REFLECTIVE LUCK AND META-EPISTEMOLOGICAL SCEPTICISM

Charles Anthony Neil

A Thesis Submitted for the Degree of MPhil
at the
University of St Andrews

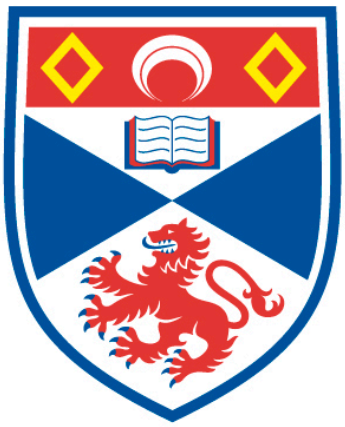

2014

Reflective Luck and Meta-Epistemological Scepticism

A dissertation submitted in requirement of the degree of

MPhil (Research)

Department of Philosophy

University of St Andrews

October 2013

Charles Anthony Neil

Supervisors: Dr Philip Ebert and Dr Patrick Greenough

Examiners: Professor Katherine Hawley and Professor Stephen Hetherington

Contents **Page**

Abstract

In this thesis, I argue that a particular type of epistemic luck, called "reflective luck", motivates a meta-epistemological challenge to externalism about justification. I argue that the meta-epistemological challenge consists in a substantive philosophical challenge to externalism that entails the rejection of a naturalized epistemology. However, I contend that the philosophical challenge to the externalist analysis of justification should be tempered with an anti-sceptical intuition that we *do* have knowledge of putatively true propositions. To this end, I argue that an externalist analysis of justification is best able to accommodate our anti-sceptical intuitions; externalism, I argue, is the best way of accommodating commonsense.

Although externalism is preserved and survives the meta-epistemological challenge, it is not unscathed. Specifically, I contend that externalism deprives us of adequate internalist epistemic grounds to think that we have externalistically justified beliefs. I identify that in principle this is not a problem for externalism, because externalism can respond in a number of ways (one way is to abandon the concept of justification as essential to knowledge), but that nonetheless an adequate epistemology ought to do justice to the legitimacy of the meta-epistemological challenge. However, accommodating this legitimacy and preserving putative knowledge is not possible within the traditional absolutist framework for doing epistemology. My conclusion is that externalism is correct, but that the kind of externalist knowledge we have is not especially fine-grained or perfect; this should put pressure on the absolutist framework for doing epistemology.

Preliminary Notes

The chapters are divided into several sub-sections for the sake of clarity; although chapter 5 has only one sub-section. Chapter 1 is largely introductory, and is an attempt to get to grips with the concept of luck, and its varieties, and to sketch some ways in which reflective luck is used by internalists to motivate the meta-epistemological challenge to externalism. Chapter 2 is exploratory, although I do argue that mentalism is not a viable form of internalism. The subsequent chapters 3 and 4 outline what I take to be the main challenges for externalism and internalism, respectively; these chapters form the core of the thesis. In chapter 3, I outline "the philosophical objection" to externalism, and evaluate internalist alternatives – I argue that meta-epistemological scepticism will also

afflict anti-sceptical forms of internalism. In chapter 4, I propose an anti-sceptical form of externalism that keeps everyday knowledge whilst preserving epistemic closure. However, I also outline the limitations of this externalist knowledge. By this stage, I should have motivated two intuitions that will make us sympathise with two diametrically opposed accounts of knowledge. Chapter 5 proposes a theory of knowledge that accommodates both of these intuitions.

Chapter 1: Knowledge, Luck, and Reflective Luck

This dissertation is to do with knowledge and justification. However, I discuss justification only insofar as it is a condition on which knowledge (or at least the best sort of knowledge) depends. I am working on the assumption that we should value justification because of its potential to give us knowledge, and that the goal of having justified beliefs is to have knowledge or, at the very least, true belief. (Although I will remain quiet on these issues, it seems a justified false belief is less valuable than a justified true belief, on the grounds that the false belief can't qualify as knowledge.)

However, the idea that all propositional knowledge *depends* on justification is something of a presupposition. This presupposition is derived from the popular idea that *mere* true belief cannot count as knowledge because it is compatible with a lucky guess, which, the intuition goes, falls short of knowledge.[1] (However, the conclusion of this thesis may be used to challenge the traditional presumption, since it emerges that true belief that satisfies some relevant externalist condition is compatible with a particular type of imperfect externalist knowledge that I advocate.)

We should also note that many philosophers think there are different types of knowledge human beings and animals may possess. Quite often a distinction is made between *propositional knowledge* (knowledge-that) and *procedural knowledge* (knowledge-how). The former propositional knowledge involves ascriptions that are typically preceded by a "that" clause, of the form *I know that I am a student,* or *that the sun will rise,* and so on. Propositional knowledge thus asserts that some fact obtains; it is therefore the type of knowledge that scientists are primarily interested in. In contrast, procedural knowledge, if there is such a thing, is non-propositional; one may have this sort of knowledge if one *knows how* to ride a bicycle, for example. Some philosophers deny the legitimacy of a distinction between propositional knowledge and procedural knowledge, but I will not explore this issue here.[2] In this dissertation, I am only interested in propositional knowledge (this will be abbreviated to "knowledge that *p*", or sometimes, "knowing *p*", but whenever I talk about knowledge, propositional knowledge is what I have in mind).

[1] I discuss the intuition that knowledge excludes luck in this chapter.
[2] For a refutation of this distinction, see: Jason Stanley and Timothy Williamson, "Knowing How", (2001), *The Journal of Philosophy*, Vol. 98, No. 8, pp. 411-444

Likewise, whenever I talk about how justification relates to knowledge, I am discussing its relation only to propositional knowledge.

In this chapter, I show that the platitude that knowledge excludes luck can be extended to show that justification – understood in terms of having grounds to believe some proposition *p* - is insufficient for knowing *p*. To capture the intuition that knowledge excludes luck we need to make knowledge depend on some external condition: the subsequent status of this external condition and its ramifications for knowledge is the primary focus of this thesis.

I start chapter 1 by introducing the notion of luck and examining ways in which epistemic luck may undermine knowledge. To this end, I advocate responding to *veritic luck* by adopting an externalist counterfactual modal requirement. However, I show that externalist requirements of any sort still fail to eliminate a further type of epistemic luck, called *reflective luck*. However, reflective luck is an altogether more elusive notion, and so I attempt to define what it is and how it might be eliminated, and, most importantly, the meta-epistemological challenge that reflective luck presents for externalist justificatory requirements.

Epistemic Luck

A very general schema for when some condition C is lucky may be represented as follows:

Condition C is lucky in the actual world α if C obtains α, but fails to obtain in a close possible world β[3]

In the above schema, "C" may denote *any* condition. However, luck has received a great deal of attention within epistemology, where "C" is often substituted for some condition upon which knowledge (that *p*) depends, such as truth or belief.

It is relatively uncontroversial to say that knowledge depends on certain conditions obtaining. For example, we may say that knowing proposition *p* depends, in part, on believing *p*. More often, epistemologists talk of knowledge depending on a conjunction of conditions – so that knowing *p* depends on believing *p and* being justified in believing *p, and p* being true, and so on. The question of what conjuncts knowledge depends on is contested; as is the question of

[3] Let α represent the actual world and β represent some close possible world – where the closeness of a possible world to the actual world is determined by its similarity to the actual world.

whether we can in principle provide a suitable conjunctive analysis of knowledge.[4] However, most accept that knowledge depends on certain positive conditions obtaining - by positive conditions I mean conditions such as having a belief and being justified.

However, epistemologists will often further insist that even satisfying some conjunction of positive conditions is still not sufficient for knowledge: in addition to a certain set of conditions obtaining, knowledge further depends on particular conditions *not obtaining as a result of luck* – when a particular condition obtains in a lucky fashion, then knowledge is sometimes said to be excluded or compromised. On this view, knowledge depends both on a conjunction of positive conditions (such as belief and truth) and some further anti-luck condition.

The idea that knowledge requires some anti-luck condition is often traced back to Edmund Gettier's seminal paper (1963), which argues against a conjunctive analysis of knowledge in terms of justified-true-belief on the grounds that it is possible to have a justified-true-belief that is luckily true. [5] By way of two examples, Gettier expressed the idea that knowledge cannot depend on luck. The following well-known Gettier case is paradigmatic of the intuition that knowledge excludes (a particular variety) of epistemic luck that pertains to the truth condition of a justified belief:

Smith and Jones have applied for the same job. Smith has evidence for the following conjunctive proposition:

(*D*) Jones is the man who will get the job, and Jones has ten coins in his pocket (Smith has been assured that Jones will be selected, and Smith has counted the coins in Jones' pocket)

D entails the following:

(*E*) The man who will get the job has ten coins in his pocket

Smith sees the entailment from *D* to *E* and <u>accepts *E* on the grounds of *D*</u>

However, unbeknownst to Smith: (1) Smith will get the job instead of Jones, and (2) Smith has ten coins in his pocket. Therefore, *E* is true, whilst proposition *D*, from which Smith inferred *E*, is false. In this case, *E* is true, Smith believes that *E* is true, and Smith is justified in believing *E* is true on the grounds of *D* (hence, if knowledge is justified-true-belief, then

[4] E.g. Williamson (2000)
[5] Gettier, Edmund, "Is Justified True Belief Knowledge?" *Analysis*, (1963), Vol. 23, No. 6, pp. 121-123

Smith has knowledge). However, Gettier's example appeals to the intuition that Smith does *not* know that *E* is true because *E* is only true in virtue of the number of coins in Smith's pocket, and Smith does not know how many coins are in his own pocket. Smith's belief is therefore luckily true because its truth-value depends on factors that are not connected to the factors that confer justification for Smith believing the proposition in question; hence Smith could easily have had a false belief.

The above Gettier case involves a subject having a justified true belief about some proposition that is luckily true. Specifically, it is lucky that the given belief is true, since in a wide-class of counterfactual circumstances in which Smith holds the same belief, with the same degree of justification (where justification is understood as defeasible), Smith fails to have a true belief. The above case demonstrates that luck is also normative – it can be good or bad. Since having true beliefs about the world is reasonably considered to be more intrinsically valuable than having false beliefs about the world, the luck in question – that allows an agent to have a true belief - is "good"; it helps us approximate the cognitive goal of having true beliefs about the world.[6]

Gettier cases are constructed out of a synthesis of good and bad luck (Zagzebski, 1994; Pritchard, 2005). More specifically, as Zagzebski notes, [7] Gettier cases involve good luck "cancelling out" bad luck to allow an agent to have a (fortuitously) true belief.[8] That is to say, Gettier cases involve some feature *F* of agent S's doxastic situation that is amiss; typically, *F* is amiss because it allows *S* to justifiably infer a true belief that *p* (via some method "M") that renders *p* false in a wide-class of counterfactual circumstances in which S believes that *p* via M (In the above Gettier case, *F* may be substituted for the fact *that Smith will get the job instead of Jones*). However, Gettier cases also involve some further feature *F** of S's situation that allows S to have a true belief that *p* via M (In the foregoing example, *F** can be substituted for *Smith having ten coins in his pocket*). The *bad* luck involved in Gettier cases is a result of *F*, which - were it not for the intervention of *F** - would allow S to have a justified false belief that *p* (via M). The good luck arises due to *F**, which unbeknownst to the agent, allows him to have a true belief in the actual world. The value we ascribe to the

[6] Zagzebski, L., "The Inescapability of Gettier Problems", *The Philosophical Quarterly*, Vol. 44, No. 174, (1994), pp. 65-73
[7] See Pritchard, *Epistemic Luck*, (OUP, 2005), p 149. Also: Zagzebski, Linda, "The Inescapability of Gettier Problems", *The Philosophical Quarterly*, (1994), Vol. 44, No. 174, pp. 65-73
[8] Ibid.,

presence of some lucky feature of our doxastic environment is a function of how that feature enhances or frustrates normative goals.

An event is lucky, therefore, only insofar as it has an *effect* on some agent, and the luck in question may be evaluated as *good* or *bad* only insofar as the effects of a given event are desirable or troublesome. As Pritchard notes, the notion of luck should therefore be distinguished from the idea of chance, which is probabilistic.[9] Luck is instead a normative concept, as its very plausibility depends on judgments or ideas about what *ought* to be the case. For example, events can occur as a result of chance; but where no one is affected by the occurrence of chance events, it would seem inaccurate to describe them as "lucky" - and incongruous to label them "good" or "bad" either way. [10] In this respect, "chance events" can be analysed probabilistically. However, luck has a subjective and response-dependent component, as something that depends primarily on the existence of agents (or more specifically, on the existence of *values*). As I will show, "epistemic luck" is an umbrella term for different kinds of luck that are of significance or effect in relation to an agent's possession of knowledge.

The intuition at stake in Gettier cases is that Smith lacks knowledge because his justified belief is true only as a result of luck, specifically good luck. The majority of epistemologists buy into the platitude that knowledge excludes luck, and hence most would be inclined to deny Smith knowledge. The view that no instance of knowledge can partially depend on luck is called "incompatibilism". Incompatibilism is a popular view, as the appeal of Gettier cases attest. As I will argue, incompatibilism about epistemic luck is motivated by the view that knowledge requires a conceptual connection between the truth of a proposition and the justificatory grounds we have for believing that proposition – and when an agent's justified belief is true in virtue of luck, the relevant connection between justification and truth is severed (I will also briefly respond to the kind of luck at stake in Gettier cases by supporting an externalist safety condition).

However, incompatibilists about epistemic luck will seldom claim that knowledge excludes luck simpliciter. Instead, incompatibilists will usually hold a differentiated position, which is to say that knowledge excludes luck that relates only to *particular* epistemic conditions on which knowledge depends. For example, Peter Unger (1968) argues that knowledge does not

[9] Pritchard, "Epistemic Luck", *Journal of Philosophical Research*, (2004), Vol. 29, pp. p195
[10] Ibid.,

exclude all luck that arises in the belief-forming-process, but only excludes luck affecting the truth of one's belief[11]:

"...a complete absence of the accidental is claimed, not regarding the occurrence or existence of the fact known, nor regarding the existence or abilities of the man who knows, but only as regards a certain relation concerning the man and the fact"[12]

Unger is allowing that it may be lucky *that someone S believes some proposition p*. At face value, this variety of luck (that pertains to an agent believing some proposition *p*) is not enough to make knowledge that *p* lucky, nor is it enough to make the belief that *p* lucky. Accordingly, just because events "might have been different", this, as Unger emphasises, does not seem enough to make a belief that *p* lucky.[13] Notably, this is because there is a distinction between (1) *it being a matter of luck that S believes p* and (2) *S's belief that p being lucky*. The first describes the fact that it is a contingent matter that an agent is related to the world in a certain way, and hence it is lucky that an agent believes some proposition *p* because it is conceivable that one's relation to the world could have been different, such that one fails to believe *p*. Given the contingency of any agent's standing in the world, it will always be lucky that someone believes some proposition. This variety of luck may be understood as "doxastic luck":

(1) Doxastic Luck: It is lucky that S believes *p*.

Doxastic luck does not compromise propositional knowledge, since the fact that it is lucky that an agent believes that *p* does not imply a disconnection between the agent believing *p* and the fact that *p*.[14] Moreover, we can apply the qualification that knowledge is compatible with doxastic luck whilst preserving the intuition, in the foregoing Gettier case, that Smith lacks knowledge. This is because Gettier cases derive their argumentative force from the second variety of epistemic luck mentioned, which is consistent with what Mylan Engel terms *veritic luck.*[15]

(2) Veritic Luck: it is lucky that S's belief that *p* is true

[11] Unger, Peter, "An Analysis of Factual Knowledge", *The Journal of Philosophy* (1968), Vol. 65, No. 6, p,161
[12] Ibid.,
[13] Ibid.,
[14] See: Pritchard, *Epistemic Luck*, p. 135. As Pritchard remarks, "genuine knowledge possession is not undermined by the fact that it is a matter of luck that the agent is in a position to know anything at all at that moment", p.135
[15] Engel, Mylan, "Is Epistemic Luck Compatible with Knowledge?", *The Southern Journal of Philosophy",* (1992), Vol. 30, No. 2, pp. 59-75

In the case of veritic luck, the *truth* of S's belief *p* is lucky. When a belief is veritically lucky then the relationship between an agent's justificatory grounds for believing *p* and the truth-value of *p* are disconnected, in a way that is often said to exclude knowledge. Incompatibilists will claim that someone such as Smith lacks knowledge because he harbours a true belief that is veritically lucky. The commonly held view that knowledge excludes veritic luck should therefore be distinguished from the bolder and less plausible view that knowledge excludes luck simpliciter.

Some types of luck matter, others don't. This appears to be the mantra. Whether any particular variety of luck matters, has been represented as a function of whether that variety of luck is harmful to knowledge possession. When a particular type of epistemic luck is harmful to knowledge possession, then that species of luck may be termed "malign". Equally, any type of epistemic luck that is harmless in relation to knowledge possession is "benign".[16] Epistemologists for the most part focus on eliminating malign varieties of luck, for benign luck is compatible with knowledge.

Once we identify a particular species of luck "X" as malign, our response will fall broadly into either of two categories. Firstly, we may accept incompatibilism, which is to say that X *excludes* knowledge; in which case, any instance of knowledge possession must necessarily be free from the variety of malign luck "X" identified. Alternatively, we may adopt a compatibilist view of X. The latter would entail the recognition that X is harmful, and perhaps severely so, whilst maintaining that knowledge is still possible even under these unfavourable circumstances. Those who adopt a compatibilist view of X may construe such luck as "harmful but not fatal". The incompatibilist, however, will regard any instance of X as fatal to knowledge. So, I suspect there are at least two invariantist approaches to understanding any particular instance of malign luck; one is incompatibilist, and the other, compatibilist. It is mistaken to think that because something is harmful it is necessarily fatal, and hence it is a mistake to think that identifying a particular variety of luck X as malign necessarily entails an incompatibilist view of X with knowledge possession.[17]

When a particular variety of luck X is harmful, then an adequate characterization of knowledge will need to postulate some further epistemic condition C that militates against X. If, moreover, we are incompatibilists about X and claim that X *excludes* knowledge, then C

[16] For a discussion of this issue, see: Pritchard, *Epistemic Luck*, (OUP, 2005), p. 131-141
[17] I argue in chapter 5 that reflective luck is compatible with a type of imperfect externalist knowledge.

will necessarily be a condition that knowledge depends on, much like knowledge depends on other conditions such as belief and truth. Typically, "X" is substituted for *veritic luck*, for this is the variety of luck anti-luck epistemologists usually have in mind when they claim that knowledge excludes luck.

Veritic luck and Epistemic Safety

The reason Smith allegedly lacks knowledge is because of the veritically lucky nature in which his belief turns out to be true. When a belief that *p* is veritically lucky in the actual world (let's call the actual world "α"), there are a sufficient number of cases (possible worlds) in which one could have falsely believed that *p* via the same method as in the actual world.[18] So, although one's belief that *p* may be true in α, when one's belief that *p* is veritically lucky in α, then it seems *easily the case* that one could have falsely believed that *p* in α. In modal terms, possible worlds (call these worlds "β worlds") in which one falsely believes that *p* are very similar to the actual world – and hence are easily conceivable as *being* the actual world. In order to ensure that a belief that *p* is not veritically lucky in α, a true belief needs to therefore be situated a certain distance from β-worlds in which one falsely believes that *p*.

The requirement that knowledge entails one's true belief being a "safe distance" from falsity may be represented in modal terms as an anti-luck safety requirement, which can be formulated as follows:

For all cases α and β, if β is close to α and in α one knows that C obtains, then in β one does not falsely believe that C obtains.[19](T. Williamson, 2000)

Safety is a modal state that concerns what could easily have happened, or what might easily have been the case. [20] Where β is close to α, and where C obtains in α but not in β, then it might easily have been that C does not obtain in α. Here, "C" may represent any condition. For example, Timothy Williamson elucidates that "to be safe on top of a cliff, a young child must be at least three feet from the edge; it is not enough to be some positive distance or other, no matter how small, from the edge".[21] There is thus nothing distinctively epistemic

[18] This is closest to Pritchard's reading. See: *Epistemic Luck*, p.156. However, as I will show, only *one* close possible world in which *p*, and where one has a false belief that *p*, is actually enough to undermine knowledge that *p* in the actual world.
[19]Williamson, Timothy, *Knowledge and its Limits*, (Oxford University Press, 2000) , p.128
[20] Ibid., p123
[21] Ibid., P125

about the above safety condition; it is a fairly mundane modal state that may apply to any condition.

However, safety is often applied to knowledge because knowledge is often thought to require a similar margin from error – it is not merely enough for one's belief that *p* to be true, there should also be no danger of one falsely believing *p*. Much like someone must be a *certain* distance from the edge of a cliff in order to be in a *safe* position, an agent's true belief that *p* must also be situated "a certain distance" from possible worlds in which that agent falsely believes *p* in order for one's true belief that *p* to be free from veritic luck. Consequently, for whenever S knows that *p* (via some unspecified method "M" as the means by which S knows that *p*) in the actual world α, there should be no cases β similar to α in which S believes that *p* (via M) and ¬*p* obtains. Importantly, safety does not require 1) that a belief that *p* is *responsive* or *sensitive* to the truth-value of *p*, such that in all cases β which are similar to α, and in which *p* obtains, one believes that *p* via M. Instead, safety only entails 2) that in all cases β one does not falsely believe that *p* via M. These two requirements, which represent sensitivity and safety respectively, do not contrapose.[22] (In chapter 4, I defend a view that means that whilst these two conditions do not contrapose, they have very similar predictions).

If knowledge entails an anti-luck safety requirement, then knowing *p* (via M) in our actual world α (Kp_α) entails that a belief that *p* (B*p*) is true in α and β, where β is a close world (CW) in which one believes *p* via M. This understanding of safety entails that there need only be *one* β-world in which the conjunction ($Bp \wedge \neg p$) obtains in order to make B*p* unsafe in α. Formally: $Kp_\alpha \rightarrow \neg\exists\beta(Bp_\beta \wedge \neg p_\beta \wedge CW(\beta))$

On this schema, a possible β-world is a close world only if it is sufficiently *similar* to the actual world. Ordering β-worlds in terms of similarity to the actual world importantly allows us to include worlds that are remote in terms of probability of occurring as nonetheless relevant to the safety of one's belief in the actual world. By ordering possible β-words in terms of *similarity* to the actual world, the fact that a particular β-world has a low statistical probability of *being* the actual world does not necessarily make it a remote possible world; closeness, on this account, is not necessarily determined probabilistically. This is not to say that the probability of an event is entirely irrelevant to the closeness of a possible world (For

[22] In short, safety holds the belief that *p* fixed and asks whether *p* is true in nearby worlds in which one believes *p*. Sensitivity varies the truth-value of *p* to see whether our belief responds to the truth-value of *p* at nearby worlds. For a good discussion of why these conditions do not contrapose, see: Sosa, Ernest, "How to Defeat Opposition to Moore", *Nous,* (1999), Vol. 33 pp. 149-150

example, I struggle to think of any worlds that have an extremely high probability of occurring yet are modally remote or very dissimilar to the actual world). Nonetheless, there is positive reason to resist the view that ordering of possible worlds should be determined probabilistically. For example, the possible world in which the proposition "I will win the lottery" is true is very close to the actual world, and hence counts as a relevant β-world, since very few conditions in α would need to differ in order to make the lottery proposition true. Determining the closeness of possible worlds in terms of *similarity* to the actual world (which is not a probabilistic notion) allows us to capture what is intuitively unsafe about the following utterance:

"I know that I will *not* win the lottery" (where the utterer has bought a lottery ticket).

A similarity ordering of close possible worlds allows us to deny that the above utterance about the lottery (call this utterance "*L*") is safe because there is a close possible world in which one falsely holds *L* – namely, the world in which one wins the lottery. Attempting to determine closeness probabilistically, however, would make the world in which I win the lottery (and *L* is false) a "remote" possible world, because it has a low statistical probability of being the actual world. Ordering nearby worlds probabilistically would therefore, perversely, make *L* a safe utterance! Consequently, there is at least a presumption that we should order close possible worlds in terms of similarity to the actual world, which should not be conflated with the statistical probability of such a world *being* the actual world. Hence, close possible worlds will be worlds that share a similar set of conditions to the actual world, and hence are easily conceivable as being the actual world.

However, *L* can only be described as an instance of an unsafe utterance, on a similarity ordering of close possible worlds, if, as stipulated, we also quantify across *all* close possible worlds and not merely *some* or *most* possible worlds. Quantifying across merely *some* close possible worlds would insufficiently capture the intuition that *L* is unsafe in the actual world. This is because failing to quantify across *all* close possible worlds would allow *L* to be safe in the actual world, because in *most* close possible worlds *L* is a true proposition; after all, although the world in which I win the lottery is close, there are more close possible worlds in which I fail to win the lottery. In order to accommodate the intuition that *L* instantiates an unsafe belief, we must therefore quantify across *all* close possible worlds, ordered in terms of similarity to the actual world, in which an agent believes *L*. On this ordering of possible worlds, it takes very little to make a given belief unsafe - just *one* close possible world in

which an agent has a false belief that *p*, (where the belief that *p* is true in the actual world and in most nearby possible worlds in which one believes that *p*), is enough to make one's belief that *p* unsafe.

However, depending on our pre-theoretical intuitions about certain cases of knowledge possession (or knowledge failure), there are potential challenges to my formulation. The formulation I have supported quantifies across *all* close possible worlds. However, let us see how the requirement that we have a safe belief functions in relation to the following example (Comensana, 2005; Sosa, 2000):[23]

[Garbage Chute Example: I throw a trash bag down the garbage chute of my condo. Some moments later I believe, and know, that the trash bag is in the basement. However, the closest possible world where my belief is false is plausibly one where, unbeknownst to me, the bag is stuck somewhere in the chute, and I still believe that it is in the basement][24]

Sosa describes this above example as an intuitive case of knowledge possession.[25] And if the above example *is* a case of knowledge possession, then we have some reason to resist a sensitivity principle. This is because a sensitivity analysis of knowledge states that to know some proposition *p*, our belief that *p* must subjunctively depend on the truth of *p*.[26] Where there is a possible world in which one falsely believes that *p*, then our belief that *p* will fail to subjunctively depend on the truth of *p*. My belief in the above case is not sensitive.[27] However, on a version of safety that quantifies across *all* nearby possible worlds, then my belief is not safe either.

We should therefore distinguish between two different versions of safety, both of which potentially have very different implications. One version of safety adopts partial-quantification across close possible worlds (PCQ) to state that S knows *p* if and only if S's belief that *p* is true in *most* close possible worlds in which S believes that *p*. The second version adopts full-quantification across close possible worlds (FCQ) and states that S knows that *p* if and only if S's belief that *p* is true in *all* close possible worlds in which S believes

[23] Comensana, Juan, "Unsafe Knowledge", *Synthese*, (2005), Vol. 146, Issue. 3, p. 396 (This Counterexample was originally proposed by Sosa, "Skepticism and Contextualism" in *Philosophical Issues*, (2000) Vol. 10, p. 130)
[24] Comensana, "Unsafe Knowledge", p. 396
[25] Sosa, Ernest, "Skepticism and Contextualism", *Nous*, (2000), Vol. p. 13
[26] For the relevant account of subjunctive dependence, See: Nozick, Robert. *Philosophical Explanations*, pp. 197-205
[27] Sosa agrees (Skepticism and Contextualism, p13), but sees this examples as a difficulty for sensitivity, rather than a virtue.

that *p*. Adopting a PCQ version of safety has a potential advantage in that, by quantifying only across *most* possible worlds, PCQ accommodates our intuitions at stake in the above garbage chute example. Conversely, because the world in which one has a false belief about the trash bag is considered close, then an FCQ version of safety, which quantifies across *all* close possible worlds, will deny us knowledge in much the same way that a sensitivity principle does.

However, this consideration should not necessarily drive us to relax our safety requirement. Given that the probability of an event occurring is often a poor guide to closeness, then, if we do not share the intuition that Sosa describes, then FCQ places us in a strong position to state that what Sosa describes is actually a case of knowledge-failure. Moreover, providing that we still want to capture the lottery intuition, then we have a reason to favour a version of Safety-FCQ or indeed sensitivity. Accordingly, on Safety-PCQ we would have to grant that the lottery case is *not close* in order to make it a case of knowledge-failure – because, after all, in *most* possible worlds in which I think I will lose the lottery my belief is true. If we therefore want to capture our intuitions about the lottery case, then we should adopt safety-FCQ, but this would have the effect of denying knowledge in the garbage chute example.[28] On either of these two versions of safety, it seems one of our intuitions has to give.[29]

The above "conundrum" is only a conundrum insofar as we actually harbour both intuitions. (I, for one, do not think the garbage chute example is a case of knowledge, but rather, a case of reasonable belief). However, those who want to keep both the intuition that we do not know that we will not win the lottery *and* the intuition that we have knowledge in the garbage chute example, may refine their version of safety to capture both these intuitions by making a further distinction between *close* possible worlds and *very close* possible worlds. By making this distinction, we could retain Safety-PCQ across close possible worlds, whilst preserving Safety-FCQ across *very* close possible worlds. On this schema, the case of winning the lottery would be classifiable as a *very close world* and would therefore be fully quantified across (of course, the probability of this world occurring is extremely low, but probability needn't determine closeness; it can be considered very close because very few features of the actual world would need to change in order to bring it about). Additionally, we could

[28] FCQ would present a problem for Sosa if he wants to preserve the lottery intuition, because he thinks the garbage chute example *is* a case of knowledge possession (Ibid., p.13).

[29] Pritchard also discusses this problem. See: Pritchard, Duncan, "Sensitivity, Safety, and Anti-luck Epistemology", in *The Oxford Handbook of Skepticism*, Edited by John Greco (Oxford University Press, 2009), p. 448

maintain that the garbage chute example is close but not *very* close – since *more* would have to change in order to bring it about that my belief about the trash bag is false.[30] Consequently, in the garbage chute example, we need only quantify across *most* possible worlds. We could therefore have a hybrid version of safety that is strict enough to accommodate the lottery intuition as a case of knowledge-failure, whilst relaxed enough to accommodate intuitions about Sosa's garbage chute example.[31]

To have a hybrid theory we need to recognize that "closeness" is a gradable concept, that is, an object can be more or less close. Some objects are close, others are closer still. Therefore, the close possible world is not necessarily the *closest* possible world. Once we recognize that closeness is a gradable concept, an important template for accommodating both intuitions is in place.

So far, I have suggested that a modal safety condition offers an analysis of knowledge that excludes veritic luck, the most commonly targeted malign variety of epistemic luck. The formulation I have proposed determines closeness of possible β worlds in terms of similarity to the actual world α, and quantifies across *all* nearby β worlds. Whether we should define safety in these demanding terms will depend on one's pre-theoretical response to particular examples of belief-formation. However, I have tentatively suggested that there is reason to quantify across all nearby worlds and not merely a wide class of nearby worlds, because it enables us to capture the intuition that we cannot know that we will not win the lottery. The lottery intuition also provides us with a presumption in favour of a similarity ordering of close possible worlds. However, my analysis may still be challenged on a number of grounds, and by no means constitutes an exhaustive discussion of epistemic safety. (Although I have suggested that if we want a more relaxed version of safety, then we can plausibly distinguish between *close* and *very close* possible worlds in order to preserve the lottery intuition). However, for the purpose of this introduction, I have wanted to note the *aim* of a safety requirement (to eliminate veritic luck) and to sketch some paradigmatic means by which this might be achieved.

[30] Pritchard opts for this interpretation: Pritchard, Duncan. "Sensitivity, Safety, and Anti-luck Epistemology", in *The Oxford Handbook of Skepticism*, Edited by John Greco (Oxford University Press, 2009), p. 51-452
[31] Ibid.,

Reflective Luck

Veritic luck has been identified as potentially harmful to knowledge possession. One way in which epistemologists have responded to veritic luck is to make knowledge depend on the aforementioned safety condition. A safety condition is an externalist condition, because whether a belief is safe does not depend on any factors internal to the cognitive perspective of the agent. Externalism insists that knowledge does not *necessarily* depend on justificatory-conferring factors that are cognitively available, or accessible. Externalism, like internalism, is therefore a meta-epistemic thesis that makes a claim about the dependence-base for knowledge, which is compatible with a number of different theories of justification – including both a safety-based approach advocated here, as well as reliabilist theories of justification, such as Alvin Goldman's.[32] The consequence of an externalist approach is that a belief that *p* may be safe irrespective of whether an agent is aware of, or has cognitive access to, the factors that make a belief that *p* safe (the factors that fix the truth of *p* across nearby worlds in which an agent believes that *p*). Furthermore, because safety is an externalist condition, whether a belief that *p* is safe does not depend on whether the agent concerned *believes* that the belief that *p* is safe. However, as I want to explore in this thesis, a belief that is safe (or satisfies some other externalist condition for knowledge) may still fall short of knowledge on the grounds that externalism has to countenance "reflective luck".

Reflective luck was first explicitly discussed as a potential challenge to externalism in Duncan Pritchard's *Epistemic Luck* (2005).[33] Pritchard defines a belief as reflectively lucky when "Given only what the agent is able to know by reflection alone, it is a matter of luck that his belief is true".[34] Reflective luck is a potential challenge to any epistemologically externalist theory of knowledge; that is, any theory that does not make knowledge necessarily depend on justificatory conditions that are cognitively available, or accessible, to a given subject.

No authors have denied the plausibility of externalist theories of justification *explicitly* on the grounds that externalism gives rise to reflective luck. Nonetheless, the idea that reflective luck excludes knowledge is invoked by several epistemologists. For example, Linda

[32] Goldman, Alvin. "What is Justified Belief", in *Epistemology: An Anthology*, Blackwell, (2008) – Originally published in *Justification and Knowledge*, Edited by G.S. Pappas (Reidel, 1976), pp. 1-23. Reliabilism is an externalist theory of justification that states that the causal ancestry (or more specifically, the antecedent belief forming process) responsible for the belief that *p* exhaustively determines whether *p* is justified.
[33] Pritchard, Duncan, *Epistemic Luck*, (Oxford University Press, 2005), p. 175
[34] Ibid.,

Zagzebski (1996) alludes to reflective luck as a sufficient reason to reject externalist process reliabilism:

"The value of the truth obtained by a reliable process in the absence of any conscious awareness of a connection between the behaviour of the agent and the truth he thereby acquires is no better than the value of a lucky guess." [35]

Richard Fumerton (1995) alludes to a similar idea:

"Perception, memory and induction *may* be reliable processes in Goldman's sense……but the skeptic can argue, we have no reason to believe that these processes *are* reliable, and thus, even if we accept reliabilism, we have no reason to conclude that the beliefs they produce are justified."[36]

Zagzebski's claim highlights the sense in which a belief's truth is reflectively lucky if we lack any cognitive perspective on the properties that make it true. Fumerton's argument is analogous, implying that *if* belief-forming processes are reliable *then* from the agent's cognitive perspective the reliability of these belief-forming processes is accidental. Accordingly, if in the actual world and most nearby possible worlds we lack any cognitively available reason to believe that our belief-forming processes are reliable, then the justificatory status of our first-order belief is *out of our control*. Implicit in these accounts is the suggestion that the correct external relation between belief and truth, alone, is not enough to know that *p*. Knowledge (that *p*) further entails that a belief is also free from reflective luck, of the kind that arises when an agent has no cognitive perspective on the factors that confer justification for a given belief.

A further (implicit) appeal to reflective luck as incompatible with knowledge possession is found in Laurence BonJour's example of Norman, the clairvoyant:

"Suppose that a person, Norman, is a reliable clairvoyant with respect to the geographical whereabouts of the president of the United States. He frequently has spontaneous beliefs or hunches, which he accepts without question, concerning the location of the president on a particular day, and in fact these are always correct. But Norman pays very little attention to news reports and other sorts of information about the president and his or her whereabouts

[35] Zagzebski, Linda, *Virtues of the Mind* (Cambridge University Press, 1996), p. 304
[36] Fumerton, Richard, *Metaepistemology and Skepticism*, (Rowman and Littlefield, 1995), p.174

and has never made any effort to check his hunches independently. Nor does he have any real conception of how such hunches might work or any general views about the reliability of such a process. Norman's beliefs resulting from his spontaneous clairvoyant hunches satisfy the reliabilist's requirements for justification, but are they really justified?"[37] (BonJour, 2002)

Norman's beliefs, if justified, satisfy a robustly externalist view of justification due to his lack of any cognitive perspective on the properties or relations that make his hunches justified. Let us also continue with the further supposition, for argument's sake, that Norman's beliefs are epistemically safe – I.e. no nearby possible world exists in which Norman has false beliefs regarding the president's whereabouts. However, notwithstanding Norman's clairvoyant beliefs meeting an externalist justificatory requirement, Norman still strikes us as merely the passive and lucky recipient of reliably true beliefs. The kind of luck in this instance does not stem from the reliability of Norman's belief-forming processes, but instead, arises from Norman's lack of any cognitive perspective on the reliability of his belief-forming processes. Accordingly, even if Norman's belief-forming processes are objectively reliable, Norman's belief is reflectively lucky because nothing within his reflective grasp could *warrant him in believing that his clairvoyant beliefs are reliable*.

Internalists such as BonJour draw our attention to how a true belief that *p* can hypothetically meet a demanding externalist justificatory requirement, and hence not be veritically lucky, whilst still being lucky *from the agent's cognitive perspective*. In the case of Norman, his belief turns out to be true (and for the sake of argument, its truth obtains across all nearby possible worlds in which Norman has clairvoyant hunches) and is not subject to veritic luck. Nevertheless, Norman's belief is still reflectively lucky because he lacks any reason to believe that his hunches are reliably true across all nearby worlds in which he has reliable hunches about the president's location.

The above accounts highlight how externalist accounts of justification have to countenance a particular variety of luck that pertains to it being lucky, given what an agent is reflectively aware of, that a given belief is true.[38] Reflective luck affects any post-Gettier analysis of knowledge, that is, any analysis of knowledge that aims to eliminate veritic luck by way of incorporating an external condition that establishes a conceptual connection between truth

---

[37]BonJour, Laurence, *Epistemology: Classic Problems and Contemporary Responses*, (Rowman and Littlefield, 2002) p.231

[38] See: Pritchard, *Epistemic Luck*, (OUP, 2005), p. 173

and justification, or makes justification in part a function of truth.[39] Pritchard states that in this regard reflective luck arguments affect both internalism and externalism, since whether an external condition obtains will not necessarily be reflectively accessible to the agent[40]. Accordingly,

"Zagzebski…is mistaken to believe that the kind of reflective luck that she describes only afflicts externalist epistemologies since, as we have seen, it is essential to any Gettier proof epistemology, whether externalist or internalist, that it incorporates an external condition on knowledge, the obtaining of which, will not be reflectively accessible to the agent. Accordingly, if this is a problem at all, it is a problem that afflicts all adequate post-Gettier epistemologies".[41]

Pritchard is right that any account of knowledge *that is free from veritic luck* must necessarily incorporate an external condition that connects truth and justification, where this condition may be reflectively inaccessible. However, internalism about justification entails that justification is necessarily supervenient on reflectively accessible factors; internalism does not allow any external condition that eliminates veritic luck to be relevant to justification unless that condition is also reflectively accessible. In this regard, we must recognize that reflective luck arguments in principle run deeper for externalism. Accordingly, internalists can always state that the commitment to eliminating reflective luck is more essential than providing a post-Gettier epistemology that accommodates everyday knowledge that is free from veritic luck (and whose dictates can be satisfied). Internalists can in principle alleviate reflective luck from their theories of justification, and in doing so can deny that subjects such as Norman possess knowledge by fundamentally denying that they are epistemically justified (of course, this internalist strategy may have an unacceptably high cost).

Reflective luck arguments imply that counterfactual externalist conditions, such as safety, are not enough to eliminate varieties of epistemic luck pertaining to the truth-status of a belief, because there exists some further ostensibly malign sense in which a true belief may be reflectively lucky. Although I have run reflective luck objections against reliabilism and safety, the same considerations apply to any externalist analysis of knowledge whereby the

[39] Ibid.,
[40] Pritchard, "Epistemic Luck", *Journal of Philosophical Research*, (2004), Vol. 29, p. 212
[41] Ibid.,

justificatory condition that converts true belief into knowledge (or *whatever* condition converts true belief into knowledge) is external to the cognizer.

The problem is that any externalist epistemology gives rise to reflective luck. But then again, there is strong reason to externalize our justificatory commitments. This is because, in order to eliminate veritic luck, we must establish a conceptual connection between the truth condition for knowledge and the justificatory condition for knowledge; this allows that knowledge becomes a function of how we stand in relation to nomological properties and processes in the world (where this standing is not something that we are necessarily able to discern "via reflection"). As Gettier cases show, we can have good internal grounds for thinking that *p* even though *p* is not knowledge because the connection between truth and justification is severed. In this respect, externalist epistemologies respond robustly to one kind of malign luck (veritic luck) that internalist conceptions of justification give rise to, but in doing so, leave us with reflective luck.[42]

However, as it stands, the concept of the truth of a belief being lucky *given what an agent is reflectively aware of* is sufficiently vague to include a variety of different challenges to externalism under the rubric of "reflective luck arguments". As I will show, reflective luck arguments can invoke different platitudes. Some of these platitudes are about local cases of belief-ownership, whereas others highlight a general meta-epistemological concern about the status of externalist knowledge.[43] In what follows, I will try to more clearly define what this elusive "reflective luck" is. To this end, I reject the idea that reflective luck arguments depend on views about doxastic voluntarism; and instead I propose a modal analysis of reflective luck. I then provide a template for what an anti-reflective luck condition might look like, before showing how it presents epistemologists with a conundrum.

Interpreting Reflective Luck

The judgment with regards to certain cases of reflective luck is that the subjects who are externalistically justified appear mere "belief machines" or "passive recipients" of reliably true beliefs. Recall that it's Norman's *passivity* in the belief-forming process that is objectionable for internalists such as BonJour.

[42]Pritchard, Duncan, "Epistemic Luck", *Journal of Philosophical Research*, (2004), p. 209 – Pritchard emphasises that reflective luck is common to all post-Gettier epistemologies and is a function of externalism
[43] See: Breyer, Daniel. "Reflective Luck and Belief Ownership", in *Acta Anal* (2010), Vol. 25, pp. 122-154

The opposite of passive receipt, then, is active acquisition. With the right sort of active processes, epistemic agents will – if the internalist intuition is well founded – be able to have beliefs that are *not* outside of one's reflective perspective? One way of interpreting reflective luck arguments is thus to understand them as platitudes about belief control. The lack of perspective on the justificatory factors appears to confer a lack of control over what is believed. The condition that knowledge that *p* precludes passive receipt of the belief that *p* would be readily satisfied if we were fully in control of our beliefs, that is, if we could *choose* our beliefs.

Cases of reflective luck are cases in which an agent appears to lack control, but so are all or at least most cases of belief acquisition - and this is the problem. The negation of control is a necessary condition for a belief being reflectively lucky. However, a lack of control cannot be a *sufficient* condition for a belief being reflectively lucky without conceding the sceptical consequence that we do not know a variety of propositions (or without having to show that doxastic voluntarism about, for example, perceptual beliefs is true). Thus, internalists who wish to avoid radical scepticism will need to offer a definition of reflective luck that captures the relevant intuition about Norman lacking control, but which imposes an internalist justificatory requirement that is neutral regarding doxastic voluntarism; that is, a condition whereby an agent's lack of control is not *alone* enough to bring about the kind of reflective luck that internalists such as BonJour think compromises knowledge. Otherwise they have an explanatory burden: namely, explaining how it is we can control all of our beliefs, including beliefs that do not appear to be possessed out of choice, such as perceptual beliefs. My objection to explaining reflective luck in terms of doxastic voluntarism would be that it has an explanatory burden: namely, how do we explain cases of belief possession that do *not* seem chosen, and yet still appear instances of knowledge? As I will show, a modal analysis of reflective luck does not come with this explanatory burden.

Duncan Pritchard (2005) proposes a modal understanding of epistemic luck. His account can be used to make a lack of control a *necessary* condition for a belief being reflectively lucky, without imposing the more controversial requirement that the absence of control is a *sufficient* condition for a belief being reflectively lucky.[44] Accordingly, Pritchard says:

[44] Pritchard, Duncan, "Epistemic Luck", *Journal of Philosophical Research,* (2004), Vol. 29, p.197

[If an event is lucky, then it is an event that occurs in the actual world but does not occur in most of the nearest possible worlds to the actual world (worlds which most resemble the actual world)][45]

Let the "event" in question represent "the event of believing that *p*". Our attempts to eliminate veritic luck by appeal to an epistemic safety condition rely on the above modal analysis of luck. According to this analysis, S's true belief that *p* is safe in the actual world if, and only if, S has a true belief that *p* in all nearby possible worlds (note: Pritchard quantifies across only *most* nearby possible worlds). A modal understanding of luck allows us to capture the idea that luck entails the absence of control, without making the negation of control a sufficient condition for an event being lucky.[46] This analysis is more consistent with our intuitions about putatively justified beliefs. Accordingly, there may be instances in which an agent lacks control over a given event E, but where E nonetheless is *not* lucky on the grounds that E occurs in a wide class (or all) nearby possible worlds. For example, we may lack control over the kind of knowledge we acquire via perception – but for as long as our perceptual faculties are sufficiently reliable, then the beliefs we thereby acquire via perception may be true in a wide class (or all) nearby worlds in which we possess the same perceptual faculty and form the same beliefs on the basis of this faculty – and hence are not *lucky*.

In the same way that the most plausible account of veritic luck can be understood in modal terms, an adequate analysis of reflective luck may also be understood in modal terms. To this end, Pritchard stipulates that we can understand the phenomena of reflective luck only when nearby possible worlds are ordered in a "non-standard way", that is, when nearby worlds are ordered solely in terms of what an agent is able to know via reflection alone in the actual world.[47] Reflective luck arises when an agent has a true belief that *p* in the actual world despite there being a wide class of nearby possible worlds, which are ordered in terms of what an agent is able to know via reflection alone in the actual world, in which that agent falsely believes that *p*.

On Pritchard's non-standard ordering of possible worlds, agents who are not in possession of reflectively accessible grounds for thinking that they continue to have a true belief in most

[45]Pritchard, Duncan, "Epistemic Luck", *Journal of Philosophical Research*, p.197
[46] Ibid, p. 196-197
[47] Pritchard, Duncan. *Epistemic Luck*, (Oxford University Press, 2005), p.173

nearby possible worlds will therefore succumb to reflective luck on the grounds that just about *any* possible world can count as a nearby world on a non-standard ordering.[48] Conversely, if an agent has reflectively accessible grounds for thinking that he continues to have true beliefs in most nearby worlds, then he will "tend to have beliefs that match the truth of the matter in most of the relevant nearby possible worlds."[49]

On an ordering of nearby possible worlds that quantifies across only *most* nearby possible worlds, ordered in terms of what an agent is able to know via reflection in the actual world, then cases of reflective luck arise when an agent has no reflectively accessible grounds for thinking that a given proposition is true. Incorporating the requirement that we have good reason to think that our beliefs are true will apparently eliminate this luck, meaning that in *most* nearby worlds, ordered in terms of what an agent knows via reflection in the actual world, the given belief will be true.[50] If our analysis of reflective luck quantifies across only most possible worlds, therefore, then the meta-epistemological challenge that reflective luck arguments present to externalist theories of justification is merely *local* – that is, they are confined to peculiar cases of knowledge possession such as Norman's, where epistemic subjects fail to have true beliefs at *most* nearby possible worlds, on a non-standard ordering of possible worlds.

However, if our analysis of reflective luck is one that quantifies across *all* nearby possible worlds, then the species of meta-epistemological scepticism that reflective luck arguments present is global. As I will show, the issue of quantification is essential to whether reflective luck is a condition that a theory of justification can accommodate. Moreover, I contend that the kind of meta-epistemological scepticism that externalism generates, and which reflective luck draws our attention to, can only be eliminated if we are able to satisfy a reflective luck condition that quantifies across *all* nearby possible worlds.

A Reflective Safety Condition

Reflective luck can therefore be understood in modal terms according to which a belief that *p* is reflectively lucky if the truth of one's belief that *p* obtains in the actual world α, but fails to obtain in at least one (or most) close possible worlds β. Such a β world is "close" when an agent possesses the same knowledge via reflection in β as in α.

[48] Pritchard, Duncan. *Epistemic Luck*, p.176
[49] Ibid.,
[50] Ibid.,

If we understand reflective luck in modal terms as a phenomena that is eliminated only once we quantify across *all* nearby possible β worlds, then reflective luck arguments entail that in order to know that *p* (K*p*) one's belief that *p* (B*p*) must be true in all nearby possible worlds that are reflectively indiscernible to the actual world (that is, in all possible worlds that are ordered in terms of what we know via reflection in the actual world). However, in order to satisfy this requirement, an agent must be able to rule out– via reflection – the epistemic possibility of inhabiting a close possible world where one falsely believes that *p* (a "B*p*∧¬*p*" world). On this interpretation, presuming that we also wish to eliminate malign veritic luck from our epistemology, a reflective luck condition on knowledge stipulates that knowledge that *p* entails that one's first-order belief that *p* (B*p*) is safe, and thus free from veritic luck, whilst further entailing a second-order requirement that one knows via reflection (K_r) that one's first-order belief that *p* is safe. A reflective luck condition, on this understanding, entails that one is able to "rule out" – via internal reflection - the epistemic possibility of inhabiting world in which one falsely believes that *p*.

However, an agent can only internally "rule out" the possibility of falsely believing that *p* by knowing via the method of reflection (K_r) the negation of such a possibility. Once we conjoin this reflective meta-condition with our aforementioned safety condition, then knowledge (that *p*) may be understood as entailing the following condition:

RL: $Kp \rightarrow K_r \neg \exists \beta (Bp_\beta \wedge \neg p_\beta \wedge CW(\beta))$

If a belief were to satisfy this schema (RL), then one's first-order belief would be safe, and therefore free from veritic luck, whilst also being free from reflective luck (it would be reflectively safe, if you like). If only RL could be satisfied by actual agents! In reality, RL is an extremely demanding condition, because it requires one to know, via reflection, that *p* is true across all nearby possible worlds in which one believes that *p*. My contention is that no such second-order knowledge is forthcoming; there are some nearby epistemic possibilities that *no* reflective capacity can rule out.

To see this, let us take the following case: "it is reflectively lucky that I am not a brain in a vat in which I am systematically deceived about the existence of an external world" (BIV).[51] By virtue of this hypothesis involving systematic deception, nothing within my reflective vicinity suggests that this hypothesis is false. That is not to say that sceptical scenarios are

[51] The brain-in-a-vat hypothesis was originally discussed in Putnam, Hilary. *Reason, Truth, and History*, (Cambridge University Press, 1981), p. 5

evidentially or even experientially the same as non-sceptical scenarios, but only that it would *seem* (phenomenally) that they are the same. Consequently, there is a possible world in which I believe that I have hands but, in fact, I am a BIV (so the empirical proposition "I have hands" is necessarily false). On this basis, our knowledge of empirical propositions will always be subject to reflective luck. Where "*E"* stands for the empirical proposition "I have hands", but which may equally designate any empirical proposition, we may frame the argument as follows: [52]

(1) If S is in a world in which she is a BIV, then E is false, so S cannot know *E*
(2) If S does not know that she is in a ¬BIV world, then S cannot know *E*
(3) S does not know she is in a ¬BIV world
(4) S does not know *E*

A BIV world would not differ reflectively to a ¬BIV world, and consequently *S* would necessarily be unable to discriminate between BIV and ¬BIV worlds via any kind of reflective process.[53] These sorts of sceptical hypotheses are not a problem for a purely externalist safety principle, because they describe possible worlds that are sufficiently dissimilar from the actual world not to be quantified across. Consequently, the sceptical template outline above is by no means a conclusive argument against everyday knowledge (as I show in chapter 3). However, because BIV worlds and ¬BIV worlds are phenomenally indiscernible, then if possible worlds are ordered in terms of what we know via reflection in the actual world, BIV worlds are *close* to the actual world; hence *S's* belief that *E* is always subject to reflective luck if we cannot know via reflection the negation of BIV (because in a nearby BIV world, ¬*E* is true). If reflective luck is only eliminated by RL, then it seems that even the most uncontroversial belief "I have hands" is also subject to reflective luck.

Implications of Reflective Luck

That reflective luck appears ineliminable has ramifications if we wish to use it as a basis on which to accommodate our intuitions about cases such as Norman's – because making knowledge depend on eliminating reflective luck *simpliciter* would exclude uncontroversial everyday knowledge of propositions such as "I have hands".

[52] Duncan Pritchard uses this example. See: Sensitivity, Safety, and Anti-luck Epistemology", in *The Oxford Handbook of Skepticism*, Edited by John Greco (Oxford University Press, 2009), p. 441
[53] Ibid. (the point is that the truth of E is not sensitive to the reflective grounds we have for thinking E obtains)

The alternative, less demanding account of reflective luck, which could accommodate our intuitions about Norman-type cases whilst preserving our knowledge of putatively true propositions, is a reflective safety condition that only quantifies across *most* nearby possible worlds rather than *all* nearby possible worlds. If we quantify across most nearby worlds, we could resist the implication that an agent's belief that "I have hands" is reflectively lucky on the grounds that *most* nearby possible worlds, on a reflective modal ordering, will not be BIV worlds. Accordingly, if an agent has "good reflectively accessible grounds" then, even in the absence of satisfying a principle such as RL, there are grounds to think that his belief will be true in *most* nearby possible worlds that are ordered in terms of what an agent knows via reflection in the actual world, even if some sceptical modal possibilities remain ineliminable.[54]

However, this response to reflective luck fails to appreciate the extent to which reflective luck arguments are sometimes motivated by a meta-epistemological challenge to externalism: that is, a challenge about the general externalist framework rather than an objection to particular esoteric instances of belief-formation such as Norman's. When reflective luck arguments are used in the context of presenting a meta-epistemological challenges to externalism, then anything weaker than quantifying across *all* nearby possible worlds would fail to eliminate, from the agent's first-person perspective, the epistemic possibility of inhabiting a nearby world in which one falsely believes that *p* - and hence, at least from the agent's reflective standpoint, would fail to eliminate the possibility of falsely believing that *p* in a nearby world, leaving the meta-epistemological challenge to externalism that Fumerton and BonJour present, unanswered.

However, by only focusing on resolving peculiar cases of belief-formation, such as BonJour's example of Norman the clairvoyant, there is no obvious sense that reflective luck arguments *do* motivate an internalist condition such as RL, because at face-value they are not necessarily meta-epistemological challenges to externalism. Accordingly, these "local cases" all involve peculiar cases of belief-formation whereby subjects possess reliably true beliefs that nonetheless do not appear attributable as their own.[55] However, let us postulate an enlightened counterpart case of Norman, who has reflective awareness of the relevant justificatory-conferring grounds for his beliefs. For example, enlightened counterpart

[54] Pritchard, D. *Epistemic Luck* (OUP), p. 78 – the point Pritchard makes is that one's reflective grounds may still be excellent, even if they are unable to transmit to provide one with knowledge that a sceptical scenario does not obtain.

[55] Breyer, Daniel. "Reflective Luck and Belief Ownership", in *Acta Anal* (2010), Vol. 25, p.135

Norman could, when questioned about the reliability of his hunches, adduce persuasive evidence in favour of the reliability of his hunches (it would be impertinent, therefore, to describe enlightened Norman as a belief machine). The temptation, therefore, is to think that reflective luck might be resolved either by constructing the right externalist model of belief ownership or cognitive agency[56], thus denying Norman knowledge whilst preserving externalism. Or alternatively by embracing an internalist justificatory requirement; where the justificatory requirement would only need to capture intuitions about *attribution* and hence would not demand a condition as strong as RL.

Local reflective luck arguments (that is, arguments about peculiar cases of belief possession) should therefore be distinguished from global reflective luck arguments to unearth the meta-epistemological challenge that reflective luck presents to externalism. As Daniel Breyer (2010) argues,

"Global reflective luck argument emphasise the irrespective failure of externalist theories of knowledge to address the higher-order dimension of the sceptic's challenge – namely, the challenge to show that, in fact, our beliefs are reliably formed. In effect, global arguments push the intuition that intellectual curiosity (in the face of scepticism) demands more than what externalism can provide".[57]

The reflective luck arguments presented by BonJour, Zagzebski and Fumerton are *global*, since local cases of peculiar belief-formation are only used in order to motivate general meta-epistemological concerns about how externalism is unable to accommodate higher-order justification.[58] For example, Fumerton is clearly presenting a global reflective luck argument insofar as he explicitly objects to externalism on the grounds that it deprives us of higher-order justification for our beliefs.[59] According to Fumerton, if we accept reliabilism we have no reason to conclude that the beliefs that result from reliable processes *are* reliable, and hence no reason to conclude that reliably true beliefs are justified.[60] These arguments use examples of reflective luck to pose a meta-epistemological challenge to externalism. Necessarily, if reflective luck arguments are used to push the meta-epistemological challenge

[56] Breyer. p. 154
[57] Ibid., p144
[58] Ibid.,
[59] Fumerton, Richard, *Metaepistemology and Skepticism*, (Rowman and Littlefield, 1995), p.175 – the point Fumerton emphasises here is that there is no higher-order justification because externalism just resorts to the very processes, such as reliability, which are under philosophical scrutiny to justify a reliance on those very processes (I discuss this further in chapter 3)
[60] Fumerton, p.175

to externalism, then the reflective luck condition that internalists think knowledge must depend on must quantify across *all* nearby possible worlds (on the non-standard reflective ordering). Short of quantifying across all nearby worlds, the meta-epistemological challenge will remain.

To summarise the challenge: the claim that externalism gives rise to reflective luck is a claim about *strong global meta-epistemological scepticism* (in other words, scepticism about the rationality of our "second-order beliefs"). The scepticism advanced is global because it does not target any particular class of proposition. Likewise, the scepticism itself can be considered strong insofar as it tells us that we *cannot rationally believe that our beliefs are justified* whenever our first-order beliefs are reflectively lucky. Strong scepticism may be contrasted with and distinguished from weak scepticism, which tells us only that we cannot *know* that our beliefs are justified.[61] Scepticism concerning knowledge is altogether weaker than scepticism concerning justification or rational belief, since it is possible to rationally believe that *p* without knowing that *p*, whereas it is not possible to know that *p* without rationally believing that *p*.[62] And so scepticism concerning what is rational to believe has more widespread implications. However, reflective luck arguments entail both the weaker position that we lack knowledge and the stronger position that reflective luck undermines the rationality of our meta-beliefs. Hence, reflective luck arguments tell us that it is not rational to believe that our beliefs are justified on an externalist analysis of justification. That is to say, without being able to know via reflection that the belief that *p* obtains in all nearby worlds, then our first-order belief that *p* is reflectively lucky – and so the meta-belief that our first-order belief that *p* is safe in the actual world amounts to an assumption.

We can see that general worries about the species of meta-epistemological scepticism, which reflective luck arguments give rise to, may potentially run both ways, to affect the credentials of both internalist and externalist views of justification. Accordingly, the kind of reflective luck internalists wish to eliminate by incorporating an awareness or accessibility condition also risks undermining our confidence in fairly uncontroversial instances of knowledge possession – of the kind that proponents of the "commonsense" approach to philosophy

[61] Ibid., p.30
[62] Ibid.,

defend.[63] If a reflective accessibility requirement (such as RL) were incorporated into our theory as a condition for propositional knowledge, then we would risk denigrating our knowledge of putatively true propositions. What started as a challenge to externalism, therefore, appears to present a challenge for internalists who wish to keep ordinary knowledge of everyday propositions within a demanding intellectualist picture of justification, one that requires reflective accessibility or awareness of the factors that confer justification for a given belief.

In the subsequent chapter, I explore whether a reflective requirement on justification necessarily commits us to an account of knowledge whose dictates cannot be satisfied - I explore ways in which an awareness or accessibility requirement could lead to scepticism, and consider whether mentalist varieties of internalism offer a way of handling that challenge. In order to give up externalism on the basis of the meta-epistemological challenge, we should want to know both what the implications of an internalist analysis would be, and what is potentially destructive about the meta-epistemological challenge. The next two chapters will explore these issues, respectively.

[63] E.g. G.E. Moore, "A Defence of Commonsense", in *G.E. Moore: Selected Writings,* (Routlledge, 1993) pp. 106-133

Chapter 2: Reflection and Epistemic Justification

In this chapter, I will outline how traditional internalist versions of foundationalism appeal to an accessibility condition. On examination, we will see that this accessibility condition presents us with a meta-justificatory regress, the result of which is scepticism in the actual world. This would appear to motivate an externalist view of justification according to which justification does not inferentially depend on having higher-order beliefs about our justified beliefs. However, to accept the externalist analysis at this stage would be too quick, and so I explore ways in which internalists may respond to concerns about meta-justificatory regress. However, I suggest that internalism is ultimately wedded to a demanding version of accessibility, and attempts to construct internalist accounts of justification that appeal to less demanding conceptions of access (such as evidentialism) ultimately succumb to concerns about reflective luck and meta-epistemological scepticism.

The result, I contend, is that we are caught between two different species of scepticism. That is to say, externalism expands the class of non-inferentially justified beliefs, thus eliminating the meta-epistemic regress associated with an internalist accessibility condition, but it has to countenance the idea that justificatory conferring factors for a belief may be opaque from the reflective perspective of the agent. In the subsequent chapter, I present a substantive philosophical challenge to externalism about justification that motivates the meta-epistemological challenge.

Meta-Justificatory Regress

The internalist theory of justification often thought to give rise to scepticism, and often seen as a motivation to externalize our epistemic commitments, is internalist foundationalism. By internalist foundationalism, I am referring to the conjunction of two claims. The first claim is *internalism*, which is the view that justification supervenes on factors that are internally accessible or available to an agent. The second claim is *foundationalism*; which is the view that inferentially justified beliefs ultimately receive their justification from non-inferentially justified beliefs, or so-called "foundational" beliefs. Following Fumerton's schema, we may say that a belief that *p* is inferentially justified if, and only if, its justification is constituted by at least one belief other than *p*, and a belief that *p* is non-inferentially justified if its justification does not consist in having any

other beliefs.[64] Foundationalism states that all inferentially justified beliefs are traceable to at least one non-inferentially justified belief.

Foundationalism therefore presents a linear picture of justification, where inferentially justified beliefs are supported from the "bottom up" by at least one non-inferentially justified belief, or possibly a whole class of non-inferentially justified beliefs.

And therein is the challenge: any linear chain of justification necessitates a non-inferentially justified belief that transmits justification to inferentially justified beliefs; otherwise there is no end to the chain of inferentially justified beliefs. A chain of inferentially justified beliefs could continue ad infinitum, resulting in a justificatory regress. Foundationalism must necessarily posit non-inferentially justified beliefs that transmit justification to inferentially justified beliefs in order to avert a justificatory regress. However, strictly speaking, foundationalism does not *entail* that non-inferentially justified beliefs necessarily exist in the actual world "α", but only that they *in principle exist.* After all, foundationalism is perfectly compatible with scepticism about knowledge in α, if no non-inferentially justified beliefs are forthcoming in α. However, *if* scepticism is false in α, and *if* foundationalism is true, then foundational beliefs necessarily exist in α.

The foundationalist who denies that scepticism is true in the actual world must therefore claim that non-inferentially justified beliefs are forthcoming in the actual world. Such foundational beliefs will be "basic"; these basic beliefs will require no further justification or appeal to further beliefs. Candidate basic beliefs are sometimes logical and mathematical axioms, or other propositions whose justificatory status is considered self-evident. Following John Heil's schema, a belief B is basic just in case it is justified in virtue of its possession of a certain property ϕ, where ϕ is a property of B that does not require any further justification.[65] However, when coupled with the internalist thesis that justification is *necessarily* available to the perspective of a believer, foundationalism is faced with a problem that could undermine the plausibility of basic beliefs by giving rise to a meta-justificatory regress.

To see this problem: pretend S has a basic belief B. In order for S's belief B to be justified on an internalist theory of justification, then the property ϕ that confers justification for B

[64] Fumerton, *Metaepistemology and Skepticism*, p. 56
[65] Heil, John, "Foundationalism and Epistemic Rationality" in *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, (1982), Vol. 42, No. 2, p.180

must necessarily be doxastically accessible to S - at least on the accessibility version of internalism. In this case, our basic belief B will be justified by some property ϕ; where ϕ will not be a property of some other belief about B, but will be a property of B *itself* – and hence B's justification will not depend on further justified beliefs about B (B is justified "in virtue of itself", if you like). However, if internalism is true then S must also have some *perspective* on the justificatory conferring property ϕ of B that allows B to be non-inferentially justified. The challenge for internalism is how to construe this "perspective" on ϕ without relying on some further belief B2 about ϕ, which would make B seem to depend on B2 for its justification; in which case B would appear to be justified inferentially by B2 and hence would not be basic at all. This dependency problem will also arise for B2, which risks potentiating a meta-justificatory regress.

Internalists such as Fumerton recognize this problem, although he states it differently: Fumerton claims that, according to access internalism, in order for a set of justificatory conditions J to constitute S's justification for believing that *p*, it seems inadequate to state that S's access to J is an analytic truth.[66] Consequently, for whenever S is justified in believing that *p* on the basis of J, S must have some access "A" to that condition, as Fumerton notes, such that (J+A). But this principle is not enough for S to have a justified belief that *p*, since one must also have further access (A^2) to A, such that (J+A+A^2);[67] but then A^2 is subject to the same considerations as A, which would require further access (A^3) to A^2; where these epistemic levels of accessibility can be expanded ad infinitum.[68] The problem for types of foundationalism that entail an internalist accessibility requirement is that one's justification for a given belief depends on a higher-order thought or representation that is propositional in character, which is *about* that belief. This higher-order thought, however, resembles a belief, with propositional content, and hence the higher-order thought demands the same kind of access as the first-level belief - which gives rise to a meta-justificatory regress.

The regress problem stems from the fact that an infinite regress of justification results in scepticism; regress leads to scepticism because it is impossible for finite beings to comprehend infinitely long chains of reasoning. The regress challenge to foundationalism therefore states a contingent psychological fact about human beings - and any agent with

[66] Fumerton, Richard, "Epistemic Internalism, Philosophical Assurance and the Skeptical Predicament", in *Knowledge and Reality*, (Springer, 2006), p. 183
[67] Ibid.,
[68] Fumerton, p.183

finite capacities - but one that appears to deny us knowledge in the actual world. This criticism is analogous to the claim that internalism commits us to scepticism because it "over intellectualizes" justification. Over-intellectualization is a potential problem for internalism because over-intellectualization makes justification in principle unattainable for intellectually finite beings. On this view, only beings with infinite capabilities, of the kind religious believers might be inclined to ascribe to God, are capable of justification and knowledge. Thus, internalism appears to come with an explanatory burden: how can we have justified beliefs in the actual world without potentiating a meta-justificatory regress?

Internalist Responses to Meta-Justificatory Regress

The most fundamental way in which internalists can cope with this explanatory challenge is to deny that internalism creates a meta-justificatory regress; if successful, this strategy concedes nothing to externalism. As stated, the explanatory challenge is generated by the internalist accessibility requirement. Therefore, versions of internalism that present a version of accessibilism that are not "over-intellectualized", and which recognize the finitude of our cognitive abilities, will be well disposed to handle the explanatory challenge.

Over-intellectualization results when accessibility requires that one's justified belief that *p* must be *based* on some meta-justification in order to be justified, such that one has a justified first-order belief that p *via* some meta-belief about one's first-order belief that *p*.[69] Framed in this way, it seems that *the justification for believing that p grounds the belief that p*, and hence S comes to have a justified belief that *p* via having justificatory grounds for *p*. [70]The internalist temptation is to think that justification for a given belief that *p* can only be answered with reference to the reasons that actually ground the belief that *p* – which entails over-intellectualization: specifically this conception of access requires us to explain how an agent can have a non-inferentially justified belief B if an agent must necessarily have a perspective on the property that makes one's belief justified.[71] The implication is that such a theory of justification demands an inferential performance that

[69] Smithies, Declan, "Moore's Paradox and the Accessibility of Justification", (2011), *Philosophy and Phenomenological Research*, p.5
[70] Ibid.,
[71] Smithies, p. 4

exceeds our psychological abilities, which will place considerable doxastic limits on the propositions one has justification to believe.[72]

A potentially more promising strategy for internalism may be to insist that first-order justification does not need to be grounded in some meta-justification, but instead obtains for whenever one's doxastic situation confers the logical possibility of forming a justified meta-belief; on this account, one's first-order belief needn't be grounded in that meta-belief.[73] This strategy potentially allows internalists to challenge the more fundamental externalist assumption that internalism faces an over-intellectualization objection. Understanding access in terms of the *logical possibility* of having meta-justification, for example, would allow that in reality we often do not possess the requisite concepts that allow us to form complex meta-beliefs, and so would allow us to avoid an account of justification that would be too demanding or altogether unintelligible.[74] On this alternative account of access, what matters is that one's situation *in principle* allows one to have a meta-justification for believing that one is non-inferentially justified.[75] I will now suggest that evidentialist forms of "internalism" are compatible with this weaker conception of accessibility. Hence, evidentialism offers support for an accessibility requirement; but the kind that averts a demanding meta-justificatory requirement. Although I ultimately contend that abandoning doxastic accessibility presents us with a meta-epistemological challenge that means that mentalist forms of internalism are not using the concept of "internal" in the correct sense – that is, they are nominally internalist, but their theories are externalist in spirit.

Evidentialism and Mentalism

Following Earl Conee and Richard Feldman (1985), we may formulate evidentialism as follows:

(EJ) Doxastic attitude D toward proposition *p* is epistemically justified for S at time *t* if, and only if, having D toward *p* fits the evidence S has at *t*.[76]

Evidentialism is the basic view that justification for S's belief that *p* depends on having evidence E that counts in favour of the belief that *p*. E is evidence in favour of S's belief

[72]Ibid.,
[73] Ibid., p. 5
[74] See: Fumerton, *Metaepistemology and Skepticism,* p.78
[75] Ibid., p. 82
[76] Conee, Earl. Feldman, Richard, "Evidentialism", in *Philosophical Studies,* (1985), Vol. 48, p310

that *p* if, and only if, E counts in favour of the truth of *p*. Likewise, E is evidence against S's belief that *p* if, and only if, E counts in favour of the falsity of *p*. Of course, only if S in some sense "possesses" E can E be relevant to whether S's belief that *p* is epistemically justified. Therefore, whether we are able to avoid the regress that results from doxastic accessibility depends on whether S can possess E without S's belief that *p* being necessarily arrived at, or grounded in, E; that is, without S having to necessarily intellectualize or consciously entertain E, and thus arrive at *p* on the basis of (or via) E. One way evidentialists can achieve this is to understand evidence in terms of one possessing certain mental states that exhaustively determine whether one's doxastic attitudes are justified; this view is called *mentalism* (Conee, Feldman, 1985).[77] We may formulate mentalism as a strong supervenience thesis, following Conee and Feldman[78]:

(*S*) The justificatory status of a person's doxastic attitudes strongly supervenes on the person's occurent and dispositional mental states, events, and conditions[79]

S is a supervenience thesis, since any two agents who are indiscernible in mental respects will necessarily be indiscernible in justificatory respects: the mental therefore exhaustively determines whether one's doxastic attitudes are justified.[80] In contrast, externalism allows that differences in justification can result from contingent non-mental differences, such as the causal or nomological connections responsible for an agent's belief.[81] Externalism therefore allows two agents A and B that are mentally alike to differ in justificatory respects, in virtue of non-mental differences between A and B. Mentalism can be contrasted with externalism because externalism is incompatible with the modal requirement that mentalism imposes: namely, that any two worlds that are indiscernible in mental respects are necessarily indiscernible in justificatory respects. When internalism is understood in these terms, as a supervenience thesis where justification depends on evidence - and where evidence is something that is exhaustively supervenient on mental properties - then we might be able to progress towards abandoning doxastic accessibility.

On this view, strong accessibility can be abandoned because the mental states that justification supervenes on can be construed broadly enough to include states that are

[77] Ibid.,
[78]Conee, Earl. Feldman, Richard, "Internalism Defended", *American Philosophical Quarterly*, (2001), Vol. 38, No. 1, p. 2
[79] Ibid.,
[80] Ibid.,
[81] Ibid.,

doxastically accessible *and* states that are non-doxastically accessible. Non-doxastically accessible mental states might include certain perceptual features within one's epistemic setting. Memories, for example, may serve as candidate non-accessible states. Surely I can be justified in believing that *p* on the basis that I have previously formed the belief that *p* with good reason, and now I continue to believe that *p* based on my disposition to believe that *p*, even if I am not currently aware of or consciously entertaining the reason for *p*? The point is: mental states often support our putative beliefs, even though we are not conscious of these states themselves; that is, without some higher-order thought about the justificatory status of our beliefs. This much seems reasonable: indeed, the idea that an agent can justifiably believe a whole host of propositions, without having to form complex meta-beliefs in every instance, is often the impetus for accepting an externalist theory of justification, one that allows a variety of agents to possess knowledge irrespective of their reflective capabilities.

It might make sense to talk about such perceptual and memorial states as being factors to which one has some kind of epistemic access to in virtue of these factors being within one's grasp or vicinity, even if we do not have some higher-order thought or representation about these states conferring justification for our beliefs. Mentalism is ostensibly still internalist because the various mental states and processes that confer justification are still *internal* to us, namely, because these mental states constitute *us*.[82] When epistemological internalism is understood in terms of justification supervening exhaustively on the mental, then this in principle opens up the possibility for the supervenience base to be enormously broad; understood in terms of supervenience, *any* mental state could in principle serve as a justifier. On this view, no a-priori constraints govern what factors can and cannot confer justification.

The conclusions of mentalism are, however, fairly reserved. For example, accepting mentalism does not tell us *which* mental states instantiate particular justified beliefs, nor how, if at all, the probabilistic or logical relations that obtain between one's evidence for *p* and one believing that *p* play a role in justifying one's belief that *p*. My concern with mentalism is that, in being so reserved, it either concedes too much to externalism and is thus subject to similar concerns about reflective luck and meta-epistemological scepticism, or else it resolves qualms about meta-epistemological scepticism only to succumb to the

[82] I will shortly challenge this view.

same concerns about accessibility (and meta-justificatory regress) as foundationalist varieties of internalism that appeal to doxastic accessibilism.

The problem for mentalism is that it appears not to require that the relations that obtain between a belief that *p* and evidence for *p* are relevant to the justification of one's belief that *p*. That is to say, having established that mentalism depends on evidentialism (the view that one is justified in believing that *p* only if one's belief that *p* fits the evidence one has for *p*) it seems mentalism needs to offer some account of what it means for one's belief to "fit the evidence".[83] The traditional foundationalist or coherentist picture is in a strong position to explicate what it means to be justified in relation to evidence for *p*. For example, foundationalists can claim that one is justified in believing that *p* if one's belief that *p* is appropriately supported by other beliefs. Similarly, coherentism can state that one is justified only insofar as one's belief that *p* coheres with other beliefs within one's belief set. Both foundationalist and coherentist conceptions of epistemic justification can therefore in principle accommodate logical relations and probabilistic relations as relevant to the inferential justification of one's belief (where the only difference between coherentism and foundationalism is that coherentism does not require non-inferentially justified beliefs). However, evidentialism (on which mentalism depends) is not able to offer an account of how these probabilistic or logical relations between our beliefs serve to inferentially justify given beliefs. Accordingly, this is because the logical relations that traditional internalists posit are *not* mental states. And, so, the view that justification supervenes exhaustively on the mental precludes the possibility of justification supervening on non-mental factors that nevertheless seem, at least on the traditionally internalist view, so manifestly relevant to why a subject is justified in believing a given proposition.[84] Framing this challenge, Goldman claims:

"Every traditional form of internalism involves some appeal to logical relations, probabilistic relations, or their ilk. Foundationalism requires that non-basically justified beliefs stand in some suitable logical or probabilistic relations to basic beliefs; coherentism requires that one's system of beliefs be logically consistent, probabilistically coherent, or

[83] Goldman, Alvin, "The Internalist conception of Justification", in *Epistemology: Internalism and Externalism*, Edited by Hilary Kornblith, (Blackwell, 2001), pp. 216-217 – shows that mentalism, in being so permissive, cannot accommodate the logical or probabilistic relations seem relevant to the internalists conception of inferential justification

[84] Ibid.,

the like. None of these logical or probabilistic relations is itself a mental state, either a conscious or a stored state. So these relations do not qualify as justifiers". [85]

Mentalism does not demand that we are necessarily conscious of these logical or inferential relations that, we might think, appear to warrant us in believing a given proposition. Mentalism can either ignore this concern, merely stating that mental states alone confer justification, and that to say otherwise would be to presuppose a commitment to strong accessibilism that mentalism rejects. However, because of this, it is not obvious how mentalism overcomes concerns about reflective luck and meta-epistemological scepticism. Accordingly, if we lack any first-person perspective on how our beliefs are logically or probabilistically connected to factors that are likely to render them true, then for whenever our true beliefs are justified, then given what we are reflectively aware of, their truth-value will be lucky (on the non-standard ordering of possible worlds). Alternatively, mentalism can incorporate the doctrine that the logical and probabilistic connections between one's belief that *p* and the evidence for *p* serve as relevant to one's justification for believing that *p*. However, in which case, we would need to harbour meta-beliefs about these beliefs in order for our thesis to be regarded as internalist. Thus the only plausible account of mentalism is one that faces the same concerns about *accessibilism* and the possibility of potentiating a meta-justificatory regress as foundationalism and analogous forms of internalism that depend on doxastic accessibilism.

To summarise, for evidentialism, an agent S is justified in believing that *p* at time *t* if and only if S's belief that *p* at *t* fits the evidence for *p* at *t*.[86] I have said that the only way of maintaining this thesis without relying on an account of access which does not entail a demanding accessibility condition is to accept mentalism. However, I have said that the mere logical possibility of access is objectionable in the context of supporting mentalist forms of internalism because it allows that the sorts of inferential or probabilistic relations that make a given belief likely to be true, to be *opaque* from the cognitive perspective of the believer.

The lesson of this is that, insofar as the presence of reflective luck might be used to pose a meta-epistemological challenge to externalism, it is not clear that any internalist theory that

[85] Ibid., p. 16

[86] This is similar to Feldman and Conee's formulation of evidentialism. See: Conee, Earl. Feldman, Richard, "Evidentialism", in *Philosophical Studies,* (1985), Vol. 48, pp. 15

tries to avert this problem by incorporating an accessibility requirement will succeed. This is because accessibility requirements lead to infinitely long chains of reasoning and so result in scepticism. As I have tried to show, attempting to overcome this problem by understanding "access" as a logical possibility rather than an entailment of justification, which is consistent with mentalist versions of internalism, and also direct acquaintance theories of justification, such as Fumerton's,[87] fails to avert concerns about meta-epistemological scepticism. Mentalism, for example, offers a framework for accommodating the intuition that internal differences confer justificatory differences. However, in being so accommodating about the factors that can confer justification, and in allowing these factors to remain reflectively inaccessible, mentalism does not speak to the relevant notion of "internal" that the meta-epistemological challenge to externalism motivates. The right concept of "internal", according to the meta-epistemological challenge, cannot be understood in terms of biological, physiological or any otherwise natural and potentially reflectively opaque states – our internal states should be restricted to states that *cannot* be analysed naturalistically (Chapter 3 examines this further).

A naturalistic analysis is compatible with the view that the relevant justifiers for a belief are a function of nomological relations in the external world, where these relations will continue to confer justification regardless of our perspective on the relevant processes that allow those states to confer justification. There is thus a close connection between the meta-epistemological challenge to externalism and views about what sorts of epistemic properties we should be invoking in our epistemology (I explore this connection in the subsequent chapter). As I will show, it is clear that many of the philosophical objections that apply to externalism will be applicable to internalists who opt for a naturalistic construal of what constitutes an epistemic property such as justification.

For the time being, we have a stand-off. In preserving an accessibility requirement, we have a theory of justification that over-intellectualizes justification, which results in a meta-justificatory regress. The alternative is that we analyse justification in third-person terms, where the factors that confer justification for a given belief may be opaque nomological relations or mental states. However, the consequence of this sort of analysis is meta-epistemological scepticism of the kind reflective luck arguments emphasise.

[87] Fumerton, *Metaepistemology and Skepticism*, p.73

However, externalists may stress that in light of the problem of over-intellectualization, and the intuition that we *do* have putative knowledge of propositions such as "I have hands", then what support could there be for the view that knowledge entails a demanding reflective requirement? The apparent simplicity with which externalists can deal with internalist sceptical regress might make it more appealing, all-things-considered. Externalists can simply deny the validity of the inference that an agent fails to have knowledge or justified belief in virtue of not having reflective access to the factors that confer justification. The result is that externalism is able to present an account of justification whose dictates can be satisfied by a wide range of subjects, with differing reflective capabilities, thus expanding the class of non-inferentially justified beliefs and satisfying the view that we can posses knowledge of putatively true propositions.

In what follows, I consider the substantive philosophical objection to externalism that the meta-epistemological challenge motives. I argue that meta-epistemological challenges pose a more fundamental challenge to the way in which externalism naturalizes the concept of justification.

Chapter 3: The Philosophical Objection to Externalism

In this chapter I will firstly argue that the meta-epistemological challenge to externalism is best understood as motivated by a "felt need" for epistemic priority; that is, the need for S to know that K is a knowledge source for S in a way that does not partially or wholly depend on K.[88] The notion of priority can give us insight into the nature of the meta-epistemological challenge to externalism – and I argue that externalism will struggle to meet this requirement. I then argue that there is no obvious way in which internalism can incorporate an account of knowledge that satisfies our felt need for epistemic priority, of the kind that Barry Stroud (1989) highlights,[89] and that the appeal of externalism is the ease with which it provides us with non-inferentially justified beliefs. The result is that we are caught between two intuitions. The first intuition is the desire to offer a philosophical account of knowledge, which motivates us to resist naturalism. The second intuition, which I defend in the subsequent chapter, is the desire to make knowledge relatively easily attainable; this latter intuition is anti-sceptical and supports externalism.

Second-Order Knowledge

We have established that the charge of meta-epistemological scepticism (MS) is typically levied against externalist accounts of knowledge because it allows that one can know that *p* without knowing, either via reflection or any other means, or having any first-person perspective, or any otherwise propositional attitude towards, one's knowledge that *p*. [90] The crux of the meta-epistemological challenge is this: our true beliefs may satisfy an externalist requirement for knowledge (such being reliably true, safe, or sensitive[91]), but the externalist analysis of knowledge gives us no reason to conclude that our beliefs *are* reliable or safe, and so has to countenance meta-epistemological scepticism.[92] How should externalists respond to the meta-epistemological sceptical challenge?

It is tempting to think that having knowledge that one knows that *p* would overcome meta-epistemological worries about whether we *have* externalist knowledge. That is to say, it may be thought that we can eliminate meta-epistemological scepticism if our theory of

[88] For a discussion of this "felt need" for epistemic priority, See: Barry Stroud in *Understanding Human Knowledge in General*, (1989), pp. 130-13. Please note, "felt need" is originally Stroud's term.
[89] Ibid.,
[90] See Fumerton (1995), Zagzebski (1996), BonJour (2003), Pritchard (2005)
[91] For our purposes, it doesn't matter which formulation of externalism we adopt.
[92] This is closest to Fumerton's account. See: Fumerton, Richard, *Metaepistemology and Skepticism*, p. 175

knowledge allows that this second-order knowledge is possible (where "allows" needn't entail that one has such second-order knowledge). If externalism can give us knowledge that we know, then can it resolve the meta-epistemological challenge? I want to show that responding to meta-epistemological scepticism in this way would be too simplistic, and would give an incomplete picture of what the meta-epistemological challenge to externalism consists in. To see how incorporating the possibility of second-order knowledge fails to resolve the meta-epistemological challenge, let us firstly see how externalists may be able to give us second-order knowledge by introducing a KK principle – I will then show that this principle can be satisfied whilst leaving the meta-epistemological challenge unaddressed; this will go some way to showing that the meta-epistemological challenge consists in a more substantive philosophical challenge to externalism.

There are different versions of KK, but let's take the following principle as paradigmatic:

(KK) K is a knowledge source for subject S at time *t* if, and only if, S knows at *t* that K is a knowledge source for S at *t*[93]

Let us further stipulate that on an adequate post-Gettier epistemology, K could only be a knowledge source for S in virtue of K consisting in the relevant sorts of nomological processes that allow S to have true beliefs that are free from veritic luck. On the above schema, the second-order knowledge can be provided by externalism by way of stating that K depends on some further knowledge ("K2") that allows S to know that K is a knowledge source for S, on the grounds that K2 (much like K) allows S to have a true belief (whose propositional content is *that K is a knowledge source for S*) that satisfies some externalist requirement – such as being reliably true, safe, sensitive, etc. On this externalist template, our putative second-order knowledge about our first-order knowledge, much like our first-order knowledge, could identically satisfy the same externalist requirements that regulate our first-order knowledge. Consequently, there is nothing necessarily internalist about the notion that knowledge entails having a second-order perspective on that knowledge. Instead, KK is merely a level requirement that can be satisfied by internalists and externalists alike.

[93] This formulation is used by Cohen, Stewart, "Basic Knowledge and the Problem of Easy Knowledge", in *Philosophy and Phenomenological Research*, Vol. 65, No. 2 (2002), pp.309. –my formulation of the KK principle is modelled on his "KR principle".

If the meta-epistemological challenge consists in the lack of second-order knowledge that one has externalist knowledge, then the externalist may therefore respond by way of introducing a KK principle, as shown. However, the problem with this response is that the externalist way of giving us second-order knowledge via KK does not look very philosophically interesting: the second-order knowledge is not of the right kind. Specifically, insofar as we want to satisfy our curiosity as to *which* particular true beliefs are instances of knowledge possession, there is nothing within S's reflective vicinity that could offer S any *assurance* that K is a knowledge source (for S) on an externalist version of KK. Insofar as KK is an attempt to meet a legitimate philosophical concern about the reliability of our belief forming processes, a concern that the meta-epistemological challenge to externalism draws our attention to, then the externalist response to such a challenge will be plainly question-begging, since it is invoking the same processes to explain the very processes that are under philosophical scrutiny with regards to our first-order knowledge. Consequently, the sort of meta-epistemic concerns that arise from a philosophical curiosity as to whether we have first-order knowledge that *p* will similarly afflict externalist second-order knowledge.[94]

When we inquire about whether our beliefs are instances of knowledge, we are trying to satisfy what Fumerton regarded as a certain "philosophical curiosity" about the legitimacy of our beliefs.[95] Insofar as the meta-epistemological challenge arises from a philosophical curiosity about the status of our first-order beliefs, then the answer to the question of whether our first-order beliefs are legitimate cannot be supplied by the sorts of processes that give rise to our first-order beliefs; such knowledge would not give us any assurance that our first-order beliefs are knowledge instantiations.[96] The idea that meta-epistemological scepticism is motivated solely by a commitment to a KK principle therefore misses something out: namely, that the thing that motivates meta-epistemological scepticism in the first place is a desire to have assurance that K is a knowledge source for S.

The sort of second-order knowledge the externalist proposes is not able to help us make the right sort of discriminations about our purported first-order knowledge, and so is unable to eliminate meta-epistemological scepticism. That is to say, externalism does not help us discern or differentiate, via reflection, particular instances of knowledge possession *as* cases of knowledge possession – that is, it cannot help us separate our true beliefs from cases of

[94] Fumerton, "Skepticism and Naturalistic Epistemology", *Midwest Studies in Philosophy*, (1994), Vol. 19, Issue. 1, p pp. 336-337
[95] Ibid., p. 338
[96] Ibid.,

knowledge. But nonetheless, there is nothing to stop the externalist responding in such a way. This is because externalist second-order knowledge can legitimately be subject to the same requirements as externalist first-order knowledge – and one of these requirements is that the factors that allow one to have such knowledge are not *necessarily* reflectively accessible to that agent. There is, in other words, no a-priori restriction on using the same processes responsible for converting our first-order true beliefs into knowledge as a way of also turning second-order true beliefs (whose propositional content is about the epistemic status of our first-order beliefs) into instances of second-order knowledge.[97] Such a strategy is bound to look viciously circular, but a purely externalist analysis of knowledge allows for it. In the same way that there are no *a-priori* restrictions on the external processes that allow us to have first-order knowledge, there are similarly no *a-priori* restrictions on what processes allow us to have meta-knowledge.[98]

The internalist will therefore state that externalism neglects legitimate philosophical concerns about the reliability or justificatory status of our first-order beliefs; questions about the reliability of our first-order beliefs, or about whether our first-order beliefs are cases of knowledge, appear to be beyond the purview of philosophical investigation. For example, William Alston thinks there is a philosophical problem with process reliabilism:

"when we ask whether one or another source of belief is reliable, we are interested in discriminating those that can be reasonably trusted from those that cannot….merely showing that if a given source is reliable it can be shown by its record to be reliable, does nothing to indicate that the source belongs with the sheep rather than with the goats".[99]

Implicit in this criticism of externalism is that it does not offer us an account of epistemic priority.[100] The notion of epistemic priority is compelling precisely because we are questioning whether we have externalist knowledge in the first instance – and we want to offer an answer to that question that does not beg the question. As philosophers, we are *not* taking that knowledge for granted and then attempting to see whether it can give us second-order knowledge. Given that this is how we approach first-order knowledge, the justificatory

[97] See Fumerton, *"Skepticism and Naturalistic Epistemology",* p.336
[98] Ibid. As Fumerton notes, "there is no a-priori reason why the conditions required for higher-level justified belief and knowledge might not be satisfied", p.336
[99] Fumerton, *Metaepistemology and Skepticism*, (1995) p.179 (Original Source: Alston, William, in *The Reliability of Sense Perception* (Cornell University Press, 1993), p.17
[100] See: Stroud, Barry, "Understanding Human Knowledge in General", in *Epistemology: Internalism and Externalism*, Edited by Hilary Kornblith, (2001), pp. 130. Originally published in *Knowledge and Skepticism*, Edited by M. Clay and K. Lehrer (Westview Press Inc., 1989)

grounds that allow us to know that we have knowledge will need to be independent of the knowledge in question. More precisely, I mean to say that if one has justification J on which one bases the true belief that one has knowledge that *p*, where J makes this belief count as knowledge, then J is only non-question-begging if it is antecedent to the justification (call this J2) that allows one to have knowledge that *p*.[101]

Barry Stroud (1989) has addressed a similar concern, and outlines that the kind of assurance we seek in epistemology requires something akin to epistemic independence, which he terms "epistemic priority", which the externalist project is unable to meet.[102] Accordingly, Stroud claims.

"The apparent dilemma is a familiar quandary in traditional epistemology. I think it arises from our complete general explanatory goal. We want to explain a certain kind of knowledge, and we feel we must explain it on the basis of another, prior kind of knowledge that does not imply or presuppose any of the knowledge we are trying to explain. Without that, we will not be explaining the knowledge in question in the proper, fully general way. This felt need is what so easily brings into the epistemological project some notion or other of what is usually called *epistemic priority* – one kind of knowledge being prior to another. I believe it has fatal consequences for our understanding of our knowledge."[103]

The concern that we offer a philosophically satisfying conception of second-order knowledge is motivated by this same "felt need" to have prior knowledge that a particular belief source is a knowledge source for a given subject.[104] As I have tried to show, we can satisfy a basic KK principle, using externalism, whilst leaving this felt need frustrated. The advantage of understanding meta-epistemological scepticism as motivated by satisfying the principle of epistemic priority rather than KK is that it allows us to state why we cannot resolve the problem of meta-epistemological scepticism by relying on the very knowledge sources that are under philosophical scrutiny, and so it allows us to stress what is philosophically unsatisfying about the externalists response to meta-epistemological scepticism. Instead, to have assurance that K is a knowledge source, of the kind we need to eliminate concerns about meta-epistemological scepticism, one must necessarily know that K is a knowledge source where this knowledge is not dependent, partially or wholly, on K.

---

[101]Pryor, James, "The Skeptic and the Dogmatist", *Nous,* (2000), Vol. 34, No. 4,, p. 524
[102] Stroud, *Understanding Human Knowledge in General*, p.130
[103] Ibid.,
[104] Ibid.,

However, finding some test for whether K is a knowledge source that does not at least partially depend on K is difficult. The kind of independence we are seeking is elusive! For example, we cannot use the external world as a basis on which to give us knowledge that we have knowledge of the external world if we want to provide an account of knowledge that gives us assurance that we *can* know that there exists an external world. But, equally, it appears impossible to offer an account of how we know some object that is not partially dependent on the very processes under philosophical scrutiny. Or as Stroud remarks,

"…if we really are restricted in perception to experiences or sense datum…which give us information that is prior to any knowledge of objects, how *could* we ever know anything about what goes on beyond such prior data?"[105]

The problem is that we are restricted in our everyday life to the very processes that give us first-order knowledge. This restriction results because there is a worldly element to knowledge: knowing that *p* depends both on believing that *p*, and perhaps, for the internalist, on reflecting that *p*, but it also necessarily depends on one's belief that *p* cooperating with the world in such a way that one can have a true belief that *p* that is reliable or safe etc. However, it is unclear how we are to have reflective knowledge that the world, and its nomological properties and relations, are functioning in a certain sort of way that allows one's beliefs to display this cooperation (and ultimately, to count as knowledge). As established, we cannot rely on those very processes to explain that the world is functioning in a way conducive to knowledge, and so an externalist explanation that those processes convert our true beliefs into knowledge would appear question-begging. However, it is equally implausible, if not more so, that we should know from the philosophical armchair that a given contingently reliable process is contingently reliable.

Denying Restriction

One strategy that could preserve externalism, whilst establishing epistemic priority, is to deny that we are restricted in the way that I suppose. In other words, it may be remarked that what I am depicting as an epistemic gulf, between two different types of knowledge, is in fact a perfectly plausible relation that obtains. This may allow us to know *a-priori* that our beliefs satisfy some externalist requirement for knowledge.

[105] Ibid., p.131

We should therefore add a caveat about a-priori knowledge. A-priori knowledge is knowledge that does not depend on experience, and so knowing a-priori that K is a knowledge source would not require that we know K is a knowledge source via any nomological or causal processes. However, insofar as K is a knowledge source that allows S to have the sort of knowledge post-Gettier epistemologists are interested in, then K must satisfy some externalist modal requirement that eliminates veritic luck; such that K is a knowledge source for S if, and only if, K allows S to have true beliefs across a wide range of counterfactual circumstances. However, a-priori knowledge (that K is a knowledge source for S) needn't depend on any such externalist requirements, and so a-priori knowledge (that K is a knowledge source for S) may be suitably independent of K to capture the relevant notion of epistemic independence that motivates the meta-epistemological challenge. Therefore, a-priori knowledge (that K is a knowledge source for S) would not succumb to the explanatory circularity that externalist attempts to accommodate meta-epistemological scepticism have to countenance.

However, there is a potential problem with the proposed solution. Namely, the sorts of "facts" that basic beliefs have as their propositional content are typically mathematical axioms or modal truths, or other non-contingently true propositions. A-priori knowledge "that K is a knowledge source" requires *a-priori knowledge of some contingent fact* (the fact that K is a knowledge source for S). Accordingly, the fact that K is a knowledge source is a purely contingent matter because K is only a knowledge source *in the actual world,* due to K satisfying some externalist requirement. It is a contingent fact that K is a knowledge source because of some possible world "β" that differs in nomological respects to the actual world, where K yields false beliefs at β.

No doubt a wealth of knowledge about contingently true propositions *is* furnished by experience. However, the charge that knowledge of contingent facts is necessarily *a posteriori* is a stronger epistemological claim that may be resisted. For example, Saul Kripke (1980) thinks that a-priori knowledge of contingently true propositions is possible.[106] The crux of Kripke's argument is that we can have contingent a priori knowledge because propositions that are knowable a-priori are not always rigidly designated, but may instead be fixed by some contingently true state of affairs.[107] By way of example, we are invited to consider the following scenario: let us suppose that "one meter" is the length of S, where S is

[106] Kripke, Saul. A. *Naming and Necessity*, (Oxford University Press, 1980)
[107] Ibid.,

a certain stick or bar in Paris.[108] We should then ask, "Is it a necessary truth that stick S is one meter long at time *t*?" [109]

Kripke doesn't think so, but he nonetheless maintains that we can have a-priori knowledge that S is one meter long at *t*.[110] That this is so, according to Kripke, is because "one meter" is intended to rigidly designate a certain length in all possible worlds, which in the actual world is the length of S at *t*. However, "the length of S at *t*" is not a statement that designates rigidly.[111] Thus, "one meter" and "the length of S at *t*" are not synonyms, even though we have, according to Kripke, "determined the reference of the phrase", by stating that "one meter" is the rigid designator of the length of S at *t*.[112] Therefore, once we have determined "one meter" as to mean "the length of S at *t*", we can know a-priori that S is one meter long at *t*, because it has been arrived at "automatically and without further investigation".[113] For Kripke, the statement that stick S is one meter long at *t* is knowable a-priori even though it has the metaphysical status of being only contingently true – there are, after all, plenty of possible worlds in which we fix the reference of "one meter" differently.

Whether a-priori knowledge of contingent facts is possible is beyond the scope of this thesis. Consequently, we ought to concede for the sake of argument that it is, and therefore consider the more essential question of whether a-priori knowledge of contingent facts (if possible) would allow us to present an account of knowledge that both eliminates MS and resists scepticism about first-order knowledge. As a variety of knowledge, second-order a-priori knowledge that K is a knowledge source may allow us to resist our initial paradox, by making this prior second-order knowledge non-inferential. That is to say, the status of K as a knowledge source for S is metaphysically contingent, and so a Kripkean analysis could in principle state that we can have a-priori knowledge of this contingent fact in a way that does not depend on K. As a dialectic move, second-order a-priori knowledge at least looks an attractive coping strategy.

However, there is a disanalogy between Kripke's example of knowing a-priori that S is one meter long at *t*, which is about first-order knowledge, and the case of knowing that K is a knowledge source for S, which is about second-order knowledge. Although both cases ("the

[108] Kripke, Saul A., *Naming and Necessity*, (Oxford University Press, 1980), p. 54
[109] Ibid., p. 54
[110] Ibid.,
[111] Ibid., p.56
[112] Ibid.,
[113] Ibid.,

length of S at *t*" and "K as a knowledge source for S") denote instances of knowing metaphysically contingent facts, the fact that S is one meter long at *t* obtains in virtue of an agreement that "one meter" rigidly designates the length of S at *t*. However, no such designation is forthcoming in the case of K being a knowledge source for S. In other words, *that* K is a knowledge source may be metaphysically contingent, but its truth-value is not something that has been decided or fixed by *us*, but is instead a function of external nomological properties and relations: the consequence is that K may be a knowledge source even if we are completely ignorant of this fact; whereas it would be nonsensical to suggest that "the length of S at *t*" could rigidly designate "one meter" in the absence of some awareness or appreciation that this is so. When the truth-maker is not *us*, but is instead *the world*, as is the case with second-order a-priori knowledge, then the idea that we can know that K is a knowledge source "automatically", or without empirical investigation, falters.

How, then, are we to bridge the epistemic gulf that would appear to be imposed by the requirement that in order for S to know via K, S must further know via reflection, that K is a knowledge source for S? Having already rejected second order a-priori as an explanation for how such knowledge is possible, to continue to support the thesis that we can know, via reflection, external matters of fact would be to commit us to an un-explained relation. This relation would, in the absence of any explanation, appear mysterious.

The Internalist Response

Internalism can agree that it is hardly surprising that we should not be able to know via reflection that some external condition obtains - indeed, this is the central reason why meta-epistemological problems arise for externalism. However, internalism and externalism will draw different lessons from this dilemma. For externalism, the difficulty with which we can satisfy our "felt need" to offer an explanation or reason for thinking that we have externalist knowledge (where that reason is not contingent on the processes under scrutiny), shows that attempting to do so is a misguided pursuit - because were knowledge to depend on satisfying this felt need, we have scepticism (and since scepticism is false, then the meta-epistemological challenge fails). Alternatively, for internalism, the apparent difficulty with which externalism can *in principle* have this "richer" second-order knowledge may show that externalism about justification is simply analysing the wrong philosophical concepts or

properties. [114] On this latter view, the fact that externalism begs the question about our first-order knowledge shows that externalism is flawed. With the right internalist analysis of justification, the internalist can offer an account of justification that satisfies our felt need, the thought goes.

There are two elements to what I am suggesting. The first is that objections such as Stroud's and Fumerton's draw our attention to a fundamental problem with the externalist analysis of knowledge or justification – one that is at the heart of the meta-epistemological challenge. The second element is that internalist accounts of justification might be able to offer us the resources to show that we can have an account of justification that does not lead to scepticism and which also depends on the "correct" conception of what constitutes an epistemic concept or property. These two elements are distinguishable, since the failure of the relevant internal perspective results only in scepticism – and a sceptical conclusion alone may not be enough to undermine the legitimacy of philosophical objections to externalism. I will consider these two arguments in turn.

1. Naturalism and The Philosophical Objection to Externalism

The objection to the externalist analysis of knowledge is that it is not "philosophically interesting" (Bergmann, 2000) because it allows us to depend on the very processes under philosophical scrutiny to vindicate a reliance on those processes.[115] However, what makes an analysis of knowledge "philosophically interesting"? And why does the correct analysis of knowledge need to be of philosophical interest? An analysis of knowledge that results in circularity is philosophically un-interesting, internalism can allege. However, the externalist thinks that questions of what we know are answerable by analysing nomological relations in the world. He, presumably, thinks that such processes are philosophically interesting insofar as they determine whether we are justified, and that any other sense of "philosophically interesting" is parochial to the question of what we know. Getting a handle on what it means for an analysis of knowledge to be "philosophically interesting" is thus easier said than done. As Bergman's analysis of this issue claims, a philosophically interesting epistemic property will be the kind of property whose exemplification is at issue in the controversy between

[114] This is the objection that Michael Bergmann explores. See: "Externalism and Skepticism", *The Philosophical Review, (2000)* Vol. 109, No. 2, 159-194. For an example of the philosophical objection to externalism, see Fumerton, *Metaepistemology and Skepticism*, p. 159-183

[115] Bergmann, Michael, "Externalism and Skepticism", *The Philosophical Review, (2000)* Vol. 109, No. 2, p.163

sceptics and non-sceptics.[116] The difficulty, however, arises in identifying precisely what *sort* of property is being exemplified in this controversy.

The particular epistemic property assumed to be at stake in turning true belief into knowledge is *justification*. However, as Gettier showed, the property of justification is not *sufficient* for turning true belief into knowledge. We may therefore need some further condition in addition to justification, or otherwise offer an analysis of justification that *is* able to turn true belief into knowledge. We may also want to call the relevant epistemic property different things - such as *warrant*. At any rate, we will have to be somewhat vague here: what matters is that there is *some* epistemic property on which knowledge depends being exemplified in the debate between internalists and externalists – belief is manifestly not that property, because it is a psychological state, and truth is more properly describable as metaphysical or semantic rather than epistemic.[117] The epistemic property that is being exemplified is *justification*, or something that approximates it. The philosophical objection to the externalist analysis of justification is that its analysis of justification is not philosophically interesting. More specifically, the externalist either has the wrong analysis of justification or ends up offering an analysis of some property other than justification.[118] The crux of the philosophical objection to externalism is that there is no philosophically interesting account of an epistemic property that analyses that epistemic property in terms of natural properties.

Rejecting a naturalized epistemology requires, however, some conception of what constitutes a "natural" property. Some have understood the relevant "natural properties" at stake in terms of any property that the sciences feel "comfortable" analysing.[119] However, this is not a fully satisfactory way of characterizing what constitutes a natural property, since scientists may on occasion feel comfortable "analysing" properties that cannot be analysed naturalistically; and so willingness for science to analyse a particular property does not thereby make it natural. Instead, I suggest, only *success* in analysing some property in natural terms is sufficient to make it a natural property. The internalist contention is that a naturalistic analysis of justification is not successful because it is not philosophically interesting.

---

[116] Bergmann, p. 165

[117] Kim, Jaegwon, "What is Naturalized Epistemology?" in *Epistemology: An Anthology*, (Blackwell, 2008), p. 539

[118] Bergmann, p. 160-163

[119] Fumerton, Richard, "Skepticism and Naturalistic Epistemology", *Midwest Studies in Philosophy*, (1994), Vol. 19, Issue. 1, pp. p. 321

The proposal of what constitutes a philosophically interesting property runs as follows: where there is some room for invoking an analysis of that property that is not exhaustively naturalistic, then that analysis is philosophically interesting; that the philosopher in principle has a particular competence or expertise in analysing that property, that the scientist lacks, should indeed make it an enticing property to analyse.[120] Where the property is of a kind that can be reducible to, or is identical with, the sorts of properties that can be analysed naturalistically, by science, or the so-called social sciences, then that property is not of a *particular* philosophical interest, and is interesting only insofar as one has a more general extra-curricular interest in that property.[121]

The statement that no naturalistic analysis of some epistemic property can be philosophically interesting has strong consequences. For one thing, it is a general objection to naturalizing all properties or concepts that fall under the philosophical purview, and thus is only an objection to *externalism* insofar as externalism depends on a *naturalistic* analysis of knowledge. A very real implication of rejecting naturalism about the sorts of epistemic properties that are relevant to knowledge would therefore be to force us to reject naturalism about other kinds of properties, such as moral or aesthetic properties; otherwise singling out *epistemic* properties as especially problematic would seem ad-hoc. For our purposes, however, I think my suggested interpretation of what constitutes a philosophically interesting account of an epistemic property – whether palatable or not – is consistent with the conception of knowledge that is often employed by internalism in motivating the meta-epistemological challenge to externalism. For example, Fumerton claims:

"The fundamental objection to externalism can be easily summarized. If we understand epistemic concerns as the externalists do, then there would be no objection in principle to using perception to justify reliance on perception…But there is no philosophically interesting concept of justification or knowledge that would allow us to use a kind of reasoning to justify the legitimacy of using that reasoning. Therefore, the externalist has failed to analyse a philosophically interesting concept of justification or knowledge."[122]

Fumerton's argument against externalism is consistent with the anti-naturalism espoused by my proposed definition. He sees that the sort of objection he proposes is a broader rejection of Quine's attempt to naturalize epistemic properties.[123] Reducing questions concerning

[120] Fumerton makes a similar remark. See: *Metaepistemology and Skepticism*, (Rowman and Littlefield, 1995) p.171-172

[121] Epistemological naturalism may also threaten any philosophically interesting conception of metaphysics.

[122] Fumerton, Richard, *Metaepistemology and Skepticism*, p.180

[123] Ibid., p180

justification to questions concerning complex nomological facts does not privilege any particular reflective perspective.[124] Fumerton is explicit in his rejection of this approach, since questions about "complex causal conditions that determine the presence or absence of justification for a belief are a subject matter of empirical investigation that would take the philosopher out of the easy- chair and into the laboratory.[125]

A naturalistic approach to epistemology, such as Goldman's or Quine's, employs very different methodological assumptions to internalist and characteristically "aprioristic" approaches to epistemology such as BonJour's. [126] For better or worse, Fumerton and his contemporaries are right: naturalism *does* take the philosopher out of the easy chair, since it emphasises the relevance, and indeed essentiality, of disciplines such as cognitive science and neuroscience in our epistemological theorizing. In contrast to the internalist approach, naturalizing our epistemic principles does not confer upon a cognizer *any* special privilege or responsibility in relation to the epistemic status of one's own beliefs.[127] Externalism is committed to a naturalistic analysis of justification precisely because it makes justification a function of these relations. Internalism allows epistemological inquiry to proceed from the armchair, whereas externalism affords no such luxury.[128]

It is difficult to see how we can have a philosophically interesting conception of "justification", or indeed any conception of justification at all, that reduces epistemological inquiry to a laboratory investigation. (This claim does not necessarily lead to scepticism, as I note in ch.5, because externalists can focus on an analysis of *knowledge* rather than justification). If knowledge depends on justification, then our understanding of justification will be circular and philosophically dissatisfying, because it will do nothing to abate meta-epistemic concerns about whether we *do* have knowledge, and whether we are justified in believing that the relevant justificatory processes for knowledge are in place. This is because naturalism denies the legitimacy of what Fumerton and Stroud appeal to in their arguments against externalism. On the naturalistic view of knowledge, our felt need for epistemic independence or priority is misguided – and radically so. Quine, for example, urges us to appeal to the very natural and nomological processes under philosophical scrutiny in order to

[124] For an analogous claim, see Fumerton, *Metaepistemology and Skepticism*, p.171
[125] Ibid., p171
[126] Kim, Jaegowon, "What is Naturalized Epistemology?" in *Epistemology: An Anthology*, (Blackwell, 2008), p545
[127] Kim, "What is Naturalized Epistemology", p. 547. Also see, Fumerton, *Metaepistemology and Skepticism*, p, 171
[128] Fumerton, p. 171

vindicate a reliance on those processes. That this is so, is because on the naturalistic analysis, epistemology *is* a branch of science.[129] Quine says,

"…such scruples against circularity have little point once we have stopped dreaming of deducing science from observations. If we are out simply to understand the link between observation and science, we are all advised to use any available information, included that provided by the very science whose link with observation we are seeking to understand."[130]

Quine's move should seem extraordinarily provocative. He thinks scientists can adequately answer what we ordinarily think are philosophical questions. On this view, questions about epistemic priority are worthless.[131] One type of knowledge A will gain epistemological priority over B for subject S just in virtue of A's causal proximity to S's sensory receptors being more direct than B's causal proximity to S. [132] Not only are we permitted to use the processes under philosophical scrutiny to determine whether the causal relationships in question give us knowledge, but we are moreover encouraged to use those processes.

The motivation for resisting naturalism is that it cannot offer a philosophically interesting account of what it means to be justified. However, we need to be careful in delineating the particular implications of rejecting naturalism. Importantly, the rejection of naturalism still allows that the scientist is in a position to make a contribution towards the question of what particular beliefs are reasonable – he may be able to grasp how our beliefs are related to facts in the world, for example – but importantly, he will not be able to construct an account of what it *means* to be justified; he simply will not be permitted to engage in this meta-epistemological question. The meta-epistemological question is about whether some subset of natural features – consisting of nomological relations etc – *does* instantiate the property of being justified, and so invoking those processes as an answer to that question would beg-the-question.

However, accepting anti-naturalism does not entail the idea that natural properties are not relevant to particular instances of justification. Accordingly, the anti-naturalist only needs to

[129]Quine, W.V. "Epistemology Naturalized" in *Epistemology Anthology*, Blackwell, 2008), p. 531. Originally published in W.V. Quine, *Ontological Relativity and Other Essays* (New York, Columbia University Press, 1969)
[130] Quine, *"*Epistemology Naturalized", p. 531
[131] Ibid.,
[132] Ibid., p. 534

claim that justification is not *reducible* to natural properties.[133] For example, an intuitively plausible principle is that natural differences confer justificatory differences (and vice-versa). In modal terms, any two worlds that are indiscernible in natural respects will be indiscernible in justificatory respects. Rejecting naturalism does not require that we reject the intuitively plausible principle that natural differences are responsible for justificatory differences. The rejection of a naturalized epistemology merely states that justification cannot be *reduced* to natural features on which it supervenes. The idea that we can carry out a *reduction* of justification to natural properties is the central tenet of a naturalized epistemology, which is altogether different from the assertion that justification depends or supervenes on natural properties, and so it is the former reductionist strategy that internalism should reject as misguided.

Naturalism is in a strong position to capture the relevance of natural features to justification: it simply says that to be justified *is* to instantiate certain natural properties in which justification consists. Externalism can rightly point out that there is no sense to be made of the idea that a belief could be justified if it had no bearing or relationship to natural properties. The point is: it seems an intuitive feature of thinking about justification that we are thinking about some subset of natural features that instantiate justification, that is, some subset of natural properties that are *responsible* for making one's belief justified. This raises a potential objection to anti-naturalism: how can anti-naturalists capture the responsibility that natural properties have in making one justified?

To answer this question it will be instructive to draw an analogy with the relationship between moral properties (if there are such entities) and natural properties. Accordingly, it is often thought an a-priori feature of our moral thought that action X is wrong *because* or *in virtue of* causing pain, or *because* of some other natural feature or set of features that instantiate the property of wrongness in X (these natural properties are *responsible* for X being wrong – that is, for X instantiating the moral property of wrongness).[134] Naturalists about morality can capture the intuition that moral properties are instantiated by natural properties by making those moral properties reducible or identical to the natural properties that are responsible for X being wrong. Consequently, any moral realist who rejects

[133] We can reject naturalism and keep a supervenience relation. See: Kim, Jaegowon, "What is Naturalized Epistemology?", *Epistemology: An Anthology*, p547-548

[134] E.g. See: Zangwill, Nick, "Moral Dependence", in *Oxford Studies in Metaethics,* (Oxford University Press, 2008), Volume 3, p.1-2

naturalism will need to posit a relation between moral properties and the natural properties responsible for instantiating moral properties. That is to say, non-natural realism will have to maintain that whenever a moral property is instantiated, there will necessarily exist some sub-set of natural (and so, on his view, non-normative) properties responsible for that instantiation, where these non-normative properties "exhaustively constitute" moral properties but are not identical to them.[135] Realism will need to show that the dependence of moral properties on natural properties should be construed in terms of supervenience, according to which there needn't be an identity relation between the base properties and the supervening moral property.

Similarly, if we reject naturalism then whatever form of internalism we subsequently adopt will need to posit something akin to supervenience if it is to preserve the intuitive feature of thinking about justification that externalism captures.[136] How we understand this supervenience or dependence relation is up for debate, and I will have to leave it as something of an open question. However, a global conception of supervenience would entail that any two worlds that differ in justificatory respects will differ in *some* natural respects. The differences between any two worlds, on the global interpretation, would not necessarily be perceptible (perhaps just one molecule is enough to bring about a justificatory change); and so no features of an agent's local doxastic situation would need to be radically altered in order to confer a justificatory difference on the global construal of supervenience. The notion that there is *some* natural difference – however slight - between two worlds that are discernible in justificatory respects could be all we need to capture the intuitive feature of thinking about justification; namely, that natural differences confer justificatory differences (the reality, however, may be that any two worlds that differ in justificatory respects differ in terms of *local* natural features[137]).

2. The Problem for Internalism

The challenge for internalism is to offer an account of internalist justification that does not lead to first-order scepticism. In rejecting naturalism as a way of analyzing justification, internalism may find this difficult (this is not the only challenge – accounting for the

[135] Shafer-Landau, Russ, *Moral Realism: A Defence*, (Oxford University Press, 2005), p.76
[136] Jaegwon, Kim, "What is Naturalized Epistemology", pp. 547-549
[137] This is likely to be a contingent rather than a logical requirement.

metaphysics of supervenience is another!). Nonetheless, if internalists cannot at least succeed in presenting an anti-sceptical theory, then internalism will come with a fundamental explanatory burden: namely, why should we find first-order internalist scepticism more enticing than a naturalistic externalist analysis of knowledge that does *not* lead to scepticism? [138]

Responding to this dilemma, we must state that for an internalist account of epistemic justification to be philosophically interesting it will need to be non-circular. That is to say, there will need to be an "anti-Quine" restriction on using the sorts of processes that are being scrutinized to justify a reliance on those processes. How we attempt to present an anti-sceptical non-circular account of justification that is philosophically interesting will also depend, in part, on our normative epistemic commitments. On the foundationalist analysis of justification, a belief is inferentially justified just in case it receives its justification from other inferentially justified beliefs. The story is familiar: a linear conception of epistemic justification necessitates at least one belief that is non-inferentially justified in order to avert a justificatory regress. So the requirement for foundationalism is to show that we can have at least one non-inferentially justified belief that depends on a philosophically interesting conception of epistemic justification – that is, a conception that does not invoke natural properties in our explanation of how we can be non-inferentially justified.

Fumerton attempts to meet this challenge by positing a direct acquaintance account of non-inferential justification.[139] According this view, "one has a non-inferentially justified belief that *p* when one has the thought that *p*, and one is acquainted with the fact that *p*, the thought that *p*, *and* the relation of correspondence holding between the thought that *p* and the fact that *p*."[140] Fumerton is explicit that being acquainted with truth requires that one is not just aware of some fact that *p*, but is also aware of the thought that *p* and the correspondence between of this thought with the fact that *p* – merely being acquainted with facts is not enough to know that *p*; since in order to know that *p* (where knowledge depends on truth) one must also be acquainted with the truth-maker (the fact that makes *p* true) and some thought about that fact (where this thought functions as the truth-bearer).[141] If this account were plausible, then being acquainted with the fact that *p* is part of what justifies one's belief that *p*. This has the

[138] I say this because it seems that naturalism is doing away with the concept of justification altogether – See: Kim, "What is Naturalized Epistemology", pp. 542-544
[139] Fumerton, *Metaepistemology and Skepticism, p. 75*
[140] Ibid.,
[141] Ibid., p.78

conclusion that having a (non-inferentially) justified belief that *p* entails that *p* is true, since the fact that *p* is partially constitutive of one's justification for *p*.[142]

The problem with this suggestion is that we often think that our beliefs about the world are infallible even when they are manifestly false. History shows that we often get it wrong. Fumerton himself admits that the fact that a certain justification is infallible does not entail that one could not have a mistaken belief that one has an infallibly non-inferentially justified belief.[143] At the very least we cannot have any assurance that another person's justification is infallible – because only *they* can be directly acquainted with *their* thought that *p,* and the correspondence that holds between the thought that *p* and the fact that *p*. However, this implication is problematic because it also extends to our own self-ascriptions of knowledge: given that we often misapprehend our own justification for *p*, how can we have any assurance that our belief that *p* is infallible? In trying to meet this concern, Fumerton retorts,

"If I am asked what reason I have for thinking that there is such a relation as acquaintance, I will, of course, give the unhelpful answer that I am acquainted with such a relation. The answer is question begging if it is designed to convince someone that there is such a relation, but if the view is true it would be unreasonable to expect its proponents to give any other answer."[144]

In the absence of being able to offer a reason to others that there is such a reason, then it is not clear what grounds we would have to think that this relation is possible, or that we stand in such a relation. If no one can tell me what it is like to be directly acquainted with some fact that *p* in a non-question begging way, or more specifically, what features of their doxastic situation make it the case that they stand in such a relation, then how do *I* know whether *I* stand in that relation? And in which case, how do I know which beliefs of mine are cases of being directly acquainted with truth? If we cannot answer these concerns, then the claim that one has infallible non-inferential justification appears an article of faith.

The potential internalist response to this challenge may be to say that in virtue of being *directly* acquainted with *p* then I should not need to offer any explanation of my relation with *p*; this is the strategy that Fumerton uses.[145] But recall the objection to naturalism: that it cannot offer a philosophically interesting analysis of epistemic properties because it analyses

[142] Ibid., p.77
[143] Ibid.,
[144] Ibid.,
[145] Fumerton, *Metaepistemology and Skepticism*, p. 76

properties that may be *opaque* from our perspective, and so naturalism deprives us of any ability to discriminate, via reflection, our true beliefs from cases of knowledge. Insofar as this criticism therefore constitutes an *argument* against externalism, the onus is on internalism to offer *internalist* justification for the view that we have non-inferentially justified beliefs, and to offer an account for how we might be able to identify particular instances of direct acquaintance; if the internalist is unable or unwilling to do so, there is no reason why they should lament externalists who engage in exactly the same strategy. Consequently, something akin to the meta-epistemological challenge is affecting Fumerton's account as well – since he concedes that there are no grounds that he can adduce to show that he *is* directly acquainted with the facts.

The consequence of presenting the philosophical objection to externalism is possible scepticism about the possibility of non-inferential justification, or otherwise accepting the same terms as externalism (I use the word "possible" since I haven't the space to explore every internalist construal of non-inferential justification). Nonetheless, it is apparent that the *ease* with which we can have knowledge is compromised once we demand that the question of whether we know *p* depends on having internalist non-inferential justification for *p*. When the question of what we know is made to depend on questions concerning the relationship between justification and our reflective awareness or access to that justification, then whilst the philosopher is in some sense "privileged", he is also burdened, for he must account for that relationship before he can proceed with confidence that we *have* knowledge – and accounting for that relationship is not easy.

The contrary appeal of externalism is that in making knowledge a function of how we stand in relation to natural properties that exist independently of our reflective perspective on those properties, the class of non-inferentially justified beliefs is in principle vast, and so externalism accommodates the possibility of having a wealth of knowledge with relative ease. It allows that a certain amount of "data" about what we know is attainable prior and independently of having to engage in any analysis of that data, for on the externalist picture, complex connections between facts and properties may be obtaining all the time to give us knowledge. This picture might not privilege philosophers and their meta-epistemic worries, but it would appear to privilege our anti-sceptical intuitions about a variety of cases in which it appears we have propositional knowledge.

Chapter 4: Defending Common-Sense

The tension between meta-epistemological scepticism (MS) and first-order scepticism is thus a tension between two competing intuitions. The first intuition, about eliminating MS, is a desire to offer a philosophical explanation of our knowledge, that is, an explanation that can allow us to identify instances of knowledge. I have said that this requires us to reject naturalism, on which externalism depends. However, in doing so, we have to give up an attractive aspect of externalism: namely, the relative ease with which externalism in principle gives us non-inferential justification. The second intuition is therefore anti-sceptical and supports externalism; this intuition attempts to preserve putatively true beliefs (henceforth PTBs) as cases of knowledge possession. What weight, if any, should we afford these intuitions respectively?

Commonsensism and Epistemic Closure

The desire to eliminate MS would only be satisfied by achieving epistemic priority, which is not possible. Consequently, the requirement that knowledge excludes MS would be a requirement that means we lack knowledge. However, scepticism about first-order knowledge conflicts with our deliberative commitments about a variety of propositions *as* instances of knowledge possession. The consequence of making our analysis of knowledge depend on eliminating MS is therefore to deny us knowledge of PTBs. The scope of PTBs is vague, and beliefs whose content gives rise to disagreement – such as moral, religious, or aesthetic, would be difficult to classify as PTBs without begging-the-question. However, the empirical proposition "I have hands" is universally assented to (amongst people with hands) and therefore qualifies as a PTB. Of course, a belief in such a proposition may not be a case of *knowledge*, but that is an entirely separate matter to the question of what sorts of propositions we are ordinarily committed to in everyday life.

One way of resisting scepticism about (first-order) knowledge, is to insist that any adequate theory of knowledge must explain our PTBs. On this view, the sorts of commonsense propositions that we are deliberatively committed to as true may be used as a "philosophical proofs" in order to refute scepticism. The view that philosophy should take PTBs as philosophical data is called *commonsensism*; and the view that an adequate account of knowledge must accommodate, or explain, our commitment to putatively true beliefs as cases

of knowledge is correspondingly termed *epistemological commonsensism.*[146] According to epistemological commonsensism, empirical beliefs such as "I have hands" should be taken as primitive evidence or data on which to construct an epistemology, and so any analysis of knowledge that fails to accommodate PTBs therefore ignores the relevant data about what we actually know. Thomas Reid adopts commonsensism, in its most ardent form, to philosophical principles generally.[147]

…The belief of a material world is older, and of more authority, than any principles of philosophy. It declines the tribunal of reason, and laughs at all the artillery of the logician. It retains its sovereign authority in spite of all the edicts of philosophy, and reason itself must stoop to its orders.[148]

The implication of the above passage is that we know or are at least justified in believing that there is an external world, which is an empirical belief, and that philosophy must take this observation as data.[149] The above view would leave little space for us to indulge in the philosophical curiosity of asking whether there *is* a material world, least of all to let our epistemic commitments be governed by meta-epistemic concerns about the status of PTBs. A philosophical analysis of knowledge must, if commonsensism is true, be founded on the sorts of things that we already know prior to our philosophical investigations. This has profound implications for the debate between internalism and externalism in epistemology: the internalist requirement that first-order knowledge depends on further knowing, via reflection, that one has first-order knowledge, would be a requirement that cannot be satisfied, and so the class of non-inferentially justified beliefs in the actual world would be no greater than zero – the result of which is scepticism about first-order knowledge. The elimination of MS results in an analysis of knowledge that therefore fails to explain the status of our PTB's. On Reid's commonsensism, we could reject the plausibility of any internalist requirement that deprives us of knowledge.

However, a problem with the above approach is that eschewing meta-epistemological scepticism in its entirety is subject to the charge of dogmatism. Accordingly, if such commonsensism is correct, then the role of epistemology is to tell us *how* we know what we

[146] Lemos, Noah, "Commonsensism in Ethics and Epistemology", in *Knowledge, Truth, and Duty: Essays on Epistemic Justification,* Edited by Matthias Steup (Oxford University Press, 2001), p. 204
[147] Ibid.,
[148] Thomas Reid, *Thomas Reid's Inquiry and Essays*, edited by Ronald E. Beanblossom and Keith Lehrer (Indianapolis, Ind.: Hackett, 1983), p. 54. Lemos uses this example. See: "Commonsennism in Ethics and Epistemology", p. 207
[149] Lemos, p. 207

do already know – and not to offer any critique of whether we actually *do* possess such knowledge. The epistemologist's role would be to locate the relevant nomological processes involved in the formation of reliably true beliefs, searching for the relevant inputs and outputs that turn true beliefs into knowledge.[150] But, as Fumerton nicely remarks, the neurophysiologist is much better suited to this task than the philosopher.[151] This may *indeed* be the role of epistemology, but Reid's claims are a statement of this thesis rather than an argument for its first premise; and so they supply no reason for the philosopher to accept the view that PTBs are data.

The brute assertion that epistemological commonsensism is true, and that we should rely on putatively true beliefs as philosophical data for our epistemology, lacks much persuasive force and, alone, will not demonstrate that the correct analysis of knowledge is one that excludes first-order scepticism. However, G.E. Moore's more sophisticated defence of epistemological commonsensism identifies that denying that PTBs are datum leads to a type of philosophical inconsistency, which he thinks philosophers ought to take seriously. Like Reid, Moore proclaims a "Common Sense view of the world" as wholly true: [152]

"We are all, I think, in this strange position that we do *know* many things, with regard to which we *know* further that we must have evidence for them, and yet we do not know *how* we know them, i.e., we do not know what the evidence was. If there is any "we", and if we know that there is, this must be so: for that there is a "we", is one of the things in question. And that I do know that there is a "we", that is to say, that many other human beings, with human bodies, have lived upon the earth, it seems to me that I do know, for certain" (Moore, A Defence of Common Sense, p118)

Moore is claiming there is something contradictory about denying the commonsense view. [153] I think there is something right about Moore's analysis, and that at the heart of the appeal to epistemological commonsensism is a desire to achieve an analysis of knowledge that is consistent with the deliberative commitments we inescapably make – and that there is something at least paradoxical, if not contradictory, as Moore says, about philosophical positions that are inconsistent with the sorts of commitments we invariably make. Moore's epistemological commonsensism, if successful, would support a form of methodological

[150] Fumerton, *Metaepistemology and Skepticism*, p. 171
[151] Fumerton, *Skepticism and Naturalistic Epistemology*, p.327
[152] Moore, G.E. "A Defence of Common Sense", *G.E. Moore: Selected Writings*, (Routledge, 1993). P.118
[153] Ibid.,

conservatism according to which there is a *presumption* that the pronouncements that epistemology makes ought to be consistent with the propositions it seems we know; this presumption would act as a constraint on any analysis of knowledge whose conclusions are radically revisionary.

The problem for Moore, much like Reid, is that such commonsensism (even if relatively non-dogmatic in implication) is notoriously difficult to argue for in a way that is non-question-begging and non-dogmatic. This is because assertions about what we *do* believe function as psychological statements.[154] Even if some belief is universally assented to, that does not provide us with proof that the proposition is a case of knowledge. For as long as truth is a real property in the world, and not a function of some contractual standard, then demonstrating that we possess everyday knowledge by citing psychological observations about particular propositions that are widely, or indeed universally, assented to does not demonstrate with any conclusiveness that these propositions are true.

To see this problem: any first-person claim that one knows the PTB that "I have hands" would be justified by the fact that it seems I have evidence for the proposition – let us say that my evidence consists of the fact that I perceive lucidly that I have hands, and that I am reflectively aware of this fact. However, sceptical hypotheses such as BIV, or Descartes' evil demon hypothesis, derive their argumentative force from the very fact that the same phenomenal experiences could still obtain even if one was being systematically deceived about their truth, that is, even if the PTB was false. To therefore show that one has knowledge of the relevant PTB, one would need to demonstrate that one is not a brain-in-a-vat, or that one does not otherwise inhabit some sceptical scenario. Let us call knowledge that sceptical scenarios (in which B$p \wedge \neg p$) do *not* obtain, "knowledge that q" (Kq). If we need to know that q in order to know that p, then knowledge would necessarily be closed under entailment, such that the following deductive closure principle would apply to propositional knowledge.

Kp

K($p \rightarrow q$)

$\therefore$ Kq

[154]Hookway, Christopher, "Critical Common-Sensism and Rational Self-Control", *Nous*, (1990), Vol. 24, p. 398

According to the above principle, if one does not know that *q*, then one cannot know that *p*. So if closure obtains, we must be able to know the negation of sceptical hypotheses whose falsity is logically entailed by propositional knowledge. The Moorean version of epistemological commonsensism states that we can know the denial of sceptical hypotheses in virtue of already possessing knowledge of ordinary propositions, which logically entail *q*. Moore therefore claims that we already possess knowledge of certain mundane propositions, and so on the basis of possessing such knowledge, we can know the negation of sceptical hypotheses.

However, the problem with this template is that any *claim* that *p* is an instance of knowledge possession will rest, for its argumentative force, on having internalist justification or knowledge for *p* as an instance of such a proposition; yet any such grounds are bound to be question-begging, because the reflective grounds that warrant *p* will continue to warrant *p* in a sceptical *q* scenario in which not-*p* obtains. The *p* and not-*p* hypotheses are psychologically indiscernible, and so the kind of warrant we have for thinking that *p* is true does therefore not transmit to provide us with warrant for the claim that we know *p*. [155]

The standard Moorean version of epistemological commonsensism is therefore question-begging. It states that we can internally know, via reflection, that *q* (where *q* is our knowledge of the negation of sceptical hypotheses in which ¬*p* obtains) on the grounds that we *already* know, via reflection, that *p*. But, as sceptical scenarios show, it does not seem that the reflective grounds we possess for *p* can transmit to provide us knowledge or justification that *q* obtains.[156] So, to preserve knowledge that *p* we either need to abandon our unwavering commitment to the deductive closure principle, thus allowing the possibility that we can know via reflection that *p* even though we cannot know via reflection that *q*, or otherwise stipulate that there is some method other than reflection that would allow us to know the negation of sceptical hypotheses without begging-the-question.

Externalist epistemologies allow that whether we can know some proposition *p* is a function of how our belief that *p* is counterfactually related to the truth of *p*.[157] On a safety analysis, the template for preserving epistemological commonsensism is already in place. Providing

[155] My analysis here is close to the account of transmission offered by Crispin Wright. See
Wright, Crispin, "Anti-Sceptics Simple and Subtle: G.E. Moore and John McDowell", *Philosophy and Phenomenological Research*, (2002), Vol. 65, No. 2, pp. 331.
[156] Ibid.,
[157] The counterfactuals needn't be modal – for example, the same concerns will apply to Goldman's reliabilism (1967, 1976).

that the world is largely as it appears to be, then it will follow that our true belief that *p* continues to be true in most or all close possible worlds in which one believes that *p*.

This re-iteration of the externalist's counterfactual requirement, however, is at risk of speaking past the sceptic: after all, the crux of scepticism is that the world may *not* be as it appears, and that we are therefore not justified in presuming that it is as it appears. Whether the world *is* largely as it appears, still remains unanswerable. We can therefore not know that the template that would give us knowledge that *q* is in place. This is ostensibly an even greater problem for sensitivity theorists, because the fact we cannot know that *q* can be shown to demonstrate that our belief that *p* does not subjunctively depend on the truth of *p* (in which case we cannot know that *p*).

Externalism Vindicated

There are two ways in which sensitivity theorists have attempted to overcome this problem. One response to this problem is to preserve closure whilst placing sensitivity within a contextualist framework. For example, a contextualist analysis of knowledge (DeRose, 1995) makes knowledge possession a function of propriety conditions under which a given knowledge utterance becomes salient. [158] The second strategy is to deny the validity of closure, and to therefore say that one can know that *p*, where the logical entailment of *p* is *q*, and moreover where I *know* that *p* entails *q*, whilst still failing to know that *q* (Nozick, 1981, Dretske, 2005). To summarize: safety averts scepticism by quantifying across only *nearby* worlds. Sensitivity averts scepticism by accepting contextualism or saying that we needn't know the known logical entailments of *p* in order to know that *p*.[159]

My problem with denying closure is that closure seems a reasonable assumption to make about knowledge. This is because the assumption of closure allows us to avoid the paradoxical conjunction that you know you have hands *and* you do not know you are a brain in a vat in which you do not have hands.[160] The kind of knowledge one preserves by rejecting

[158] The thought is, we can have everyday knowledge, but only in everyday *contexts*, because closure is satisfied because everyday knowledge utterances will be salient in everyday contexts. (De, Rose, 1995)[158] On the contextualist view, we cannot claim to know that *p* in the sceptical dialectic because the propriety conditions for knowledge possession are "raised" in this context. So, whilst knowledge is closed under entailment, in the everyday context of possessing everyday knowledge the propriety of our knowledge claims does not depend on adducing reflective evidence to show that sceptical hypotheses are negated. Contextualism would therefore allow us to maintain everyday knowledge, although only in "everyday contexts": thus, the knowledge that one can saliently claim in everyday contexts is immune from concerns about sceptical hypotheses obtaining.

[159] DeRose, Keith (1995) opts for contextualism. Nozick (1981) and Drestske (2005) deny closure. An argument for keeping closure can be found in Hawthorne (2005) – see bibliography

[160]DeRose, "Solving the Skeptical Problem", *The Philosophical Review,* (1995). Vol. 104, No. 1,, *p 28*

closure is of a fairly impotent variety that cannot be adduced in any conversational context without having to countenance what has been termed the "abominable" conjunction that you know that *p*, but *p* may be false.[161] Denying closure may have the desired result, but it requires us to abandon a commonplace and entrenched intuition about knowledge: if I have knowledge (of an everyday type) that I have hands, then I know that I am not a brain in a vat. Denying closure as a way of preserving everyday knowledge therefore achieves the desired result but comes with an explanatory burden: it would need to accommodate our widely held intuitions about sceptical dialectics undermining our knowledge claims. The only impetus for denying closure is therefore to say that whilst closure is an intuitively appealing principle, it commits us to scepticism, and so because scepticism is even *more* objectionable than the rejection of closure, closure must be false. On the basis of this predicament, Dretske thinks that the choice is easy: we must deny closure.[162] He presents the rejection of closure as a painful but ultimately necessary choice.

There are of course independent reasons for or against the contextualist view that I will not discuss here; similarly, there are independent reasons for or against accepting safety as opposed to sensitivity. But we should not feel *forced* to abandon invariantist forms of externalism, or to embrace a safety principle instead of a sensitivity principle, in order to accommodate our intuitions about knowledge being closed under entailment. This is because neither safety nor sensitivity require the denial of closure. Rejecting closure because it is the *least bad option* is too concessive; we can keep closure *and* invariantism, whilst preserving the essentially externalist framework for knowledge. This possibility can be realized when we stipulate that the method by which we come to know some proposition *p* is individuated externally.

Let us recall that closure threatens us with scepticism on the basis that the evidence we have for *p* does not get us knowledge of the known entailment of *p* (*q*). For whenever we know *p*, there must be something to indicate that *p* is true.[163] But the very thing that indicates *p* – or gives us evidence for *p* – cannot indicate, or give evidence for, *q*. Global scepticism looms

[161] DeRose, "Solving the Skeptical Problem", p28
[162]Dretske, Fred, "Is Knowledge Closed Under Entailment? The Case against Closure" in *Contemporary Debates in Epistemology*, Edited by Mathhias Steup and Ernest Sosa (Blackwell, 2005), p. 23
[163] Dreteske, "Is Knowledge Closed Under Entailment? The Case against Closure" p.22

because this failure of closure applies to all propositions about the material world. [164] As Dretske remarks,

"…there is nothing in the world – either mental or material – that indicates that there is a material world. Nothing in the present that indicates that there is past….that is why fuel gauges indicate, and therefore tell you, that you have gas in your tank without telling you that there is a material world (something implied by there being gas and gas tanks. That is why rings in the stump can indicate the age of a tree without indicating that the past is real"[165]

Dretske is surely right in this regard: nothing within our reflective vicinity could allow us to be assured that there exists an external world – hence the sort of evidence we often have for fairly banal empirical propositions, such as the gas being in the tank, will not transmit to give us evidence in favour of the existence of an external world. The problem for advocates of closure is that any two subjects can have the same evidence in favour of the proposition *q* (that there exists an external world), but it can be the case that one of them doesn't know that *q* on the basis of having a false belief that *q*, and so necessarily the other doesn't know that *q* even if he inhabits a *q* world. When the method by which we get evidence in favour of *q* is individuated in terms of what an agent is experientially or reflectively aware of, it seems that closure cannot be satisfied; hence it is no surprise, as Dretske rightly points out, that looking at one's fuel gauge will not tell one that there exists an external world.

This same problem arises for sensitivity theorists if the method (M) by which we come to know some proposition *p* is individuated internally. Nonetheless, sensitivity theorists still often individuate methods internally. For example, Robert Nozick says of methods:

Usually, a method will have a final upshot in experience on which the belief is based, such as visual experience and then a) no method without this upshot is the same method, and b) any method experientially the same, the same "from this inside", will count as the same method (Nozick, Philosophical Explanations, *Knowledge and Skepticism, p184*) [166]

If we understand methods in this way, then we cannot know that *q* (*q* = we are not being systematically deceived about our belief that there exists an external world) on the basis that the method M we would use to arrive at a true belief that *q* would be experientially the same, or the same "from the inside", if you will, as the method we would use in the sceptical world

[164] Ibid.,
[165] Ibid.,
[166] E.g. Nozick, Robert, *Philosophical Explanations*, (Oxford University Press, 1981), p184

in which we have a false belief that *q*.[167] Indeed, on Nozick's analysis, M would be the very same method that we would be employing to have a false belief that *q* in a ¬*q* world.[168]

However, we can overcome this problem by individuating method externally. Once we allow that the method M by which we come to know proposition *p* is a function of nomological properties and processes that are not necessarily reflectively accessible, we can have a sensitive or safe belief that *p* that is a case of knowledge that *p*, and in virtue of having such knowledge, we can also have knowledge that *q*. The external individuation of methods will allow us to say that any two subjects X and Y that are alike in internal/reflective and phenomenal respects, but where X has a true belief that *q* and Y has a false belief that *q*, may nonetheless inhabit evidentially dissimilar doxastic situations on the basis of employing different methods to arrive at the belief that *q*. X may know *q* even though there is a counterpart version of X (namely, Y) who is fed the same phenomenal experiences as X yet has a false belief that *q*. Y's doxastic situation evidentially differs to X's, which explains why Y fails to know *q*.

For example, let us suppose that S believes that *p* via M: then in order for S's belief to be a case of knowledge, it must be the case that: for whenever *p* obtains, and S considers (via M) whether or not to believe that *p*, then S comes to believe that *p* via M. Counterfactually, it must also be the case that if *p* did not to obtain and S were to consider, via M, whether or not to believe that *p*, S would not believe that *p* via M.[169] Once the method M is individuated externally, then the sceptical world in which one has a false belief that *p* would be a world in which one does *not* believe that *p* via M – accordingly, the method we would be using to arrive at our beliefs in the sceptical world, when individuated externally, would be different to the method used to arrive at true beliefs in the actual world. Internally, the method M* used in the sceptical world would be phenomenally indiscernible from M, but from an external perspective, M and M* are two quite different ways of arriving at beliefs about the world: one method depends on the existence of a physical world in order to furnish us with true beliefs about that world, the other has no dependence on the physical world. In this respect, sensitivity allows that we can have everyday knowledge that resists sceptical hypotheses – providing we do not succumb to the internalist temptation to individuate

[167] E.g. Ibid
[168] Ibid.,
[169] My use of counterfactuals is similar to Nozick's. See: Nozick, *Philosophical Explanations,* , p. 174-177

methods in terms of what an agent is experientially or phenomenally aware of.[170] But there is no reason why externalism should be tempted to individuate methods *internally* for our knowledge of closure (i.e. our knowledge that we do not inhabit a sceptical world) and *externally* for our first-order knowledge. In trying to reject closure, Nozick therefore concedes too much and makes an unnecessary move.[171] His subjunctive conditionals do all the work he needs them to do.

An interesting implication of preserving closure within an invariantist framework and keeping sensitivity is that, insofar as we accept the anti-sceptical intuition that external methods can give us knowledge, sensitivity and *some* versions of safety give us identical predictions for the purpose of preserving the externalist template.[172] The particular version of safety that gives us the same result as sensitivity will be one that quantifies across *all* and *only* nearby possible worlds. On this version, my belief that I will win the lottery is neither safe nor sensitive. But likewise, my belief that I am not a BIV is both safe *and* sensitive.

However, safety and sensitivity only give us similar predictions when we do *not* share the intuition that we need a more relaxed version of safety *and* we individuate methods externally for both safety and sensitivity. For example, a version of safety that incorporates closeness as a gradable concept in order to capture the intuition that there are very close worlds that should be fully quantified across, whilst preserving a "ring" of worlds that are only partially quantified across, will give us different results to sensitivity. On this relaxed conception of safety, I will be able to have everyday beliefs that are safe but not sensitive, and so safety and sensitivity will yield different results. However: so long as we are happy with safety as a condition that quantifies across *all* and *only* nearby worlds, and we make no further concessions in this regard, then safety will preserve everyday knowledge in much the same way as a version of sensitivity that individuates methods *purely* externally will do so.

Whatever counterfactual requirement we opt for in our normative epistemology, the same considerations will apply: that is, as long as the method by which we come to know some proposition satisfies some relevant externalist counterfactual requirement, then we may have

[170] On Nozick's understanding of methods, if any two methods M and M* are internally alike, then M and M* are alike methods. For his account, See: Nozick, *Philosophical Explanations*, p.184

[171] Pritchard, Duncan, "Sensitivity, Safety, and Anti-luck Epistemology", in *The Oxford Handbook of Skepticism*, Edited by John Greco (Oxford University Press, 2009), p.444 (Pritchard notes that the source of Nozick's problem is that he understands methods "internalistically")

[172] Pritchard, "Sensitivity, Safety, and Anti-Luck Epistemology", p. 453

everyday knowledge. On the externalist analysis, we can therefore know the negation of sceptical hypotheses *externally*.

The Limitations of Externalist Knowledge

The apparent ease with which externalism preserves everyday knowledge should frustrate and infuriate the internalist: it seems externalism only needs to re-iterate its first premise that belief-forming-methods can be individuated externally, and so, if it follows that the world is largely as it appears to be, then the sort of everyday knowledge we seek is forthcoming. To have everyday knowledge, we only need cooperation with the world; not a reflective perspective on that cooperation.

However, the limitation of this externalist everyday knowledge is that, from the reflective perspective of the agent, it will be a matter of luck that one's belief is a case of knowledge. We have denied that it is *necessarily* the case that one does not know that some sceptical hypothesis (SH) obtains. However, denying this necessity is not the same as *showing* that it is the case that one *does* know that not-SH obtains. It may of course be the case that not-SH obtains, but the externalist has not *shown* this merely by rejecting the internalists' premise that we cannot know not-SH. Whether one *has* everyday knowledge is therefore only possible, and so, from the reflective perspective of an agent, it will be a matter of luck (reflective luck) for whenever one *does* know the negation of SH. The internalist needs to articulate what is damaging about this luck.

The most obvious sense in which this reflective luck is restrictive is because it would make it improper to claim everyday knowledge in a sceptical dialectic, because the grounds that convert out true beliefs into knowledge will not be of the reflectively variety that could be adduced in a conversational context as an argument for the premise that we *do* possess knowledge.[173] Externalism may respond to this challenge by noting that on an externalist epistemology, which is non-intellectualist, knowledge possession need not be a function of what can legitimately be claimed. Externalism may also respond that this sort of independence between knowledge possession and knowledge claims is desirable for independent reasons.

However, there is a further related challenge that should not be so easily dismissed: how can we have assurance that we have externalist everyday knowledge? After all, if it is improper to

[173] See: Pritchard, *Epistemic Luck* , p.79-86

claim this knowledge to others, then surely it is improper to self-ascribe knowledge to ourselves? This question arises because, from our reflective perspective, it is only ever possible that one has externalist everyday knowledge. The dissatisfaction with externalism as a response to scepticism is that externalism is only ever a congenial possibility that scepticism is false and nothing can be offered to put us in a position to claim that it, rather than one of the uncongenial sceptical scenarios, actually obtains (C.Wright, 2007).[174]

So we need some reason to think that we inhabit a congenial situation. However, externalism only *allows* that we have externalist everyday knowledge, which falls short of an endorsement that we *do* have externalist knowledge. In order to make this latter endorsement we would need to operate on the prior assumption that some aspects of reality *are* as they appear, from which we could infer that certain beliefs we hold about the world are liable to be true. However, as I have shown, no amount of externalist justification could give us reason to think that our beliefs are externalistically justified, or satisfy some externalist requirement. It follows from this that externalism cannot be used to motivate or endorse the view that we have externalist everyday knowledge. Subsequently, in order to be *assured* that we have externalist everyday knowledge then we need to buy into the notion that there are certain internalist grounds for thinking that the world is as it appears.

The proposal is therefore that having grounds for thinking that externalism is true (of the kind that would allow us to claim externalist everyday knowledge) would depend on having internal grounds for thinking that our beliefs meet an externalist requirement. However, the same features of one's doxastic situation that make internalist knowledge claims question-begging are liable to make our internal grounds for thinking that we have externalist knowledge similarly question-begging. That is to say, internal grounds for thinking that the world is as it appears are not indubitable, as sceptical hypotheses show, and so any anti-sceptical claim to have this knowledge will be question-begging, due to transmission failure. I have shown that Reid's and Moore's claims are paradigmatic of this. The problem with Reid's and Moore's versions of commonsensism is that their arguments function as merely psychological constructions about what we *do* believe, which fail to show that scepticism is false, due to the psychological indiscernibility of sceptical hypotheses.

[174]Wright, Crispin, "The Perils of Dogmatism", in *Themes from G.E. Moore: New Essays in Epistemology and Ethics*, Edited by Susana Nuccetelli and Gary Seay, (Oxford University Press, 2007), p. 31

Pragmatic Justification

Sceptical hypotheses appear to undermine our confidence that the externalist template for knowledge is in place. They make us question the reliability of our reflective processes as a method of arriving at true beliefs about the world. We want to be able to rely on reflection as a way of having knowledge of the world, but we can't, and so the risk is that we cannot offer an endorsement of the view that externalism is true. To offer justification for externalism, we therefore need to allow that the reflectively accessible features of our doxastic situation may justify the belief that the externalist template for knowledge is in place: if we are not justified in believing that the externalist template is in place, then what accessible reason could we have to think that externalism gives us everyday knowledge? What reason could we have to think that our beliefs about the world are safe or sensitive?

The proposal is that, in the absence of having an internally epistemically justified belief that the externalist template is in place, we may still be able to have pragmatic justification for externalism. Provisionally such a view is plausible: after all, justification is non-factive and even if the externalist template is *not* in place, that need not necessarily undermine the possibility of being justified in thinking that externalism is true – we can in principle have *prima facie* justified beliefs about false propositions.[175] To be justified in thinking that externalism is true does not, therefore, require that one *knows* that externalism is true. Justification may hinge instead on the sorts of pragmatic considerations that make one's belief reasonable and responsible.

On a pragmatic account of justification, whether some proposition *p* is justified is a function of the effect that abandoning a commitment to *p* would have. This may be formulated as a non-instrumentalist and negative requirement: pragmatism about justification does not state that we are justified in believing whatever is advantageous, but rather, claims that for whenever abandoning a deliberative commitment to *p* is impossible to do without one being irrational, or is impossible metaphysically, then we have an involuntary basis to think that a commitment to *p* is justified.

Our ordinary deliberative commitments take it as given that there exists an external world that is largely as it appears to be – that is to say, in our ordinary lives, our belief that something akin to the externalist template for knowledge is in place, is psychologically

[175] James Pryor, "The Skeptic and the Dogmatist", p. 519

irresistible insofar as we invariably *do* believe propositions of the form "I have hands". Only in philosophy departments do we think that it is possible that we are being systematically deceived by evil demons, and the ilk. In everyday life, we just don't have such thoughts (indeed, the presence of such thoughts is seen as a mark of irrationality).

However, descriptive claims about *what* we believe are not going to give us the normative conception of justification that allows us to prescribe that for whenever one is justified in believing that *p* one thereby *ought* to believe that *p*. We need to articulate the connection between justification and what is responsible to believe. In relation to this connection, Hans Reichenbach (1948) gives the following example:[176]

A blind man who has lost his way in the mountains feels a trail with his stick. He does not know where the path will lead him, or whether it may take him so close to the edge of a precipice that he will be plunged into the abyss. Yet he follows the path, groping his way step by step; for if there is any possibility of getting out of the wilderness, it is by feeling his way along the path.[177]

The blind man chooses to follow the path in this instance because he really has no choice: the choice is either certain death of the adoption of a method that might possibly allow him to survive. There is no evidence to suggest this method – "groping his way step by step"[178]- will yield the desired result, but it seems that *he ought* to give it a go, and that it would be irresponsible on the basis of any sceptical doubts about where the path may lead, to take no further action and allow the inevitable to take its course – indeed, it would seem irresponsible in this scenario to even entertain sceptical doubt. There is a clear sense in which the blind man would be rationally criticisable had he not employed the aforementioned method in order to save his life.[179]

Reichenbach's examples illustrates that sometimes one's situation makes reliance on a certain method responsible, even though there is no evidence for that method – and this responsibility is enough to make one epistemically responsible in using that method. The agent, here, seems justified in virtue of pragmatic considerations. Likewise, even though reflection cannot help us differentiate non-sceptical scenarios from sceptical scenarios, our

[176] For a discussion of this, See: Enoch, David; Schechter, Joshua. "How Are Basic Belief-Forming Methods Justified?", in *Philosophy and Phenomenological Research,* (2008), p. 554. Reichenbach's original example can be found in *The Theory of Probability,* (University of California Press), p. 482
[177] Enoch, "How Are Basic Belief Forming Methods Justified?", p. 554
[178] Enoch, p. 555
[179] Enoch, p. 554

epistemic predicament may be one in which we are justified in using reflection as a basis for thinking that there exists an external world that is largely as it appears – and that sceptical scenarios are false.

The pragmatist implication is that the nature of the project, and the degree of freedom we can exert over a project, is relevant to whether one is justified in relying on a method that is essential to that project. If it is not rational to pursue the project, then reliance on some method that is essential to the successful fulfilment of that project will not be justified. Of course, we may still be *instrumentally justified* in such cases because it is advantageous to rely on that method – but in which case, we would have some reason to think that the method will give us the results we want – where this reason explains why it is we have opted for relying on that method. However, my focus here is on how we can be justified in relying on a method M where there is nothing within our reflective grasp that could discern that M is going to yield true beliefs, but where relying on M may nonetheless be pragmatically justified because it is essential to engaging in some project that is rational.

Talk of projects and their relationship to pragmatic justification is vague, and it would be difficult to offer a comprehensive list of the sorts of projects that are rational to engage in. The pragmatic account offered by David Enoch and Joshua Schechter (2008) argues, likewise, that any answer to the question of what it means for a given project to be rational will not be terribly informative: and so he concedes that rational projects may have a "basic normative status, unexplainable in more fundamental normative terms". [180]

Nonetheless, Enoch and Schechter assert that there are certain projects that are *rationally required*, and there is a clear sense in which discarding the project is not an available course of action.[181] They give the example of our reliance on our normative intuitions as required in order to decide what it is one *ought* to do. Deciding what we ought to do is a project that is rationally required for us –since we unavoidably ask what it is we *ought* to do, what *ought* we to believe, and how *ought* we to behave. They are surely right in this regard: there would be something disturbingly a-rational about someone who says he fails to engage in this project; but it would, moreover, be an insincere claim, since it would appear *impossible* not to engage in the project of deciding what we ought to do – even if the conclusions that one reached

[180] Enoch, p.557
[181] Ibid., p. 557

were objectionable. In this respect, the deliberative project is rationally required of us.[182] On the basis of this project being rationally required, Enoch and Schechter argue:

"…it is plausible that we can only successfully engage in this deliberative project if we have some way of coming to know normative truths. Given our constitution, our only hope for coming to know such truths is if our normative intuitions are at least reasonably reliable; if relying on these intuitions is not at all reliable, the project of deliberating about what to do is doomed to systematic failure."[183]

Enoch and Schechter emphasise the relationship between a project being rational and a project being unavoidable. When there is some method that is essential to engaging in a rationally required project, then we have no choice but to rely on that method. And for as long as we cannot abandon the particular project, then talk of abandoning our reliance on the method is nonsensical because it would require us to abandon the project – but where a project is rational it cannot be abandoned, and so our reliance on the method is pragmatically justified.[184]

I am likewise not sure how we could engage in a number of rationally required projects if we were to renounce our commitment to the idea that there exists an external world, which is largely as it appears, and which we can have knowledge of through our senses; indeed, if we did not believe that something like the externalist template for knowledge was in place, then we would not be able to engage in a number of rationally required projects.

Take, for example, the project of explaining some physical phenomenon.[185] When we explain some physical phenomenon we start with an initial observation, typically a perception, and we subsequently invoke an explanatory ontological criterion, according to which we have a basis to think that some property exists if it serves an explanatory role in our understanding of that phenomenon.[186] That is to say, where we cannot explain some physical phenomenon without invoking a commitment to some physical property, then it is fundamentally rational to think that that property exists – we would have, therefore, a basis to be ontologically committed to the existence of that property, and subsequently, a basis to think that the

[182] Enoch, p.556
[183] Enoch, p.557
[184] Ibid.,
[185] See Harman, Gilbert, "Ethics and Observation", in *Foundations of Ethics: An Anthology"* (Blackwell, 2007), p334-335. Originally published in "Ethics And Observation, *The Nature of Morality: An Introduction To Ethics*, (OUP, 1977) pp. 3-10
[186] Ibid.,

explanatory criterion for ontological commitment is indispensible. The project of explaining physical phenomena through science, and invoking a commitment to properties on the basis that they serve a role in our explanations, seems rational; even if it does not always yield correct answers. (Of course, the explanatory project is not the *only* rational project and the *only* basis to be committed to the existence of some property.)[187]

How could we possibly engage in the explanatory project – or other indispensible projects - without at least being justified in relying on reflection as a method of acquiring beliefs about the external world? That is, without thinking that our reflective perspective on the world is largely correct, then we could not proceed with any confidence beyond our initial observations about physical phenomena. Specifically: we need some basis to think that our initial observations, grounded in our experiential and reflective awareness of the world, are justified and likely to yield true beliefs; it seems we have to be committed to their justification in order to engage in the subsequent project of explaining them. Whilst we may question the reliability of our reflective processes in our putative philosophical thought experiments, there is no plausible sense in which we can *abandon* a commitment to the reliability of such processes, because it would require us to abandon a rational project – the project of explaining the world that we are reflectively aware of and confronted with. We are at least pragmatically justified, therefore, in relying on reflection as method for thinking that we know certain putatively true propositions. This pragmatic justification extends to the proposition that we are not being systematically deceived about the existence of an external world that is largely as it appears.

Enoch's and Schechter's account goes further. They think that we are *epistemically* justified in relying on certain belief-forming methods that are not justified inferentially; this is argued for by emphasising the connection between epistemic justification and epistemic responsibility, and it is taken to be a fundamental advantage of a pragmatic account that it can connect epistemic justification to epistemic responsibility.[188] My suggestion is that pragmatic justification differs from a pragmatic account of epistemic justification, and whether we are committed to the latter depends on the conception of "justification" we are employing. Insofar as epistemic justification is a *personal* and fundamentally normative notion, connected to what is epistemically responsible or rational to believe, then pragmatism may

[187] See Enoch, *Taking Morality Seriously: A Defence of Robust Realism*, (Oxford University Press, 2011) pp. 71-84
[188] Enoch. "How Are Basic Belief-Forming Methods Justified?, p. 569

give us the framework for the kind of conception of epistemic justification that Enoch and Schechter support – because it connects the concept of justification to what is responsible to believe [189] So where the domain of evaluation is the agent, as is the case with personal rather than propositional accounts of epistemic justification, then pragmatic constraints may ground epistemic justification.[190]

Whether pragmatic justification can give us epistemic justification is beyond the scope of this thesis and somewhat irrelevant in relation to our philosophical "felt need" to have assurance that the externalist template for knowledge is in place. This is because any pragmatist conception of epistemic justification for relying on a given method will necessarily be an externalist conception,[191] since there is nothing within our reflective vicinity that could allow us to know that the method *is* reliable; hence it seems we are dependent on pragmatic considerations to ground the rationality of our reliance on a given method of arriving at our ordinary beliefs about the world – in our case, the method of reflection.

Pragmatism about justification can show that it is rational to believe that we are not being systematically deceived about our everyday beliefs; but pragmatism will not eliminate reflective luck and its consequent, meta-epistemological scepticism. This is because it is not indubitable that there is an external world that is largely as it appears. In modal terms, there will be close possible worlds, ordered in terms of what we are reflectively aware of, in which the method of reflection delivers false beliefs.

The internalist about epistemic justification has the aforementioned "felt need" for certainty that the externalist template is in place – it must be beyond *all* doubt. The type of indubitableness he seeks displays the psychological incorrigibility that pragmatists highlight, in addition to something altogether stronger and unattainable for our knowledge of contingent propositions: namely, the very act of doubting would affirm the proposition's truth - considering *p* will be an act that will be able to perfectly discriminate *p* as true.[192] The existence of an external world and a pragmatic reason to believe that there exists an external world will, therefore, not give us a fully satisfactory internalist account of whether we *know* that there exists an external world – and will therefore not give us conclusive grounds to think that the externalist template is in place. What the internalist wants, which is altogether

[189] Ibid.,
[190] For an account of the difference between personal and propositional justification see: Engel, Mylan. (1992), pp. 133-150
[191] Enoch, "How Are Basic Belief-Forming Methods Justified?", p.569
[192] This perfection seems like an entailment.

unattainable, and perhaps what we *all* want in thinking about knowledge in a philosophical way, as Stroud remarks, "is to know or have good reason for thinking that what we believe about the world is true".[193] Our sceptical doubts about whether the externalist template is in place do nothing to affirm the truth of our pragmatically justified belief that the template *is* in place.

[193] Stroud, Barry, *Understanding Human Knowledge*, (Oxford University Press), p. 146

Chapter 5: Reconciling Externalism and Internalism: an anti-absolutist framework

Pragmatist justification will not provide internalism with the philosophical assurance internalists so crave, even if it shows that believing that we have externalist knowledge is rational. Insofar as internalist philosophical assurance is not available, the externalist epistemology that preserves commonsense everyday knowledge means that we cannot have internalist knowledge or justification for the view that we have externalist knowledge. Externalism can respond to this predicament by maintaining that so long as a given belief is externally related to truth in the appropriate way, then it is a case of knowledge simpliciter. Externalism can thus defend itself against all of the aforementioned meta-epistemic concerns just by reiterating its first premise, which entails that knowledge does not depend on having internalist knowledge or justification for one's knowledge; externalism allows that knowledge is completely opaque and completely uncompromised by being opaque. In what sense, therefore, might internalist meta-epistemic concerns still affect the prospects of externalism about justification?

In one sense, our internalist meta-epistemic concerns about the status of externally justified beliefs puts pressure on the externalist construal of "justification"', and forces us to realize that the internalist is more interested in *justification*, whereas the primary focus of externalism is constructing a template for how it is we have *knowledge.* The externalist can take our firm intuitions that we *do* have knowledge and show how such knowledge is possible; and, as I have suggested, externalism may be buttressed by pragmatic considerations that make it rational to think that the template for externalism is in place and that sceptical hypotheses are false. However, it is not obvious how externalism preserves epistemic justification as a condition on which this knowledge depends, other than using the term "justification" to denote *some* condition that converts true belief into knowledge. No doubt there is such a condition, or a set of conditions, that perform this essential task, but it is not clear in what sense we should invoke the concept of "justification" to describe whatever condition or set of conditions perform this function – accordingly, externalism allows that the factors that turn our true-beliefs into knowledge will consists in natural properties or nomological relations that may be completely opaque.

Implicit in this criticism is the idea that in naturalizing our epistemic commitments, the externalist could be said to be abandoning justification as a normative notion – and in doing

so, abandoning the concept of justification altogether.[194] For example, Quine argues that epistemology is "contained in natural science, as a chapter of psychology;[195] but there is no plausible sense in which psychology could tell us what one *ought* to believe – psychology can only tell us the conditions surrounding what we *do* believe, at best. However, if justification is a normative concept, which tells us what is permissible or reasonable to believe, that is, what we *ought* to believe, from a first-person point of view, then naturalism (on which externalism depends) cannot accommodate the normativity of justification because it is reductionist.[196] Some philosophers argue that knowledge is also a normative notion and that what makes knowledge normative is *justification*.[197] The other features of knowledge, belief and truth, are psychological and semantic/metaphysical concepts, respectively.[198] If epistemology abandons justification (the thought goes), then epistemology must abandon knowledge.[199]

This criticism of externalism is partially true: in naturalizing our epistemic commitments (and in therefore accepting a reduction of justification to natural properties, rather than positing a supervenience relation) the externalist repudiates normativity which in turn requires him to abandon the notion of justification.[200]Justification is, in this respect, an essentially internalist notion. But the extension that an epistemology that is free from normativity could have nothing to do with the concerns of epistemology or what is reasonable and responsible to believe, does not follow. This is because externalism can naturalize knowledge without naturalizing justification – he can instead simply ignore the concept of justification altogether. Granted he will subsequently not be able to meet the concerns of traditional epistemology, but he will be able to offer an account of knowledge that accommodates our intuitions about putatively true beliefs – and so he will, he thinks, be presenting a new paradigm for doing epistemology. If the criterion for an adequate epistemology is that it accommodates Moorean-type intuitions about what we know, then an epistemology "purged of normativity" is the means of achieving that.[201] "Justification", it seems, is an essentially

[194] Kim, Jaegwon, "What is Naturalized Epistemology?" in *Epistemology: An Anthology*, (Blackwell, 2008), p. 542
[195] Quine, W.V. "Epistemology Naturalized" in *Epistemology: An Anthology*, (Blackwell, 2008), p. 534
[196] I suggested earlier that we need anti-reductionist supervenience to get normativity, of the kind moral realists use
[197] Jaegwon Kim, p.542
[198] Kim, p.539
[199] Kim, p.542
[200] Kim,
[201] Kim, p.543

internalist ideal that captures our philosophical intuitions but fails to capture our pre-theoretical intuitions about what it is we actually know.

Imperfect Knowledge

However, defending externalism by way of abandoning the internalist commitment to epistemic justification, and focusing instead on *knowledge* (and on presenting an account of knowledge whose dictates satisfy our intuitions about what we know) still fails to offer a fully satisfactory way of responding to the internalist meta-epistemic challenge. That is to say, it seems important for the persuasiveness of externalism as a philosophical solution to scepticism that we are able to offer a conception of knowledge that is able to acknowledge the legitimacy of the meta-epistemic challenge to externalism, which has to do with whether we can have *epistemic justification* for the view that the externalist template is in place; re-iterating the externalists' first premise should do little to persuade *philosophers* to abandon their internalist concerns about whether the externalist template for knowledge is actually in place. Let me be clear: we should want to maintain externalism in order to keep everyday knowledge (in particular, we should want the variety of externalism that I defended in chapter 4, which keeps closure and individuates methods externally), but we should also want to acknowledge the legitimacy of meta-sceptical concerns about whether the externalist template is in place; otherwise the resulting epistemology will be something of a stand-off between two competing intuitions.

Some philosophers have argued that the predominant focus on justification is anti-externalist, because it leads to a meta-epistemic requirement.[202] Accordingly, once justification is made purely a function of what we have reflective grounds to believe, then questions about whether it is responsible to believe that we have knowledge will inevitably arise: the externalist necessarily cannot fully answer these questions in a satisfactory way, since the externalist focus is our relationship with the nomological processes in the world that give us knowledge, and not our internal reasons for thinking that we *do* have such knowledge. If we focus on justification, then externalism cannot meet the meta-epistemological challenge; but insofar as our general motivation for epistemological theorizing is to offer an account of knowledge that accords with our intuitions about what we know, then externalism can adequately do this – and so it is not obvious why, unless we are internalists, we should have to offer an account of justification.

[202] See: Stroud, *Understanding Human Knowledge,* p.145

Stroud remarks that, as theorists, philosophers should aspire to understand human knowledge in a much more general way than by merely iterating that *if* we are suitably related to the world *then* we can have knowledge.[203] A satisfactory philosophical conception of knowledge should ask *how* we know the sorts of things we think we *do* know about the world. To say that we know such things because they are true may be correct – and I have argued that it is, and that an appeal to epistemological commonsensism and pragmatism can vindicate this position - but, as Stroud remarks, the implication of re-iterating the externalist position in this way is that it implies there is nothing initially problematic about our knowledge. I think he is correct, since to explain knowledge in this way does not explain our philosophical motivation for scrutinizing our knowledge in the first place. Accordingly, Stroud claims:

"To say simply that we see, hear, and touch the things around us and in that way we know what they are like, would leave nothing even initially problematic about that knowledge. Rather than explaining it, it would simply state that we know. There is nothing wrong with that; it is true, but it does not explain how we know even in those cases in which (as we would say) we are in fact seeing or hearing or touching an object. That is what we want in a *philosophical* explanation of our knowledge." (Stroud, Understanding Human Knowledge, p145 * my italics)[204]

My proposal is this: to keep the externalist framework in place and to accommodate the legitimacy of the philosophical challenge presented by internalism, we need to allow that the type of knowledge we have is qualitatively different to the sort of knowledge that we would possess were it not afflicted by meta-epistemological scepticism. Our externalist knowledge is, in some respect, affected by our inability to fully speak to, and answer, internalist philosophical objections to externalism. We must allow that externalist knowledge is still possible, but we should concede that the absence of a suitable internalist philosophical perspective on that knowledge still *affects* our externalist knowledge.

However, the possibility of reconciling both legitimate intuitions in this way is difficult to accommodate within epistemic absolutism, which is normally presupposed in our inquiries. According to the absolutist picture, there are no degrees of propositional knowledge that *p* (See Ryle, 1949, Dretske, 1981). For example, Dretske offers the following characterization of knowledge that is paradigmatic of epistemic absolutism:

[203] Ibid.,
[204] Stroud, p.145

If we both know that today is Friday, it makes no sense to say that you know this better than I….In this respect, factual knowledge is absolute. It is like being pregnant: an all or nothing affair[205]

Absolutism is largely assumed in our epistemological inquiries and is an endemic feature of our philosophical thought experiments, which typically proceed by way of examining the presence of certain conditions of one's doxastic situation that either defeat, or fail to defeat, one's knowledge that *p*. It is part of our philosophical curriculum that absolutism is true, because it forms the foundation for so much philosophical research. The absolutist presumption, if true, would mean that propositional knowledge is not a matter of degree.[206] If there are no degrees of propositional knowledge, then meta-epistemological scepticism can only *affect* externalist everyday knowledge if meta-epistemological scepticism *excludes* externalist knowledge.

However, the difficulty for our absolutist template is that meta-epistemological scepticism *does* appear to affect knowledge, yet everyday knowledge still appears possible. We should be disinclined to accept that meta-epistemological scepticism excludes knowledge because this would commit us to global scepticism on the grounds that MS is an unavoidable feature of our epistemic condition.[207] For example, Duncan Pritchard thinks the lesson of reflective luck arguments is *not* that scepticism is true, but rather, than "a certain form of fine-grained knowledge which meets the relevant internalist conditions is impossible, though this is problem enough."[208] However, it is not clear how one can accommodate a more "fine grained" conception of knowledge when the question of what an agents knows is seen *purely* as a question of whether that agent's doxastic situation instantiates a set of invariantist conditions required for having absolutist knowledge that *p*.

However, not all philosophers accept the absolutist framework. Some philosophers adopt a gradualist conception of knowledge, which allows that in addition to having knowledge that *p*, one's knowledge that *p* can be good or bad, better or worse, depending on the quality of knowledge at stake (Hetherington, 2001).[209] The gradualist conception rejects epistemic

[205]Dretske, Fred, "The Pragmatic Dimension of Knowledge", *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, (1981), Vol. 40, No. 3, p. 363
[206] Dretske, "The Pragmatic Dimension of Knowledge", p363-364
[207] Pritchard, *Epistemic Luck*, (Oxford University Press, 2005), p.247
[208] Ibid.,
[209] Hetherington, Stephen, *Good Knowledge, Bad Knowledge,* (Oxford University Press, 2001)

absolutism. In rejecting absolutism, we can arrive at a template that accommodates our intuitions about meta-epistemological scepticism without necessarily denying us first-order knowledge in the actual world. On this view, I suggest, we may speak both to the internalist intuition *and* to the epistemological commonsensism that motivates externalism.

The difference between absolutism and non-absolutism is *not* to do with what conditions are required for knowledge, but instead hinges on how one's knowledge can be qualitatively affected by the extent to which the conditions for knowledge are satisfied; in temporal terms, the difference between absolutism and non-absolutism consists in what can subsequently happen to one's knowledge once a given set of conditions are met. Epistemic absolutism asserts that once a subject S meets a set of conditions C for knowing some proposition *p* in the actual world α, then S knows that *p* in α in virtue of S's belief that *p* satisfying C. Here, "C" functions as a knowledge threshold, and once C is met by S, S knows that *p*. Importantly, once C is met, then exceeding this threshold by having *more* justification for *p,* or a firmer belief that *p*, would have no effect on S's knowledge that *p*. This can be framed in modal terms: absolutism entails that if there is a possible world β in which S knows that *p,* but where S's degree of justification for *p* is stronger (or S's belief is firmer) in β than in the actual world α in which S knows that *p*, S's knowledge that *p* is nonetheless indiscernible across β and α.

In contrast, anti-absolutism preserves the same knowledge threshold as absolutism, but additionally allows that if S has a greater degree of justification for *p* in β than in α, then even though S knows *p* in both α and β, the knowledge that S possesses in β will be nonetheless qualitatively improved and therefore more desirable than the knowledge S has in α. S will have "better" knowledge in β on the anti-absolutist framework.

We need, importantly, to further qualify the anti-view: crucially, anti-epistemic absolutism does *not* imply relativism about knowledge, nor contextualism. On the anti-framework, there can still remain a set of invariantist conditions that are required to know *p*. For example, the anti-framework can still allow that one only knows *p* if *p* is true, if *p* is believed, and if there is justification for believing *p*, and so on. Anti-absolutism can still be strict about what counts as knowledge, stating that mere true-belief is not enough to know that *p*. For example, on an externalist version of anti-absolutism, we may state that knowledge requires a true belief that is produced by a reliable belief-forming process, or a belief being safe, or sensitive to the

facts. Either way, anti-absolutism will say that S knows that *p* if and only if S *absolutely* satisfies a set of conditions for knowing *p*, even if S's knowledge that *p* is *non-absolute*.[210]

The possibility of non-absolute knowledge allows for gradualism only if we accept that certain conditions for knowledge may be possessed in different degrees. For example, belief may be understood in terms of credence, and justification can be possessed to greater or lesser degrees. Some beliefs can be more justified (even if their justification is not any more sufficient - of course - for the purpose of having knowledge simpliciter) than others, on the grounds that the evidence for believing the relevant proposition is stronger. On the anti-view, exceeding the knowledge threshold by having *more* justification than suffices for knowledge will strengthen the knowledge possessed. Importantly, anti-absolutism may allow that there are *other* conditions that do not necessarily have to be met in order to have propositional knowledge, but where satisfying them can nonetheless improve the quality of knowledge at stake. A certain resistance to sceptical reasoning may strengthen one's knowledge that *p* by strengthening one's justification for *p*. [211]

Ernest Sosa (2009) recognizes the advantages of admitting that knowledge can be *better* or *worse* in the dialectic between internalism and externalism. Sosa claims that one can know some proposition in a way that is "better" than how other people know that same proposition.[212] To accommodate this intuition, Sosa distinguishes between two different types of knowledge, "animal knowledge and "reflective knowledge".[213] The former category describes a purely externalist knowledge that is attainable if an agent has a belief that originates from a truth-conducive source.[214] Sosa likens animal knowledge to "lucking into some benefit in the dark". The type of luck that arises in the case of animal knowledge is reflective, because animal knowledge does not require that the knower has an epistemic perspective on the belief. In contrast, "reflective knowledge" is of a type that arises when an agent can justifiably *endorse* the source of one's knowledge.[215] Sosa thinks that reflective knowledge is an altogether higher quality than animal knowledge, and that humans aspire to this "higher" reflective knowledge.[216] Sosa's distinction between animal and reflective

[210] See: Hetherington, *Good Knowledge, Bad Knowledge*, p.7
[211] Ibid., p30
[212] Sosa, Ernest, "Reflective Knowledge: Apt Belief and Reflective Knowledge: Vol 2", (Oxford University Press, 2009)., p136
[213] Ibid., p. 135
[214] Ibid., p. 136
[215] Ibid.,
[216] Ibid.,

knowledge is an attempt to couple our internalist intuitions about what human beings demand from philosophical knowledge, with the preservation of epistemological commonsensism that militates against scepticism.

Rejecting absolutism allows us to take meta-epistemological scepticism seriously, as having ramifications for our knowledge, whilst simultaneously allowing us to have the very knowledge that scepticism tries to deny.[217] Applied to our conundrum, the anti-absolutist understanding of knowledge allows us to show that the internalist places unreasonable demands on knowledge simpliciter, but that nonetheless, his objections are illuminating because they show that we lack a perfectionist or fine-grained type of knowledge.[218] If meta-epistemological challenges aim to make a point about the externalist analysis, namely, that the externalist analysis offers us a conception of knowledge that speaks past important philosophical questions, then internalism would be well-advised to embrace anti-absolutism, otherwise it has to explain how it is we can have internalist non-inferential justification whilst preserving our common-sense intuitions. Likewise, if the externalist only wants to offer an account of knowledge that accommodates our commonsense intuitions, and if naturalizing our epistemic commitments is the best way of achieving this, then he would also be advised to accept anti-absolutism; namely, this is because anti-absolutism can give us externalist knowledge – and that is all he needs in order to accommodate our anti-sceptical intuitions. In other words, our externalist knowledge does not need to be perfect or fine-grained, or indeed, philosophically interesting; it merely needs to vindicate the commonsense intuition *that we have knowledge*. By accepting anti-absolutism, externalists can accommodate this intuition, conceding nothing to the internalist, without having to engage in the project of refuting the *legitimacy* of the internalist quest for a perfect and fine-grained conception of knowledge.

On the grounds that any adequate account of knowledge should, where possible, attempt to unify competing intuitions about what knowledge demands, rather than press for a particular conception of knowledge that renders competing intuitions parochial, the concession that our externalist knowledge is not as complete as we might think seems an appropriate one to make. This would give us externalist knowledge whilst acknowledging that the internalist meta-epistemological challenges plays an important role in helping us understand our epistemic predicament and the quality of knowledge we possess in the actual world.

---

[217] Hetherington, *Good knowledge, Bad knowledge*, p. 34

[218] Hetherington, p. 65. See also: Pritchard discussion of us lacking a "fine grained" conception of knowledge in *Epistemic Luck* (Oxford University Press), p.247

The meta-epistemological challenge is in this respect a challenge for internalism and externalism alike, because it challenges *any* analysis of knowledge that is wedded to absolutism. It challenges absolutists who accept externalism because it means they cannot accommodate the internalist intuition that there is something wanting about the externalist conception of knowledge. And it challenges internalist forms of absolutism because it means that, insofar as the meta-epistemological challenge is legitimate, internalist absolutism must accept the unwelcome conclusion that all putatively true beliefs are cases of knowledge failure. However, I have argued that anti-absolutism can satisfy the intuition that there is something deeply wanting about the externalist analysis whilst allowing that we can at least possess imperfect externalist knowledge about the world. In recognizing that our externalist knowledge is imperfect, we can vindicate the legitimacy of the internalist meta-epistemological challenge without being driven to scepticism.

Bibliography

BonJour, Laurence, *Epistemology: Classic Problems and Contemporary Responses*, (Rowman and Littlefield, 2002)

BonJour, Laurence. Sosa, Ernest, *Epistemic Justification: Internalism vs. Externalism, Foundationalism vs. Virtues*, (Blackwell, 2003)

Bergmann, Michael, "Externalism and Skepticism", *The Philosophical Review,* (2000) Vol. 109, No. 2, pp, 159-194

Breyer, Daniel, "Reflective Luck and Belief Ownership", *Acta Anal* (2010), Vol. 25, pp. 122-154

Cohen, Stewart, "Basic Knowledge and the Problem of Easy Knowledge", *Philosophy and Phenomenological Research*, Vol. 65, No. 2 (2002), pp. 309-329

Comensana, Juan, "Unsafe Knowledge", *Synthese*, (2005), Vol. 146, pp. 395-404

Conee, Earl. Feldman, Richard, "Internalism Defended", *American Philosophical Quarterly*, (2001), Vol. 38, No. 1, pp. 1-18

Conee, Earl. Feldman, Richard, "Evidentialism", *Philosophical Studies,* (1985), Vol. 48, pp. 15-34

DeRose, Keith, "Solving the Skeptical Problem", *The Philosophical Review,* (1995). Vol. 104, No. 1, pp. 1-52

Dretske, Fred, "Is Knowledge Closed Under Entailment? The Case against Closure" in *Contemporary Debates in Epistemology*, Edited by Mathhias Steup and Ernest Sosa (Blackwell, 2005), pp. 13-24

Dretske, Fred, "The Pragmatic Dimension of Knowledge", *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, (1981), Vol. 40, No. 3, pp. 363-378

Engel, Mylan, "Personal and Doxastic Justification in Epistemology", *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, Vol. 67, No.2 (1992), pp. 133-150

Engel, Mylan, "Is Epistemic Luck Compatible with Knowledge?", *The Southern Journal of Philosophy",* (1992), Vol. 30, No. 2, pp. 59-75

Enoch, David; Schechter, Joshua, "How Are Basic Belief-Forming Methods Justified?", *Philosophy and Phenomenological Research,* (2008), Vol. 76, No. 3, pp. 547-579

Enoch, David, *Taking Morality Seriously: A Defence of Robust Realism"*, (Oxford University Press, 2011)

Lemos, Noah, "Commonsensism in Ethics and Epistemology", in *Knowledge, Truth, and Duty: Essays on Epistemic Justification,* Edited by Matthias Steup (Oxford University Press, 2001)

Fumerton, Richard, *Metaepistemology and Skepticism*, (Rowman and Littlefield, 1995)

Fumerton, Richard, "Skepticism and Naturalistic Epistemology", *Midwest Studies in Philosophy*, (1994), Vol. 19, Issue. 1, pp. 321-340

Fumerton, Richard, "Epistemic Internalism, Philosophical Assurance and the Skeptical Predicament", in *Knowledge and Reality*, (Springer, 2006), pp. 179-191

Gettier, Edmund, "Is Justified True Belief Knowledge?" *Analysis,* (1963), Vol. 23, No. 6, pp. 121-123

Goldman, Alvin I, "A Causal Theory of Knowing", in *The Journal of Philosophy*, (1967) Volume. 64, No. 12,

Goldman, Alvin, "The Internalist conception of Justification", *Midwest Studies in Philosophy,* (1980), Vol. 5, pp. 27-52. Republished in *Epistemology: Internalism and Externalism*, Edited by Hilary Kornblith, (Blackwell, 2001), pp. 36-68

Goldman, Alvin, "What is Justified Belief", in *Epistemology: An Anthology,* (Blackwell, 2008), pp. 340-354. Originally published in *Justification and Knowledge*, Edited by G.S. Pappas (Riedel, 1976), pp. 1-23

Hawthorne, John, "The Case for Closure", in *Contemporary Debates in Epistemology*, Edited by Matthias Steup and Ernest Sosa (Blackwell, 2005), pp. 26-40

Heil, John, "Foundationalism and Epistemic Rationality", *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, (1982), Vol. 42, No. 2, pp. 179-188

Hetherington, Stephen, *Good Knowledge, Bad Knowledge,* (Oxford University Press, 2001)

Hookway, Christopher, "Critical Common-Sensism and Rational Self-Control", *Nous*, (1990), Vol. 24, pp. 397-412

Kim, Jaegwon, "What is Naturalized Epistemology?" in *Epistemology: An Anthology,* (Blackwell, 2008). Originally published in J. Tomberlin *Philosophical Perspectives,* 2. Epistemology, pp. 381-405

Kripke, Saul. A., *Naming and Necessity*, (Oxford University Press, 1980)

Moore, G.E, "A Defence of Common Sense", in *G.E. Moore: Selected Writings*, (Routledge, 1993). Originally published in *Contemporary British Philosophy* (2nd series), ed. J. H. (Muirhead, 1925)

Nozick, Robert, *Philosophical Explanations*, (Oxford University Press, 1981)

Pritchard, Duncan, "Anti-Luck Epistemology", *Syntheses*, (2007) Vol. 158, pp. 277-297

Pritchard, Duncan, *Epistemic Luck*, (Oxford University Press, 2005)

Pritchard, Duncan, "Epistemic Luck", *Journal of Philosophical Research*, (2004), Vol. 29, pp. 193-222

Pritchard, Duncan, "Sensitivity, Safety, and Anti-luck Epistemology", in *The Oxford Handbook of Skepticism,* Edited by John Greco (Oxford University Press, 2009)

Pryor, James, "The Skeptic and the Dogmatist", *Nous*, (2000), Vol. 34, No. 4, pp. 517-549

Putnam, Hilary, *Reason, Truth, and History*, (Cambridge University Press, 1981)

Quine, W.V., "Epistemology Naturalized" in *Epistemology: An Anthology,* (Blackwell, 2008) Originally published in W.V. Quine, Ontological Relativity and Other Essays (New York, Columbia University Press, 1969), pp. 69-90

Rosenthal, David. M., "Self Knowledge and Moore's Paradox", *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, (1994), Vol. 77, No. 2/3, pp. 195-209

Shafer-Landau, Russ, *Moral Realism: A Defence* (Oxford University Press, 2005)

Smithies, Declan, "Moore's Paradox and the Accessibility of Justification", *Philosophy and Phenomenological Research*, (2011), Vol. 85, Issue 2, p. 273-300

Sosa, Ernest, "How to Defeat Opposition to Moore", *Nous*, (1999) Vol. 33, pp. 141-153

Sosa, Ernest, "Reflective Knowledge: Apt Belief and Reflective Knowledge: Vol 2", (Oxford University Press, 2009)

Sosa, Ernest, "Skepticism and Contextualism", *Nous*, (2000), Vol. pp. 1-18

Stroud, Barry, *Understanding Human Knowledge: Philosophical Essays*, (Oxford University Press, 2000)

Stroud, Barry, *Understanding Human Knowledge in General*, in *Epistemology: Internalism and Externalism*, Edited by Hilary Kornblith, (2001), pp. 126-146. Originally published in *Knowledge and Skepticism*, Edited by M. Clay and K. Lehrer (Westview Press Inc., 1989)

Steup, Matthias, "Belief Control and Intentionality", *Synthese*, (2012), Vol. 188, pp. 145-16

Unger, Peter, "An Analysis of Factual Knowledge", *The Journal of Philosophy* (1968), Vol. 65, No. 6, pp. 157-170

Wright, Crispin, "The Perils of Dogmatism", in *Themes from G.E. Moore: New Essays in Epistemology and Ethics*, Edited by Susana Nuccetelli and Gary Seay, (Oxford University Press, 2007). pp. 25-49

Wright, Crispin, "Anti-Sceptics Simple and Subtle: G.E. Moore and John McDowell", *Philosophy and Phenomenological Research*, (2002), Vol. 65, No. 2, pp. 330-348

White, Roger, "Epistemic Permissiveness", *Philosophical Perspective*, (2005), *No. 19*

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

Zagzebski, Linda. “The Inescapability of Gettier Problems”, *The Philosophical Quarterly*, Vol. 44, No. 174, (1994), pp. 65-73

Zagzebski, Linda, *Virtues of the Mind* (Cambridge University Press, 1996)