the commonsense explanation; and, finally, that (4) when we consider other abductive criteria—e.g., explanatory power—abductivism doesn't hold water. Another contemporary author to discuss here is Duncan Pritchard, who has opposed both conservatism and simplicity as possible effective abductive criteria against skepticism.

A third author relevant to this debate is James Beebe, who has pointed out several problematic ways of conceiving abductivism. First of all, it is not clear if, and how, the commonsense scenario satisfies the explanatory criteria better than the competing skeptical scenarios. Secondly, these skeptical scenarios may not necessarily have less explanatory power than the commonsense one. Thirdly, there is not one single resolutive argument in favor of the thought that the skeptical hypotheses are less simple in certain respects. Furthermore, some abductivists, while describing the features of our sensory experience, have appealed to something they take to know about the external world, thereby making question-begging presuppositions. Moreover, some abductivists—e.g., James Cornman and Alan Goldman—have provided simply instrumental or pragmatic arguments as opposed to epistemic arguments. In addition, Beebe considers the path some others have taken (e.g., Bertrand Russell), who denied that abduction is capable of providing any epistemic justification—and also raises his own objection, based on Arthur Fine: that abductivists may violate a relevant principle.

Another important author who has argued against abductivism is Bas Van Fraassen.[44] I will discuss some of his well-known arguments against IBE, contained in his book *Laws and Symmetry* (1989). Finally, I will outline another possible objection, which is my own: that the notion of "inference to the best explanation" itself seems ambiguous.[45] That is, a certain explanation may be regarded as the best among several because it is the most probable (to be true), or because it would provide the greatest degree of understanding. In order for IBE to work, this ambiguity has to be resolved.

4.1 Fumerton

In his *Skepticism and Reasoning to the Best Explanation* (1992), Richard Fumerton argues that the abductive argument is effective if and only if it is presupposed that what we want to explain actually has a possible explanation. He writes:

> Unless we have an antecedently justified belief that most things have explanations it is difficult to see why we would take the fact that one potential explanation seems better than the others as a reason for supposing that the potential explanation is true. (Fumerton 1992, 162)

Fumerton maintains that, whereas making such an assumption might be legitimate in the process of scientific enquiry, it seems at least problematic in the context of an anti-skeptical reply.

Moreover (Fumerton 1992, 163–5), he claims that there is a sort of ambiguity in the notion of best explanation. But, unfortunately, the passage where he presents his objection is rather unclear. I shall reproduce it in full below, and discuss it afterward:

> When we talk about one of a number of competing explanations being the best by certain criteria *C*, we may mean one of at least two different things. First we may mean only that the explanation in question is better by criteria *C* than each of the competing explanations taken individually. E_1 might be better than E_2, than E_3, than E_4, and so on.

[44] Beebe deals with Van Fraassen as well, but I have decided to consider him separately both for presentation purposes and because I will discuss his view in more detail, relying on his book.

[45] This sort of ambiguity has nothing to do with the one Fumerton talks about, as we will see shortly.

> This fact by itself does not imply that even if we are justified in believing that one of these explanations is correct, we are justified in believing that E_1 is the correct explanation. It is true that if we are forced to "bet" the rational bet would be E_1. If we are forced to make plans based on some explanation, E_1 is the rational explanation to choose. But being the rational explanation to act on is not the same as being the rational explanation to believe. […] I can cheerfully admit that E_1 is the most attractive explanation while admitting that the disjunction of propositions asserting alternative explanations is more likely to be true. In such a situation I stress again that there may be nothing irrational about acting as if E_1 is the correct explanation, particularly if one must do something and one must make plans based on some particular explanation. But as an epistemic agent one certainly doesn't have to fallaciously infer that E_1 is likely to be the correct explanation. (Fumerton 1992, 164)

I think the best way to make this argument intelligible is the following. When we say that among several possible explanations one is the best given a certain criteria C, we could mean two different things. Firstly, we could mean that, given C, one explanation is better than each of the other *individual* explanations. That is, considering a set of explanations $\{E_1, E_2, E_3, \ldots, E_n\}$, E_1 is better than E_2, and all others up to E_n. Expressed in probabilistic terms,[46] this means that $P(E_1) > P(E_2)$, $P(E_1) > P(E_3)$, $P(E_1) > \ldots$, $P(E_1) > P(E_n)$. But, secondly, we could also mean that, given criteria C, one explanation (e.g., E_1) is better than *all* the other explanations taken together—i.e., a sort of disjunctive explanation; or as Fumerton says above, "the disjunction of propositions asserting alternative explanations".

Considering again a set of explanations $\{E_1, E_2, E_3, \ldots, E_n\}$, this means that E_1 is better than $E_2 \vee E_3 \vee \ldots \vee E_n$. In particular, this would mean that $P(E_1) > P(E_2 \vee E_3 \vee \ldots \vee E_n)$. Such distinction is very important because, if we mean that one explanation is "better" in this second sense, the disjunction of alternative propositions $E_2 \vee E_3 \vee \ldots \vee E_n$ is typically going to be more likely true than E_1 alone, i.e., it is likely that the numerical values of probabilities are such that we have that $P(E_1) < P(E_2 \vee E_3 \vee \ldots$

[46] Fumerton does not make his point in these terms, but I think that my reconstruction captures its meaning.

$\vee E_n$).[47] Then, such a disjunctive explanation—*if* we accept it as a genuine explanation, which is far from clear—will likely be the best one, since its probability is likely higher than that of E_1.

Furthermore (Fumerton 1992, 165–7), considering simplicity as a possible criterion leads to the problem that some skeptical hypotheses might be simpler than the ordinary real-world hypothesis (RWH), even when we consider the *number of kinds* of entities. As Fumerton says, «we can compare our commonsense hypothesis about physical objects to […] a Berkeleyean hypothesis about a very complex mind orchestrating the comings and goings of sensations» (1992, 165). Whereas a Berkeleyean-style skeptical hypothesis would have minds, mental states, and causation as postulated entities, an ordinary hypothesis would have all these three kinds *plus* physical objects. So, it is not so straightforward to say if, and how, commonsense beliefs do satisfy the abductive criterion of simplicity better than the skeptical scenarios. In addition, it is by no means evident that, other things being equal, simpler hypotheses are more likely to be true than more complex ones.

Finally (Fumerton 1992, 167–9), we could consider other abductive criteria, such as explanatory power. Nevertheless, the explanation provided by a Berkeleyean-style skeptical hypothesis, where minds are the causes of our sensations, has the same explanatory power of—i.e., explains as much as—the commonsense hypothesis where physical objects are the causes of our sensations. As Fumerton says:

> Explanations involving minds as the causes of our sensations seem to explain just as much as explanations involving physical objects [as causes], however, […] we at least have experience of our own minds producing mental states to rely on in understanding how some other mind might be able to produce our sensory images. (Fumerton 1992, 168)

---

[47] The reason behind this claim is the probabilistic (*a priori*) truth that $P(E_i \vee E_j) = P(E_i) + P(E_j) - P(E_i \cap E_j)$. So, when $P(E_i \cap E_j) = 0$—i.e., when two explanations are probabilistically independent, or, roughly put, do not have anything to do with each other—the probability of the disjunction of the alternative explanations will be the sum of their probabilities, i.e., $P(E_i \vee E_j) = P(E_i) + P(E_j)$. But this means that the probability of the disjunction of the alternative explanations will likely be higher than the probability of a single explanation, i.e., $P(E_i \vee E_j) = P(E_i) + P(E_j) > P(E_i)$ (or $P(E_j)$), since probabilities are always positive. In my reconstruction of Fumerton's argument, it is the sum $(P(E_2)+P(E_3)+\ldots+P(E_n))$ that, I take, he says it is likely to be higher than $P(E_1)$.

Now, let me answer Fumerton's objections. To the first objection, my reply is that it is indeed already presupposed, in a certain sense, that there is a possible explanation for our sense perceptions of an external world. Nevertheless, I believe this assumption is not really problematic for two reasons. First of all, it seems rather self-defeating to think that such an assumption is false: without it, basically everything in the scientific field would be worthless to study. Moreover, why should we regard, in this specific context, the scientific enquiry as massively different from the anti-skeptical enquiry? Indeed, if we accept this objection, not only the abductive response will fail, but every anti-skeptical reply will be doomed from the very beginning. Secondly, note that the skeptic makes the very same assumption, so the skeptical explanation would be undermined as well. That is, the skeptic too believes that there is an explanation for the fact that we have the sense perceptions that we have, the only difference is that they explain this fact in terms of the BIVs scenario.

To the second objection, I reply that, even though Fumerton's point is cogent, we shouldn't consider the concept of a better explanation in his second sense. To start off, the list of alternative explanations could be as long as we want, even infinite. In fact, we could always add one more explanation to the list (the disjunction), even just out of our imagination.[48] For this reason, the second sense seems to be epistemically weaker—since it involves some arbitrariness—and, therefore, the first one ought to be preferred. In addition, one can ask when it is really the case that competing explanations of the same phenomenon do not have *anything* to do with each other.[49] In fact, they will likely share many elements. For example, the BIV-skeptical explanation shares several elements with the RWH, e.g., the existence of our brains.

To the third objection, my reply is that to consider simplicity as a possible criterion does not mean to consider it as the *only* possible criterion. We should always consider several different criteria working together to decide which is the best explanation.[50] In addition, the radical skeptical scenario of BIVs is different from the "Berkeleyean-style" skeptical hypothesis. They posit different elements and they

---

[48] I don't see any problem in adding an explanation just because we imagine it, for it would still have its own probability and it would still represent a possible state of affairs.

[49] This condition is needed since it assures that the probability of a disjunction is higher than the probability of each disjunct. See fn. 47 above.

[50] See Beebe (2009, 609–11) and the third chapter for examples of other possible criteria.

require different considerations. Moreover, the idea that, other things being equal, a simpler hypothesis is more likely to be true than a more complex one, could in fact be justified. I will discuss a possible line of argument in more detail in the next chapter.

Finally, to the last objection, I reply that explanatory power is not the only other possible criterion to consider. I stress, once again, that we should always consider several criteria, and that we should consider the BIV-skeptical scenario specifically.

4.2 Pritchard

Another author who has very recently rejected abductivism is Duncan Pritchard (2015, 25–8). He first attacks the abductive principle called conservatism, according to which a belief that is more in accordance with our established set of beliefs is more epistemically grounded than another belief that is less in accordance with this set of beliefs. Pritchard argues that such a principle may be rational in our everyday scientific practice, but it appears as significantly problematic when applied to radical skepticism. The reason is that, in a radical skeptical context, the very same set of beliefs may be dramatically mistaken, so we can't rely on it to justify another belief. Indeed, since we can't distinguish a radical skeptical hypothesis from the commonsense hypothesis, we don't know whether our ordinary beliefs are true or not.

He then argues against simplicity as an anti-skeptical abductive principle. For him, to put it crudely, simplicity is in the eye of the beholder: someone might think of the skeptical scenario as actually simpler than the everyday scenario, especially if simplicity is understood as ontological simplicity.[51] In addition, as for conservatism, simplicity is relative to a previously acquired set of empirical beliefs, so when we consider the skeptical alternative that those very same beliefs are mistaken, how could those beliefs—and hence simplicity—provide us an epistemic justification for choosing the commonsense hypothesis over a skeptical one?

In my opinion, Pritchard is right in rejecting conservatism as a possible abductive criterion. Even if it is not completely useless, conservatism is epistemically very weak in the context of an anti-skeptical reply to the BIVs scenario. And yet, I believe

[51] The ontological simplicity of a theory is here viewed as depending only on the number of entities needed by that theory and not on the number of *kinds* of entities postulated by the theory.

that Pritchard is wrong—or at least too hasty—in rejecting simplicity. First of all, as we have seen in Chapter 3, ontological simplicity ought to regard the number of *kinds* of entities and not just the number of entities in general. Thus, if understood correctly, the application of ontological simplicity in the BIV-skeptical scenario might indicate that the RWH is in fact simpler. Moreover, simplicity is not necessarily relative to our previously acquired set of empirical beliefs: it might be possible to conceive simplicity as an *a priori* abductive principle. Hence, it might be possible to acquire epistemic justification. I will propose such an *a priori* account of simplicity in the following chapter.

4.3 Beebe

In *The Abductivist Reply to Skepticism* (2009), James Beebe discusses several problematic ways of conceiving abductivism. To start off (Beebe 2009, 613–4), the most obvious problem is that it is not clear if, and how, the commonsense scenario —or RWH—satisfies the explanatory criteria better than the skeptical scenarios. For example, the skeptical hypothesis of the BIVs, which posits the existence of only a supercomputer and some brains in vats, could seem ontologically simpler than the real-world scenario. In addition, this skeptical hypothesis may not necessarily have less explanatory power than the commonsense one. Abductivists argue that, even if skeptical hypotheses could be simpler in certain respects, they are definitely less simple in others.[52] However, Beebe maintains, none of these arguments seem resolutive.

Moreover, Beebe (2009, 615–9) discusses two ways the abductivists have used to defend their position, which he thinks are actually detrimental. Some abductivists have made question-begging presuppositions in describing the features of our sensory experience, which the commonsense hypothesis has to explain. In particular, while describing those features, they have appealed to something they take to know about the external world. Nevertheless, as he points out, it is not possible to defend the RWH by appealing to any consideration which assumes the truth of some of our beliefs about the external world, for this would be a circular argument. For example, as stated by Michael Williams (1999, 49), when abductivists argue that our sensory data have a sort of coherence that is explainable through the

[52] An example of this view is Vogel. See Chapter 3.

RWH, they do so by means of an appeal to some regularities in the appearance of physical objects with similar conditions of perception.

The second problematic way to defend abductivism consists of providing simply instrumental or pragmatic arguments as opposed to epistemic arguments. As Beebe writes:

> Because epistemological skepticism is the thesis that we lack epistemically justified beliefs about the external world, any response to this thesis that tries to show that our external world beliefs are merely pragmatically, instrumentally or otherwise nonepistemically justified will fail to address the central skeptical challenge. (Beebe 2009, 618)

What makes these arguments nonepistemic in nature is their reliance on explanatory criteria that are not themselves truth-linked. Indeed, Beebe argues, in order for some explanatory criteria to provide us with an *epistemic* reason, they must be truth-linked. In fact, an objection the skeptic could rise is that simplicity cannot be a discriminating factor for truth, since simplicity and truth are actually independent from each other. Simpler does not necessarily and by itself means truer.

For example, James Cornman (1980) believes that IBE is more consistent with «pretheoretical, commonly and firmly accepted beliefs» (Cornman 1980, 15) than any other competing principle. Abduction represents a better principle for making more «effective, efficient, and […] moral decisions to guide us through life» (Cornman 1980, 10). Thus, for its role in helping us in making decisions and achieving certain practical goals, he thinks it is reasonable to utilize IBE. Nevertheless, Beebe maintains, not only Cornman's practical justification is not enough if we have to face a skeptical challenge, but our accepted beliefs are exactly what the skeptic is challenging and, therefore, we cannot appeal to them.

Furthermore, Beebe (2009, 619–20) discusses Goldman's defense of abductivism. In his *Empirical Knowledge* (1988), Goldman offers a sort of inductive justification in place of an *a priori* justification. He starts from the view that «all epistemically interesting connections between beliefs and the world that ultimately verifies or falsifies them are contingent» (Goldman 1988, 7). All the *a priori* conceptions, he states, fail to respect the contingency of this connection and the logical independence of physical reality from our thought. He then offers an inductive

justification, appealing to the survival advantage of those who use abductive reasoning.

First, he argues that it is possible to reconstruct how the ability to use IBE has been selected for the survival of our species, showing its contribution to it. Second, he takes the explanatory coherence of this evolutionary account and uses it to support the truth-conduciveness of abduction reasoning. Unfortunately, in Beebe's opinion, Goldman's argument is plainly circular, since he «uses IBE to infer the truth of RWH and then justifies IBE using theories premised on the truth of RWH» (Beebe 2009, 619). Goldman himself was well aware of this circularity, but he found it not so problematic, for circularity, so he maintains, is inevitable: every principle must be justified by other principles, but ultimately basic principles cannot be justified through an appeal to even more basic principles.

Moreover, Beebe considers another possible path some abductivists have taken, with Bertrand Russell among them. That is, to deny that there is any good reason to believe that abduction is capable of providing us with an epistemically justified reason whatsoever. In Russell's opinion, «there is no very good ground for supposing that a simple law is more likely to be true than a complicated law, though there is good ground for assuming a simple law in scientific practice, as a working hypothesis» (Russell 1927, 132). Believing, he continues, in the existence of an external world is and has to be antecedent to any evidence. Such assumption is practically justified in the context of scientific enquiry, but represents a mere prejudice—that is, amounts to a flawed theory—in the context of skepticism.

In reply, I will briefly note here that, if our goal is to provide an epistemic justification for IBE-based beliefs, then we should contend that there is a connection between simplicity and truth, but only *a priori*. This link is what might *epistemically* justify believing in the truth of the RWH. I will try to find such an *a priori* connection in the following chapter.

Moreover, another possible objection is raised by Beebe himself (2009, 625–6), based on Fine (1984). Basically, abductivists may violate the following principle:[53]

[53] Fine's original principle (Fine 1984, 85) is the following: Metatheoretic arguments must satisfy more stringent requirements than those placed on the arguments used by the theory in question. Beebe decided to conceive his own principle because he doubted «both the cogency and the clarity of [Fine's] general principle» (Beebe 2009, 626). I have called it (EC) for presentation purposes.

(EC) In the context of a debate about the epistemic credentials of a particular belief source one should not appeal to a belief source whose credentials are far more widely in dispute than the source one seeks to defend.

Defending the belief in the truth of the RWH by means of IBE is problematic because the epistemic reliability of our sensory data as a source of belief for the existence of the external world is much less questionable than the epistemic reliability of IBE. In fact, Beebe argues, no one really doubts the RWH, but many doubt IBE. As admitted by Beebe himself, this objection is not decisive at all, but it shows that there seems to be «something unsatisfying about defending RWH» (Beebe 2009, 626).

My reply to Beebe's objection is the following. First, the fact that many object to IBE but not so many object to our sensory data should not count for too much in this debate. There are many different ways to conceive IBE, and if many raise objections, it may certainly be because a really satisfying account of IBE has still not been found. But what if we succeed in finding one? The principle would not tell against abductivism anymore, or at least not so easily.

Secondly, appealing merely to the amount of debate around certain belief sources does not tell us anything about their credibility. It could happen that one thing we strongly believe as true turns out to be completely false and another thing we have highly discussed for decades—or even centuries—is finally proved right. Finally, that not many question our sensory data might be true for an everyday or even scientific context, but in *our* context the reliability of our sensory data is precisely what is at stake.

4.4 Van Fraassen

In his book *Laws and Simmetry* (1989), Van Fraassen raises several objections towards IBE. One of his arguments against is the argument of the "bad lot":

> We can watch no contest of the theories we have so painfully struggled to formulate, with those no one has proposed. So our selection may well be the best of a bad lot. To believe is *at least* to consider more likely to be true, than not. So to believe the best explanation requires more than an evaluation of the given hypotheses. It requires a step beyond the comparative judgment that this hypothesis is better than its actual rivals […]. For me to take it that the best of set *X* will be more likely to be true than not,

> requires a prior belief that the truth is already more likely to be found in *X*, than not. (Van Fraassen 1989, 143)

So, given that we can only grasp a part of all the possible hypotheses, we cannot acquire a justification for believing that the true hypothesis is among the small amount we know. Therefore, IBE is undermined.

According to Van Fraassen (1989, 143–6), there are three possible ways to try to escape the problem. The first way—which he calls the "reaction of Privilege"—is to claim a sort of privilege for the humans' intellect, since humans are by nature predisposed to hit on the right range of hypotheses. The second way consists in arguing that we must choose among the given hypotheses, for we have no other alternative. Accordingly, the reason why we make the choice is not that we have some special or peculiar beliefs, as for the first reply, but because we *must* choose from that "bad lot". Van Fraassen opposes both of them and, since I agree with him, I will not report his argumentation. Indeed, it is rather easy to see why they don't work for science, let alone for any anti-skeptical strategy. Briefly: that we are "by nature predisposed" to hit on the right range of hypotheses, has to be demonstrated, and, thus far, it has been in fact proved wrong many times; that we must choose among the given hypotheses because we have no other alternative, might work practically in our daily life, but it's hard to see how it could become a rational epistemic principle.

The third way to escape the problem could be to re-elaborate IBE so that «it is not the rule to infer the truth of the best available explanation [..., but it is the rule] to allocate our personal probabilities with due respect to explanation» (Van Fraassen 1989, 146). Accordingly, explanatory power is a mark of truth; it might not be infallible, but it is a characteristic symptom. There are two forms of implementing this view: the first one claims that what makes a hypothesis the best explanation is its being (more) likely to be true; the second one claims that the notion of rationality itself requires these features to function as relevant factors in the rules for a rational response to the evidence. Whilst the first one «must lean on *intrinsic explanatoriness*, which can be discerned prior to empirical observations», the second one «[must lean] specifically on explanatory success after the observational results come in» (Van Fraassen 1989, 146).

Against the first form, Van Fraassen advances the following argument:

> I believe, and so do you, that there are many theories, perhaps never yet formulated but in accordance with all evidence so far, which explain at least as well as the best we have now. Since these theories can disagree in so many ways about statements that go beyond our evidence to date, it is clear that most of them by far must be false. I know nothing about our best explanation, relevant to its truth-value, except that it belongs to this class. So I must treat it as a random member of this class, most of which is false. Hence it must seem very improbable to me that it is true. (Van Fraassen 1989, 146)

So, recalling the previous argument of the bad lot, not only am I unjustified in believing that the true hypothesis is among the few ones I selected, but the hypothesis I believe to be true is actually most likely false. One way to answer this objection, Van Fraassen argues, is to claim that «we do have further knowledge of our own best explanation, relevant to its truth-value, beyond how well it explains». «I'm afraid», he continues, «that this will bring you back to the reaction of Privilege, to glory in the assumption of our natural or historical superiority» (Van Fraassen 1989, 146). Significantly though, he doesn't go any further in this respect.

Moreover, Van Fraassen deals with the second form in the seventh chapter of his book (Van Fraassen 1989, 151–182). Since it would take us much space and it would not be so relevant to our discussion after all, I have decided to not address it here.

Now, I believe a possible answer to Van Fraassen might be found, indeed, in his words. He maintains that a way to save IBE is to view it as the rule to "allocate our personal probabilities", and that a possible form of conceiving such an allocation is the claim that a hypothesis is the best explanation—among others—because it is (more) likely to be true. This sort of likeliness must have the nature of an *intrinsic explanatoriness*, discernible prior to empirical observations. What he objects against this view is, simply, that it would imply the claim of a sort of privilege for the humans' intellect, assuming our natural or historical superiority.

He shows two ways of defending such a claim (Van Fraassen 1989, 143–4): on one side, naturalism, which defends it on the basis of our adaptation to nature and of our evolutionary success that must be due to a certain fitness; on the other side, rationalism, which appeals to beliefs about God—e.g., that we are made in his image—thereby arguing that it is only reasonable to believe that we are peculiarly adapted to discover the truth when we come up with our guessed explanations.

Clearly, there might well be other ways to defend that claim, other than evolution and God. I believe another possible way consists of discerning *a priori* the (higher) probability of the best explanation. If we will be able to provide such an account, not only will we be able to provide a relevant objection to Van Fraassen's own argumentation, but we will also find a significant help in our anti-skeptical quest.

4.5 A final objection

Finally, let me discuss another possible objection to abductivism. The notion of "inference to the best explanation" itself seems ambiguous, in a way that is different from Fumerton's: a certain explanation could be regarded as the best among several because it is the most probable (to be true), or because it would provide the greatest degree of understanding. In fact, those two different views of "best explanation" may lead to choosing different explanations as "the best" one. In other words, an explanation which is the most probable to be true may not be the one that provides the greatest level of understanding, and vice versa.

Consider the following example. Someone reads in the newspaper that a Japanese person committed suicide and, not knowing anything else, one starts wondering what was the reason. The most common cause of suicide is by far depression, so one might be justified in thinking that the explanation which is most likely to be true is that the person who committed suicide was depressed. Nonetheless, given the Japanese culture of shame, another explanation could be advanced: that this person committed suicide because they were feeling ashamed for something. So, one might be at the same time justified in thinking that the explanation which provides the highest degree of understanding is, everything considered, that the person committed suicide because they were ashamed for something they did.

This example is not supposed to be decisive, but I hope it will help to shine some light on the distinction introduced above. Indeed, I have found it particularly difficult to present a cogent example. I am aware that the explanation dealing with depression may also be the one providing the highest degree of understanding and, vice versa, the one dealing with the Japanese culture of shame may be the most likely to be true. However, I believe the example still helps us to highlight the difference.

I believe that some criteria, especially simplicity, deal first of all with the truth-probability of certain explanations. For instance, a simpler explanation is supposed to be, first of all, more probable to be true; moreover, a simpler explanation does not necessarily provide a better understanding. Other criteria, e.g., explanatory power, deal mostly with the degree of understanding provided by certain explanations. Some other criteria, conservatism for example, look at both the truth-probability of some explanations and the level of understanding they provide. Yet, in the context of an anti-skeptical reply, we should prioritize truth-probability over understanding. Although things could be different, and surely more complicated, in the context of scientific enquiry, here we are facing theoretical scenarios where the degree of understanding different hypotheses provide is roughly the same. What makes the difference, first and foremost, between the BIV scenario and the commonsense scenario (RWH) is the likelihood of one scenario compared to the other. Indeed, what the skeptic is objecting to us is that we are not entitled to believe in the truth of our commonsense beliefs about the eternal world, since it's underdetermined.

5. *A priori* simplicity and the skeptic's defeat

As should be clear from the fourth chapter, the only possible appeal to abduction we are left with, especially when facing skepticism, is and has to be *a priori*. Indeed, Bertrand Russell (1912, 21–2; 1927, ch.18) argued that an abductive response to skepticism cannot appeal to anything we think we know about the external world. He writes:

> We must therefore, if possible, find, in our own purely private experiences, characteristics which show, or tend to show, that there are in the world things other than ourselves and our private experiences. (Russell 1912, 22)

But, as we saw above, this is, in Russell's opinion, (epistemically) impossible. In fact, on several occasions, I have referred to a hypothetical *a priori* account of abductivism, particularly of simplicity. Let us recall those occasions. First, to Fumerton's objection that it is by no means evident that, other things being equal, simpler hypotheses are more likely to be true than more complex ones, I have replied that this idea could be actually justified. Second, to Pritchard's objection that simplicity is relative to a previously acquired set of empirical beliefs, as for conservatism, I have replied that simplicity is not necessarily relative to our previously acquired set of empirical beliefs and that it might be possible to conceive simplicity as an *a priori* abductive principle. Third, to Russell's objection, I have argued that if we want to find an epistemic justification through IBE, we should conceive an *a priori* connection between simplicity and truth. Finally, to Van Fraassen, I have replied that a way to save abduction is to discern *a priori* the (higher) probability of the best explanation.

In general, the particular kind of abductivism which appeals to an *a priori* justification is usually called "rationalist abductivism". According to this view, there are some abductive principles, simplicity *in primis*, which provide an *a priori* justification to take a certain explanation as the best among several others. In the context of the skeptical challenge, this strategy aims at providing an *a priori* justification to reject the radical skeptical hypothesis and to believe in the truth of the commonsense one.

In what follows I will first discuss some suggestions made by some of the authors we have seen in the fourth chapter, in particular Fumerton and Beebe. I will then critically present a defender of rationalist abductivism, that is, Laurence BonJour. After that, I will present my own rationalist abductivist approach, based on an account of *a priori* simplicity. Such an account will provide, together with some additional considerations we have encountered in the third chapter, an epistemic justification to positively answer the skeptical challenge. Finally, I will assess its anti-skeptical strength and answer some possible objections.

5.1 Some suggestions

A first hint has been given by Fumerton himself (1992, 167). He states that «there is at least one plausible *a priori* appeal to simplicity», adding:

> Suppose that I start out with two hypotheses, H_1 and H_2 and that relative to my evidence E, H_1 and H_2 are equally likely. Suppose further that I acquire some additional evidence which requires me to either abandon H_2 or *add* to it an additional hypothesis T_1 while H_1 remains fine as it is. Clearly, the conjunction (H_2 and T_1) has a probability lower than the probability of the conjunct H_2, and if the probability of H_2 is the same as the probability of H_1, it follows that the probability of the more complex hypothesis (H_2 and T_1) is lower than the probability of H_1. But all this presupposes that we started out with H_1 and H_2 having equal probability relative to the original evidence base. There is nothing in any of this which suggests that other things being equal simpler hypotheses are more likely to be true than complex hypotheses. (Fumerton 1992, 167)

Unfortunately, this passage seems to be put in the text as a brief suggestion, and nothing else is said about this kind of *a priori* appeal. Moreover, I believe it is not so clear how the last sentence of the passage is related to the rest of the quoted text. My guess is that Fumerton here wants to say that, although this *a priori* appeal seems reasonable, it is premised on the fact that the two competing hypotheses have the same basic probability of being true. However, it is very hard to say whether the RWH has the same basic probability of a skeptical competing scenario because it is very hard to say if any additional T is added to any of the two basic hypotheses. Therefore, we are not able to compare the two competing scenarios in

the context of an anti-skeptical discussion. I believe that Fumerton's suggestion might be helpful, even if I disagree with his statement that there is nothing which suggests that simplicity is not truth-linked. It might be true that in this elementary formulation—the one he has provided—simplicity is, in fact, not a mark of truth-likeliness. Nevertheless, if we could develop more the conception of an *a priori* appeal to simplicity, we may manage to escape the problem.

Beebe (2009, 623) suggests that any justification for abductivism must satisfy at least the following constraints:

(1) The justification must be epistemic.
(2) The justification cannot be circular.
(3) The justification cannot be transmitted from any justified belief about the external world, since the justification of this entire class of beliefs is called into question by the skeptic.
(4) The justification must be capable of underwriting inferences from beliefs about our sensory experiences to beliefs about the external world.

Only an *a priori* form of justification, Beebe continues, seems to be able to satisfy these criteria. Moreover, he writes:

> Abductivists who wish to go the *a priori* route must face the challenge that practically everyone who works on abductive inference believes that such inferences are justified empirically and that the theoretical virtues are broadly empirical and contingent marks of truth. Rationalist abductivists must also formulate and defend—in the face of strong opposition—a conception of probability according to which statements of probability are necessary and can be known *a priori*. Finally, they must also find a way to handle formidable challenges to IBE that have been raised in recent years. (Beebe 2009, 625)

These formidable challenges to IBE that Beebe is referring to are some of the objections we have seen in the fourth chapter. He also maintains that even if rationalist abductivists think that RWH is more likely to be true than the skeptical hypotheses, they also acknowledge that two claims are nonetheless true (Beebe 2009, 628):

(1) It is both metaphysically and epistemically possible that RWH is false.
(2) It is both metaphysically and epistemically possible that we inhabit a world where simpler (deeper, broader, less *ad hoc*, etc.) explanations are true less often than more complex (shallower, narrower, more *ad hoc*, etc.) explanations.

Still, rationalist abductivism maintains that the RWH, despite these two claims, is more probable than the skeptical competing hypotheses and we are justified in believing its truth.

5.2 BonJour

In the contemporary debate, Laurence BonJour (1985, 1998, 1999, 2003) is one of the defenders of rationalist abductivism. He has offered his own *a priori* account of abductivism based on a foundationalist view.[54] According to BonJour (2003, 92), we can *a priori* see that the best—i.e., the most likely—explanation is that the causes of our sensory data are those physical objects reflected in our sensory experiences. He argues that the spatial features of our sensory experience are «such as to be easily and naturally explainable by supposing that they are systematically caused by a relatively definite world of mostly solid objects arranged in three-dimensional space, and by no other hypothesis that is not […] essentially parasitic on that one» (BonJour 2003, 88). These features are:[55]

(i) The continuity that exists between the varied sensory experiences we can have of an object through one sensory modality.
(ii) The coordination between the contents of experiences from different sensory modalities.
(iii) The regular, repeatable and unified sensory experiences we have of different objects as we move through space.

---

[54] According to the view called "Foundationalism", our knowledge rests upon some non-inferential justified beliefs which do not depend on anything else and which act as a foundation for our further justifications and beliefs. For a more detailed overview, see Hasan's and Fumerton's entry "Foundationalist Theories of Epistemic Justification" in *The Stanford Encyclopedia of Philosophy* (Fall 2022 Edition), URL=<https://plato.stanford.edu/archives/fall2022/entries/justep-foundational/>.

[55] See Beebe's reconstruction (Beebe 2009, 607–8).

(iv) That features of our sensory experiences are correlated in just the way we would expect them to be if some of the things we experience were the causes of other things we experience.
(v) That families of sensory experiences and their relations to each other seem to change over time in ways that intuitively reflect changes in both the experienced objects and the movements of observers.

A possible explanation of these experiential patterns is what he calls (BonJour 2003, 92) the "*quasi-commonsensical* hypothesis", according to which our sensory experiences are caused by a realm of three-dimensional objects, having at least approximately the same shapes, spatial relations, and causal properties as the objects we perceive in our sensory experience.

He then further proposes a distinction between two general kinds of explanations (BonJour 2003, 93). The first kind of explanation, which includes the quasi-commonsensical hypothesis, is called the "*analog* explanation". Such an explanation accounts for the features of our sensory experience by appealing more or less directly to the basic characteristics of the objects in the hypothesized world. These objects are directly reflected or represented in our sensory experience.

The second kind of explanation, which includes all the skeptical hypotheses, is called the "*digital* explanation". This second type of explanation makes sense of our sensory experience with a conjunction of two things: first, a representation of the world depicted in the quasi-commonsensical hypothesis; second, a more or less abstruse machinery which produces in those who perceive it an experience that mirrors the experience that they would have if the depicted world was actually real, though it is not true that this world exists, for the true features of the world are utterly different from the ones we perceive in our sensory experience.

BonJour maintains that the quasi-commonsensical hypothesis is the best among the analog explanations. The reasons he provides are the following:

> (1) It is very hard to see how the spatial features of experience could be explained in analog fashion by anything other than a spatial world: certainly a multi-dimensional world of some sort seems needed to account for the various sorts of experiential paths that return to the same experiential starting point. (2) For essentially the same reason, a two-dimensional world does not seem to have enough internal complexity to

> account for all of the possible experiential sequences and variations. (3) Thus an alternative analog hypothesis would arguably have to involve a world of at least three dimensions containing objects whose shapes and relations differ systematically from those that are actually reflected in our experience, with the character of experience thus involving something like a systematic distortion. (BonJour 2003, 93–4)

BonJour continues by noting that if the differences between the quasi-commonsensical hypothesis and any other analog explanation are large enough to be significant, the characteristics of the quasi-commonsensical hypothesis which make it work really well in explaining our sensory experience will make any other analog explanation not work so well, or even at all. He doesn't go further in explaining those differences, stating that «no one has ever actually described such a possible alternative in even an approximate way, and here it seems to me quite reasonable to place the burden of proof on those, if there are any, who seriously believe that such analog alternatives genuinely exist» (BonJour 2003, 94).

Obviously, there are still many other *digital* explanations that could explain as well as the quasi-commonsensical one, or even better. Nevertheless, BonJour argues that such digital explanations are, explanatory speaking, worse than the quasi-commonsensical explanation. He provides two main reasons (BonJour 2003, 94–5). First of all, there seems to be something rather arbitrary about it, partially reflected in the underdetermination of the relevant competing explanations. As he says, it is «intuitively unreasonable to adopt such a further, arbitrary mechanism until and unless there is some specific feature of the data to be explained that requires it» (BonJour 2003, 95).

Secondly, he provides the following argument:

> The explanatory success of a digital explanation of experience depends in effect on the truth of two claims: first, that the corresponding analog explanation could indeed account for the experience in question (that a material world could produce experience of the sort that we in fact have), since the digital explanation works by emulating the action of the cause or causes postulated by the analog explanation; and, second, that the specific translating mechanism postulated by the digital explanation in question can indeed successfully do the job of emulation (that God or the computer can indeed systematically produce the sort of experience that would be produced by the represented material world). But the explanatory success of the

> analog, quasi-commonsensical hypothesis depends on the truth of only the first of these claims. (BonJour 2003, 95)

Hence, finally, he concludes that there is an *a priori* good reason to believe that the quasi-commonsensical explanation is the best possible explanation and we are *a priori* justified in accepting it.

As BonJour himself acknowledges (2003, 95–6), there are still three problems with his view. First, even if the quasi-commonsensical hypothesis might be more likely to be true than any of the digital alternatives, it is not so obvious that it is also more likely to be true than the disjunction of all the possible alternatives, nor it is clear whether it satisfies our commonsense intuitions or the requirements for knowledge. Second, we are at best justified to attribute to physical objects only the characteristics required to explain our sensory experience (e.g., spatial relations and causal properties), but these characteristics do not include the secondary qualities of our sensory experience (e.g., colors). Third, when ordinary people make claims about the physical world they do not think of anything close to the quasi-commonsensical hypothesis and the argument in its favor. Consequently, most ordinary people will not have a full and explicit justification for making such claims.

I believe some objections may be raised against BonJour's approach. Granted his distinction between analog and digital explanations, and granted his thought that the quasi-commonsensical hypothesis is the best (and probably the only possible one) among the analog explanations, I think it's not so straightforward to see how the quasi-commonsensical hypothesis could be better than all the digital competing explanations.

The first reason he provides, which to me resembles some of Pryor's remarks above, might be problematic. Even if there might be something arbitrary—*ad hoc*, I would say—in the skeptic's move, still being "intuitively unreasonable" to accept the skeptical point—that we could be BIVs—is not enough if what we want is an *epistemic* reason. It might work on a practical or pragmatic ground, but how relevant is it on an *epistemic* ground? Moreover, in our anti-skeptical context, where our quasi-commonsensical explanation is underdetermined, being reasonable to take it as better *is* what has to be (epistemically) proved. Indeed, the skeptic could certainly reply to BonJour that, since we are in an underdetermined situation, unless he

provides a relevant good reason to reject the competing skeptical explanation, there is no intuitive unreasonableness in such an explanation.

In fact, the second reason BonJour provides is an attempt to provide a relevant reason, which nonetheless, as I understand it, seems to me rather problematic. There is, to start off, a misunderstanding at the base of his argument: he implicitly takes the digital hypothesis to depend on the quasi-commonsensical one, whilst in actuality it is the exact opposite. First, he says that any digital explanation undersigns the claim that the corresponding analog explanation could indeed account for the experience in question. The problem is that *of course* it does; it follows by definition and there is no claim to undersign: any digital explanation is constructed in such a way that the ones who are in that scenario do think to be in the quasi-commonsensical scenario instead. Significantly, the quasi-commonsensical hypothesis depends on the digital hypothesis in question, and not vice versa. For, those inside the virtual simulation would just think to be in the real world, no matter how the virtual simulation is constructed.[56]

Second, the supercomputer's mirroring mechanism does successfully its job of emulation because there is no other way. It can't but systematically produce the experience that would be produced by the represented material world, for it's written in its program. Nevertheless, again, this is not because the digital hypothesis depends on the quasi-commonsensical hypothesis, but rather, it's true the opposite.

BonJour's argument resembles, indeed, Vogel's. So, it could be possible to save such an argument relying, as Vogel does, on the explanatory structure. I believe that that was BonJour's attempt after all, but, in the way he puts the matter, it is rather unclear how the quasi-commonsensical hypothesis is supposed to be better than all the digital explanations.

Now, I want nonetheless to defend BonJour's account from the three further issues he has raised, for they could represent some possible objections to any *a priori* abductivist account. To the first issue, I reply that, following what I have already said regarding Fumerton in the fourth chapter, the disjunction of alternative explanations ought not to be considered a genuine explanation. Since the number of alternatives, as stated by BonJour himself, is infinite, we can always add or remove arbitrarily some explanations, thereby changing the possible probabilistic outcome.

[56] They wouldn't know how the (true) real world may be outside the virtual simulation!

Moreover, whether such an approach satisfies the requirements for knowledge will be dealt with later in the chapter.

To the second issue, I reply that secondary qualities are epistemically either unnecessary or (at worst) irrelevant, for they more or less depend on the subjects' individualities. Consider the example of colors: most humans view the same colors, but some others—e.g., daltonic people—don't. What for many of us is green, and a specific kind of green, for daltonic people is nothing but a shade of grey. Significantly, this does not depend on any consideration of the characteristics *required* to explain our sensory experience, so they are not pertinent factors in assessing the explanations. Finally, to the third issue, my reply is that, in short, what ordinary people do or think is not our concern here. Besides, just as BonJour admits, the argument is at least in principle available to them. That, practically, most people wouldn't know it or think about it does not mean *per se* that, epistemologically, the argument doesn't work or it's not cogent.

5.3 *A priori* simplicity

In contrast to both Vogel and BonJour, I will now present my own approach to abductivim. It will be a rationalist approach and its core will be an *a priori* account of (ontological) simplicity. Earlier in the thesis, I have several times argued that the principle that, other things being equal, a simpler hypothesis is more likely to be true than a more complex one, can in fact be justified. As we have seen, such a principle has to possess an "intrinsic explanatoriness", which can be discerned prior to any empirical observation. In other words, it must be truth-linked, and it must be so *a priori*. In addition, we have also seen that it has to be able to discern *a priori* the (higher) probability of the best explanation.

So, the basic idea is reminiscent of Ockham's Razor, namley that a hypothesis which posits a smaller number (of kinds) of entities is *a priori* probabilistically more likely to be true than a hypothesis postulating a larger number. In other words, the more kinds of entities are postulated, the more claims about their existence are made, and the less likely will be that all these claims are true. To say this as clearly as possible, even if just *one* existential claim about one of these kinds of entities is false, then the hypothesis is false (since the hypothesis is a conjunction). So, the more a conjunctive hypothesis posits, the more it is exposed to failure.

The development of this idea consists of employing basic probability theory. Roughly put, we have that the probability of a conjunction—e.g., $P(A_1 \wedge A_2 \wedge A_3 \wedge \ldots A_n)$—is always less, or at most equal, to the probability of any of the conjuncts A_i, where A_i is a proposition postulating the existence of an entity of kind i. In fact, if all these propositions A_i are probabilistically independent of each other, the probability of their conjunction is the product of the probabilities of the conjuncts, i.e., $P(A_i \wedge A_j) = P(A_i) \times P(A_j)$.[57] And, since probabilities are sub-unitary, this product is less than the probability of any of the conjuncts. Accordingly, since a theory positing fewer (kinds of) entities has a higher probability to be true than a theory positing more (kinds of) entities, we are justified—in fact, *a priori* justified—in considering an explanation to be better when it posits fewer (kinds of) entities.

Let us now imagine two competing hypotheses H_1 and H_2. Moreover, whereas the first one posits the existence of (say) four kinds of entities K_1, K_2, K_3, and K_4, the second one posits the same four kinds and an additional fifth kind K_5, different from the previous four. In the case of H_1, its probability of being true can be calculated by taking into account (i.e., depends upon) 4 variables, $P(A_i)$, with i = 1 to 4, where $P(A_i)$ = the probability that a kind of object K_i exists: $P(H_1) = P(A_1) \times P(A_2) \times P(A_3) \times P(A_4)$.[58] In the case of H_2, its probability of being true can be calculated by taking into account (i.e., depends upon) five variables, P(Ki), with i = 1 to 5. So, $P(H_2) = P(A_1) \times P(A_2) \times P(A_3) \times P(A_4) \times P(A_5)$.[59] Now, since probabilities are sub-unitary numbers, it is clear that $P(H_1) > P(H_2)$. These considerations follow from the theoretical calculus of probabilities, hence they are *a priori*. Note that although it may be that H_2 and not H_1 is the true hypothesis, H_1 remains the one with the highest *a priori* probability of being true.

Therefore, I shall suggest the following *a priori* account of ontological simplicity:

---

[57] When independence doesn't hold, we have that $P(A_i \wedge A_j) = P(A_j) \times P(A_i \mid A_j)$, where the second factor of the product is a *conditional* probability, of A_i given A_j. So, the probability of the conjunction is still less or at most equal to the probability of each conjunct.

[58] I assume, to streamline the presentation, that the existence of the four kinds is probabilistically independent. This is not an implausible assumption, although I admit it needs further elaboration.

[59] As for H_1, the existence of the five kinds has to be probabilistically independent as well.

"*A Priori* Ontological Simplicity": Whenever, while considering two or more competing hypotheses, one hypothesis posits the existence of fewer kinds of entities compared to the number posited by the other(s), we are *a priori* justified in taking the hypothesis which posits the existence of fewer kinds of entities as true.

Importantly, it seems that my account does satisfy Beebe's requirements. The justification it provides is epistemic, it is not circular, it is not transmitted from any justified belief about the external world, and is capable of underwriting inferences from beliefs about our sensory experiences to beliefs about the external world. Moreover, my account possesses an "intrinsic explanatoriness", which can be discerned prior to any empirical observation. Also, not only is it *a priori* truth-linked, but it does discern *a priori* the (higher) probability of the best explanation. Finally, I believe it does provide a pure and full epistemic justification: "pure" because, being *a priori*, it does not and cannot involve any practical or pragmatic considerations; "full" because, if such an account is successful, it does not provide only a some-degree kind of justification, but a full-degree justification.

Let us recall the relevant objections left untackled in the fourth chapter. First, to Fumerton's objection, I reply that my *a priori* account of simplicity shows that, other things being equal, simpler hypotheses are more likely to be true than more complex ones. Second, to Pritchard's objection, I reply that my *a priori* account of simplicity is not relative to a previously acquired set of empirical beliefs, as for conservatism.

Third, to Russell's objection, I reply that my *a priori* account of simplicity does provide an *a priori* connection between simplicity and truth. Finally, to Van Fraassen, I reply that my *a priori* account of simplicity does find a way to discern *a priori* the (higher) probability of the best explanation.

5.4 The end of the quest?

Let us consider now the two competing scenarios RWH and CSH, as in Vogel's. In the third chapter, I have advanced two considerations which I have utilized to argue against Vogel's arguments. These two considerations are, briefly, that (1) the CSH posits, even if just implicitly, the existence of an external physical world; and that (2) since (1), the CSH posits the existence of more kinds of entities, and, therefore, it is ontologically less simple. Thus, it may well be that the skeptic is right, but it is,

according to my account, *a priori* more likely to be true that the RWH is the actual scenario we are in. Thus, are we *epistemically* justified in believing in the RWH? Yes, I believe we are.

To put my argument in more formal terms:

(1) The BIV skeptical hypothesis posits, implicitly but necessarily, the existence of an external world.

Therefore,

(2) The BIV skeptical hypothesis posits the existence of more kinds of entities than the competing commonsense hypothesis.

Hence,

(3) The BIV skeptical hypothesis is less (ontologically) simple than the competing commonsense hypothesis.

Thus, since we are *a priori* justified in rejecting a less (ontologically) simple hypothesis,

(4) We are justified in believing in the commonsense hypothesis.

But if so,

(5) Radical skepticism is defeated.

According to such an anti-skeptical argument, we are able to find some rational support to favor our belief in an everyday proposition over the incompatible radical skeptical hypothesis, tackling the UP-based skeptical argument. Moreover, that rational support will also provide us justification to (claim to) know the denials of radical skeptical hypotheses, thereby tackling also the CP-based skeptical argument.

At this point, one may worry that, after so much analysis of this topic, my approach may sound less sophisticated and refined than others. Nevertheless, for our purposes—i.e., finding a radical anti-skeptical strategy—I do believe that my account provides a new approach (even if based on known ideas), which is, at the very least, rather interesting, and, hopefully, even successful. Therefore, I now turn to some possible objections.

5.5 Some possible objections

I will now discuss some possible objections.

(1) Some may argue—the skeptic in particular—that, even granted its success, the justification that my account provides is not fully and purely epistemic; in fact, it may not even be a justification at all. Appealing to a sort of *a priori* probabilistic ontological simplicity is too weak an argument, and doesn't really amount to a move against skepticism. The truth about the existence of the external world, so they argue, cannot be justified by such a "superficial" approach, for there are too many factors which this account doesn't consider, and which might change the outcome.

In reply, I will say that the fact that there is an external world simply follows logically from the skeptic's own argument. Again, even the simplest BIV scenario will posit by definition the existence of an external world: the supercomputer and the brains in a vat have to be physically located somewhere after all. I don't see how this basic idea could rationally be challenged. Therefore, what is truly at issue is not the existence of an external world *per se*, but the existence of *our* world, as we know it. What we need is a justification for believing that the external world is reasonably how we commonly think it is. I believe that my account provides such a justification.

(2) Some may object that, even granted that my account works for BIV skepticism, it might not work as easily and straightforwardly for other types of radical skepticism (or even not work at all!). For example, how is it supposed to work for the classic Cartesian case of the Evil Genius? In this case, I believe, we wouldn't even need any account of simplicity: we are epistemically justified in rejecting *a priori* the evil genius scenario because the argument in its favor posits the existence of a metaphysical/mythological creature whose existence is completely outside the realm of epistemology. The skeptic, in this case, is not epistemically legitimized in presenting such an argument and we don't owe them any epistemic answer. However, this is not the place for an extensive discussion about the evil genius scenario.

In addition, as stated in the fourth chapter with respect to the "Berkeleyean-style" skeptical hypothesis, every skeptical hypothesis ought to be analyzed separately, for it might require some peculiar considerations. This thought is, indeed, consistent with abductivism. I do think my account of ontological simplicity might be useful, or even resolutive, for tackling all the skeptical hypotheses, but someone may still

come up with a different scenario where neither simplicity nor abductivism work. So, I will nonetheless concede the point and rest content with arguing that my account raises a successful reply to only BIV skepticism.

(3) Another objection is that it is not so straightforward that my account provides a justification which is not transmitted from any justified belief about the external world. Admittedly, this point seems problematic, but not so much to completely disqualify my account. Upon reflection, it can be said that my account does indeed utilize some kind of justified beliefs about the external world in two senses. First, when I argue that the BIV skeptical argument posits the existence of an external world, even if just implicitly; second, when I state that that very argument postulates the existence of two kinds of entities: the virtual objects and the non-virtual ("real") ones, i.e., the objects outside the simulation. In reply, I maintain that, even if these two points do in fact involve the external world, they are inferred directly from the skeptical argument and not from the external world or our sensory experience. In other words, they are justified because of the logic of the skeptic, not because of what we perceive.

(4) Another objection, probably the most compelling, is that even granted my *a priori* account of simplicity, it is not true that the objects in the CSH simulation are conceivable as an *ontologically* different and additional kind of entities. So, in essence, this is an objection to how we count the relevant entities. The problem, the arguer may say, is that I have gone too far. In fact, the arguer may continue, it is not even clear how to interpret such an ontological distinction. I am aware of the cogency of such an objection, but I still want to reply to it. First, the problem is that to be rather unclear is first and foremost the skeptic's construction of their scenario. Indeed, the skeptic's statements—as I maintained several times—have some unintended consequences the skeptic wasn't fully aware of. My attempt has been to make these collateral implications finally explicit. Second, this very discussion highlights the fact that the skeptic has left some room for interpretation and that their augment is, after all, near to an argumentative short circuit. Significantly, so I believe, the way the skeptic has constructed their hypothesis leads to both this ambiguity and this sort of "ontological impasse". Accordingly, we are entitled to make such a move because of the way the very skeptic has constructed their alternative scenario.

(5) Let us now consider another objection. This criticism stems from the fact that knowledge does not come from justification *per se*. In other words, having justification for believing that *p* does not imply that we know that *p*. In particular, justification for believing that *p* by itself is not sufficient for possessing knowledge that *p*. Consider the traditional analysis of knowledge as justified true belief (Gettier 1963, 121):

S knows that *p* if and only if:
(i) *p* is true;
(ii) *S* believes that *p*; and
(iii) *S* is justified in believing that *p*.

Or in a more sophisticated form (BonJour 2010, 23):

In order for *S* to know that *p* at time *t*:
(i) *S* must believe or accept *p* at *t* without any doubt;
(ii) *p* must be true;
(iii) *S* must have at *t* a reason or justification that guarantees that *p* is true.

Accordingly, in order to know that *p*, not only do we have to be justified in believing that *p*, but *p* must be true. Thus, we might even be justified in believing that we are not BIVs, but we still don't know whether we are or not. Such lack of knowledge is precisely what undermines our success.

However, things may not be so simple for the objector. Even granted the traditional definition of knowledge provided above, which is a significant concession considering all the criticisms raised to it throughout the years, it defines possessing knowledge and does not say anything about when we can *claim* to have knowledge. To *have* knowledge and to be in the position to *claim* to have knowledge are two different things. Arguably, having knowledge is stronger than claiming to have knowledge: we could claim to know that *p* even without truly having knowledge that *p*—and this may happen because *p* is false and we ignore such falsehood. Hence, we can still claim to know that we are not BIVs even without knowing that it is false,

and that claim, or *right*, comes from the *a priori* abductive justification we have acquired. The truth of the BIV scenario *per se* exceeds both us and epistemology.

(6) Finally, one could argue that the abductivist approach requires, by definition, the application of several criteria and not just simplicity, which is only one of the possible abductive principles. Accordingly, an abductivist approach based on only simplicity will be much weaker than an abductivist approach based on several criteria. The challenge here is to provide an account for each possible criterion which can satisfy the relevant constraints. Not only must the justification be *a priori* truth-linked, but also it has to be fully and purely epistemic, not circular, not derived from any justified belief about the external world, and it has to underwrite inferences from beliefs about our sensory experiences to beliefs about the external world.

However, even if we could find such an account, still the commonsense hypothesis about the external world may not satisfy the criteria in question. Let's recall the other abductive principles I have introduced in Chapter 3:[60] explanatory (consequent and primitive) simplicity, explanatory elegance/straightforwardness, and explanatory power. At first glance, the only criterion that looks promising is explanatory primitive simplicity. With regards to explanatory elegance/straightforwardness, it is not clear at all how to provide an *a priori* account of it, nor how a more elegant/straightforward hypothesis is supposed to be more likely to be true. On the other hand, explanatory power and explanatory consequent simplicity do share the same problem: even granted an *a priori* account, the things a hypothesis is capable of explaining (or the further explanations involved) are subject to significant changes according to our amount of knowledge. As we have seen, even just reflecting on a hypothesis may require further explanations involved in it, or postulated things that that hypothesis is able to explain. Therefore, I don't see any way to discern, based on these two principles, whether the commonsense hypothesis or the BIV hypothesis is a better explanation.

Explanatory primitive simplicity may be saved because the primitive explanatory notions instrumental in a hypothesis are easier to grasp and understand. Also, they are not subject to any change, for they are presupposed like the axioms of a theory. Now, could an *a priori* account based on probability work for this criterion as well? Is

[60] I am excluding conservatism, for, as stated in Chapter 4, I embrace Pritchard's criticism on its weakness as an anti-skeptical principle.

it possible that a hypothesis that posits less primitive explanatory notions is probabilistically more likely to be true? I leave this issue for further work.

Conclusion

Let's recapitulate. In this thesis, I have dealt with a specific kind of philosophical skepticism, called radical or Cartesian skepticism. First, I have first presented the argument in its two major forms: the one relying on the Closure Principle and the one based on the Underdetermination Principle. I have also argued that the second form is the one to be preferred. Second, I have analyzed some of the best-known anti-skeptical strategies: Moore and Wittgenstein, Williams and Wright, Pryor, and, finally, Putnam. Third, I have presented the particular anti-skeptical strategy of abductivism. I have outlined abduction and abductivism in general; then, I have discussed Vogel's own abductive approach, raising several objections against it. Fourth, I have examined some of the most important criticisms of the abductive approach, trying to rebut them. Finally, I have considered the particular kind of abductivism called "rationalist abductivism". I have addressed the suggestions made by some authors and I have looked closely at Laurence BonJour's position. After that, I have presented my own *a priori* account of simplicity, based on probabilistic grounds. In light of this account and some considerations already advanced in the third chapter, I have also proposed my own anti-skeptical argument. Finally, I have addressed some possible objections.

Bibliography

Audi, Robert. 1988. *Belief, Justification and Knowledge*. Belmont, CA: Wadsworth.

Beebe, James. 2009. “The Abductivist Reply to Skepticism.” *Philosophy and Phenomenological Research* 79(3): 605–636.

BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

——. 1998. *In Defense of Pure Reason*. Cambridge: Cambridge University Press.

——. 1999. “Foundationalism and the External World.” *Philosophical Perspectives* 13: 229–249.

——. 2003. “A Version of Internalist Foundationalism.” In BonJour and Sosa (2003), 3–96.

——. 2010. *Epistemology: Classic Problems and Contemporary Responses* (2nd edition). Lanham, MD: Rowman & Littlefield.

BonJour, Laurence and Ernest Sosa. 2003. *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Malden, MA: Blackwell.

Brueckner, Anthony. 1986. “Brains in a Vat.” *Journal of Philosophy* 83: 148–67.

Child, William. 2011. *Wittgenstein*. New York: Routledge.

Coliva, Annalisa. 2015. *Extended Rationality: A Hinge Epistemology*. London: Palgrave-Macmillan.

Comesaña, Juan. 2017. “On Sharon and Spectre’s Argument against Closure.” *Philosophical Studies* 174(4): 1039–1046.

Conant, James. 1998. “Wittgenstein on Meaning and Use.” *Philosophical Investigations* 21: 222–250.

Cornman, James W. 1980. *Skepticism, Justification, and Explanation*. Dordrecht: Reidel.

Darwin, Charles. (1962). *The Origins of Species*, edited by George G. Simpson. New York: Collier.

DeRose, Keith. 1992. “Contextualism and Knowledge Attributions.” *Philosophy and Phenomenological Research* 52(4): 913–929.

——. 1995. “Solving the Skeptical Problem.” *The Philosophical Review* 104(1): 1–52.

DeRose Keith and Ted A. Warfield (eds.). 1999. *Skepticism: A Contemporary Reader*. New York and Oxford: Oxford University Press.

Descartes, René. 1986. *Meditations on First Philosophy*, edited and translated by John Cottingham. Cambridge: Cambridge University Press.

Dretske, Fred I. 1970. "Epistemic Operators." *The Journal of Philosophy* 67(24): 1007–1023.

Fine, Arthur. 1984. "The Natural Ontological Attitude." In *Scientific Realism*, edited by Jarrett Leplin, 83–107. Berkeley: University of California Press.

Fumerton, Richard. 1992. "Skepticism and Reasoning to the Best Explanation." *Philosophical Issues* 2: 149–169.

Gettier, 1963. "Is Justified True Belief Knowledge?." *Analysis* 23(6): 121–123.

Goldman, Alan. 1988. *Empirical Knowledge*. Berkeley: University of California Press.

Harman, Gilbert. 1965. "The Inference to the Best Explanation." *Philosophical Renew* 74: 88–95.

——. 1992. "Induction: Enumerative and Hypothetical." In *A Companion to Epistemology*, edited by Jonathan Dancy and Ernest Sosa, 200–206. Cambridge, MA: Blackwell.

Huemer, Michael. 2001. "The Problem Of Defeasible Justification." *Erkenntnis* 54(3): 375–397.

Jenkins, Caroline S. 2007. "Entitlement and rationality." *Synthese* 157: 25–45.

Lavoisier, Antoine. 1862. *Oeuvre* (II). Paris: Imprimerie Impériale.

Lycan, William G. 2002. "Explanation and Epistemology." In *The Oxford Handbook of Epistemology*, edited by Paul K. Moser, 408–433. Oxford: Oxford University Press.

McGlynn, Aidan. 2014. "On Epistemic Alchemy." In *Scepticism and Perceptual Justification*, edited by D. Dodd and E. Zardini, 173–189. Oxford: Oxford University Press.

Moore, George E. 1925. "A Defence of Common Sense." In *Contemporary British Philosophy* (2nd series), edited by J. H. Muirhead. London: Allen and Unwin.

——. 1939. "Proof of an External World." *Proceedings of the British Academy* 25: 273–300.

Moyal-Sharrock Danièle. 2016. "Wittgenstein Today." *Wittgenstein-Studien* 7(1): 1–14.

Moyal-Sharrock Danièle and William H. Brenner (eds.). 2005. *Readings of Wittgenstein's On Certainty*. Basingstoke: Palgrave Macmillan.

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

Pedersen, Nikolaj J. 2009. "Entitlement, Value and Rationality." *Synthese* 171(3): 443–457.

Pritchard, Duncan H. 2005. "Wittgenstein's On Certainty and Contemporary Anti-Scepticism." In Moyal-Sharrock and W. H. Brenner (2005), 189–224.

——. 2011. "Wittgenstein on Scepticism." In *The Oxford Handbook on Wittgenstein*, edited by O. Kuusela and M. McGinn, 522–549. Oxford: Oxford University Press.

——. 2012. "Wittgenstein and the Groundlessness of our Believing." *Synthese* 189(2): 255–272.

——. 2015. *Epistemic Angst: Radical Skepticism and the Groundlessness of our Believing*. Princeton and Oxford: Princeton University Press.

——. 2016. "Putnam on BIVs and radical skepticism." In *The Brain in a Vat*, edited by S. Goldberg, 75–89. Cambridge: Cambridge University Press.

——. 2021. "On hinge epistemology." *Sképsis* 12(23): 1–17.

——. 2022. "Hinge Commitments and Common Knowledge." *Synthese* 200 (3): 1–16.

Pryor, James. 2000. "The Skeptic and the Dogmatist." *Noûs* 34(4): 517–549.

——. 2004. "What's Wrong with Moore's Argument?." *Philosophical Issues* 14: 349–378.

——. 2014. "There Is Immediate Justification". In *Contemporary Debates in Epistemology* (2nd edition), edited by Steup, Matthias, John Turri, and Ernest Sosa, 202–221. Chichester: Wiley Blackwell.

Putnam, Hilary. 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.

Russell, Bertrand. 1912. *The Problems of Philosophy*. Oxford: Oxford University Press.

——. 1927. *The Analysis of Matter*. London: George Allen & Unwin.

Sandis, Constantine and Danièle Moyal-Sharrock (eds.). 2022. *Extending Hinge Epistemology*. Anthem Press.

Sharon, Assaf and Levi Spectre. 2017. "Evidence and the Openness of Knowledge." *Philosophical Studies* 174(4): 1001–1037.

Thagard, Paul R. (1978). “The Best Explanation: Criteria for Theory Choice.” *The Journal of Philosophy* 75(2): 76–92.

Van Fraassen, Bas C. 1989. *Laws and Symmetry*. Oxford: Clarendon Press.

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

———. 1993. “Dismissing Skeptical Possibilities.” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 70(3): 235–250.

——. 2004. “Skeptical Arguments.” *Philosophical Issues* 14: 426–455.

———. 2005. “The Refutation of Skepticism.” In *Contemporary Debates in Epistemology*, edited by Steup, Matthias and Ernest Sosa, 72–84. Chichester: Wiley Blackwell.

Williams, Michael. 1991. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Oxford: Blackwell.

——. 1999. “Skepticism.” In *The Blackwell Guide to Epistemology*, edited by John Greco & Ernest Sosa, 35–69. Maiden, MA: Blackwell.

Wittgenstein, Ludwig. 1969. *On Certainty*, edited by G. E. M. Anscombe and G. H. von Wright, translated by D. Paul and G. E. M. Anscombe. Oxford: Blackwell.

Wright, Crispin. 1985. “Facts and Certainty.” *Proceedings of the British Academy* 71: 429–472.

——. 1992. “On Putnam’s Proof that We Are Not Brains in a Vat.” *Proceedings of the Aristotelian Society* 92: 67–94.

——. 2004. “Warrant for Nothing (and Foundations for Free)?.” *Proceedings of the Aristotelian Society* 78(1): 167–212.

——. 2014. “On Epistemic Entitlement (II): Welfare State Epistemology.” In *Scepticism and Perceptual Justification*, edited by D. Dodd and E. Zardini, 213–247. Oxford: Oxford University Press.