

EPISTEMIC SECURE REALISM

Joost Ziff

ABSTRACT

I argue for epistemic secure realism, the view that a successful conscious act to form the knowledge that p , where p is a worldly proposition, involves coming into secure cognitive contact with the fact that p by exercising our capacity for rational (worldly) thought. I do this through a transcendental argument against a thesis whose truth would preclude epistemic secure realism: the ‘common kind justification thesis’, the view that indistinguishable epistemic states must have the same epistemic status. This thesis constrains the possible ways justification can be plausibly structured. I argue against two of these positions: common kind foundationalism and common kind coherentism. Against the foundationalist, I marshal a Wittgensteinian-style private language argument. Against the coherentist, I make a similar rule-following argument using a new brain in a vat scenario I call ‘nonsense in a vat’, in which our existing epistemic practices turn out to seem coherent but instead are nonsensical. Both views, I argue, violate indubitable phenomenological theses and thus will have to be discarded. In response, I set out epistemic secure realism as an intuitive way to avoid the paradoxes that I developed against foundationalism and coherentism while still maintaining the phenomenologically plausible notion of judgment. To help develop this position, I weave together conceptual materials from Ernest Sosa and William James having to do with judgment and conscious volition. Finally, I sketch out potential responses to objections that epistemic secure realism might face concerning justification. Specifically, I address brain in a vat cases, socially justified beliefs/affirmation, and skeptical worries about the correct procedure of judgment.

LAY SUMMARY

I give an argument for what I call “epistemic secure realism.” By this I mean that knowing, just like seeing, is a kind of ability that we humans have. This contrasts with, for example, the position that knowing is more like ‘gambling’. Instead of knowing for sure, you just try your best to make sense of how the world ‘seems’, in the hopes that if the world does exist you would get it right. So, in a way, if the world exists you would “know” as much as when the world does not exist because you would have the same information. On the contrary, I suggest that your knowledge of the world would be drastically different if the world exists. To show this, I analyze the concept of knowledge to figure out what beliefs about knowledge we are committed to and thus have to believe, even if we tried our hardest not to. As a result, I argue that we must view successful knowledge as being about how the world actually is rather than about how the world looks to be the case. Thus, if we have to view successful knowledge as being about the actual world, then the best way (I argue) to make sense of that commitment is if we understand knowing as a real power we have in the actual world.

CONTENTS

1	INTRODUCTION	1
1.1	WHAT EPISTEMIC SECURE REALISM IS	1
1.2	WHY I AM ADVOCATING FOR EPISTEMIC SECURE REALISM	2
1.3	HOW I INTEND TO ADVOCATE FOR EPISTEMIC SECURE REALISM	3
2	THE BEGINNING TRANSCENDENTAL AND PHENOMENOLOGICAL MATERIALS	5
2.1	RATIONALITY, AFFIRMATION, AND JUDGMENT	5
2.2	KANTIAN SKEPTICISM AND THE PHENOMENOLOGICAL METHOD	7
2.3	FROM INDUBITABLE TO SENSE TO INDUBITABLE POSSIBILITY	10
2.4	RECAPPING THE ARGUMENT	13
3	THE MYTH OF THE GIVEN PARADOX	15
4	THE BASICS OF SOSA'S THEORY	19
4.1	SOSA ON BELIEF AND KNOWLEDGE	19
4.2	SOSA ON JUDGMENT AND AFFIRMATION	20
4.3	SOSA'S COHERENTISM	21
5	THE COHERENTIST PARADOX	25
5.1	A FIRST PASS	25
5.2	DEFUSING THE TRIANGULATION ARGUMENT	28
5.3	A SECOND PASS	32
6	SETTING OUT EPISTEMIC SECURE REALISM	35
6.1	DIAGNOSING THE PROBLEM	35
6.2	CONSCIOUS ATTENTION AND THE WILL	36
6.3	REASONING AND FULLY CONSCIOUS JUDGMENT	37
6.4	J-ENDEAVORS AND THE CRITICAL DOMAIN	38
6.5	THE WORLD-ORIENTED JUDGMENTAL COMPETENCE	40
7	ROUNDING OFF EPISTEMIC SECURE REALISM	45
7.1	EVIL DEMONS AND VIRTUAL REALITIES	45
7.2	PROBLEMS IN SOCIAL JUSTIFICATION	46
7.3	THE CORRECT PROCEDURE FOR GLOBAL JUDGMENT	48
7.4	CONCLUDING REMARKS	49

1 INTRODUCTION

The purpose of this work is to advocate for a certain view of justification, rationality, and judgment that I call *Epistemic Secure Realism*. I call the position such because I intend it as a very explicit nod at the use of term ‘direct realism’ in the philosophical literature on perception. I will first give a short description of the position. After that, I will explain my motivation for the project and how the position fits within the literature. Finally, I will give the overall plan of the work.

1.1 WHAT EPISTEMIC SECURE REALISM IS

Perceptual direct realism or ‘naive realism’ is the idea that experience presents mind-independent objects as its constituent parts. In other words, direct realism simply takes the ‘naive’ view of experience, namely that it is exactly what we think it is. Perceptual direct realism was a mostly unpopular position in analytic philosophy until it was advocated for in various forms by John Hinton, Paul Snowdon, and John McDowell¹. Mostly this was done through the use of a particular bit of conceptual machinery known now as ‘disjunctivism’. Disjunctivist argumentation proved so effective within philosophy perception literature because it purported to defuse what is known as the ‘argument against illusion/hallucination’. Since veridical perceptions of objects were indistinguishable from illusory or hallucinatory episodes, the argument goes that these different types of episodes share some sort of (metaphysical/epistemological) commonality. Disjunctive argumentation rejects this line of reasoning (by, e.g., providing counter examples) in the hopes of opening up dialectical space for direct realism. What originally seemed to be one of the same type of experience (‘seeming to see’) could now actually be two distinct ‘disjuncts’ with radically different natures (‘merely seeming to see’ vs. ‘actually directly seeing’).

Disjunctivism is typically viewed as a strategy exclusive to the philosophy of perception. However, this is not necessarily the case. For one McDowell’s version of disjunctivism has always had an epistemic tint to it. For another, Duncan Pritchard has more recently advocated for disjunctivism of perception within a solely epistemic context. But both of these positions still only understand disjunctivism and direct realism as closely related to perceptual knowledge. While epistemic knowledge in general is not ‘direct’, it should be seen as ‘secure’; as an exercise of a worldly competence akin to sight. Perceptual direct realism thinks that successful seeing just involves seeing ‘directly’ by exercising our capacity for sight. *Epistemic* or *intellectual* secure realism thinks that successfully coming to know that p just involves *consciously cognizing* p ‘securely’ by exercising our capacity for rational (worldly) thought.

Now it is tempting to think that I am advocating for this view as a way to extend a ‘common sense’ metaphysical view of the world, one on which there are trees, tables, geese, fish, and human beings. However, this is not necessarily the primary concern of the work. Rather, the argument for epistemic secure realism merely requires that we see

¹See the introduction and articles of Haddock and MacPherson 2008 for more on the history of this literature

a successful judgment that p function as coming to cognize that p in a way that is akin to seeing in its metaphysical connection to what there is. Here, we can fill in whatever interpretation of the metaphysics of ‘ p ’ being the case that we like; so long as the contents of our ‘comings to know that p ’ is exactly the kind of thing that is the case. In other words, there is no difference between the picture successful knowings provides us and the way the world actually is.

1.2 WHY I AM ADVOCATING FOR EPISTEMIC SECURE REALISM

The original intent for this work was to reconstruct an argument that McDowell presented in his book *Mind and World*. McDowell’s work has had a looming influence on disjunctivist thought, and much of the epistemic side of these contributions can be found in there. McDowell’s work is so interesting because it aims to solve a certain type of skeptical attitude in philosophy that seems pervasive after the successes of enlightenment science. The attitude as such seems to be that our perception and our common sense thoughts about the world are in some ways deceptive or ‘false’. McDowell’s answer to this view is novel because it points at epistemic rather than metaphysical considerations. As long as we hold very ‘basic’ ideas about thought’s answerability to the world, we will be forced to take on board these epistemic considerations and thus accordingly have to revise our metaphysical views. Unfortunately, as Crispin Wright said in his review of *Mind and World*, “McDowell is a strong swimmer, but his stroke is not to be imitated” (1996, 252), which is to say, McDowell’s work is densely presented and hard to follow. As a result, the argument I present is more of a parallel reconstruction of his point rather than a faithful interpretation. With this work, I hope to make plausible the same type method that McDowell uses, while expressing what are certainly wholly different theses. At the most, this work could be called McDowell ‘inspired.’

That said, I will not be able to truly spell out the metaphysical implications within this work. Instead, the focus here will be to spell out epistemic secure realism as a systematic epistemological theory, which (I hope) will not be any the less consequential. Epistemic secure realism as motivated will imply a view according to which justification and rationality flow forth from specific worldly powers of thought. Put in more concrete terms, this means that any philosophical view where any subject S is necessarily as rational as a subject T because S and T justifiers are internally indistinguishable from each other, will simply be wrong. This does not mean that something seeming the case cannot provide any justification. If the world is there, and our perceptual faculties are generally reliable, then something’s seeming the case certainly does provide justification for it being the case, and it would be rational to think so. But, it is incorrect to presume that something’s *mere* seeming the case provides justification. In fact, in a world where we had no power to get at the world through some metaphysically integrated power like that allowed by epistemic secure realism, seeming the case’s would provide no justification at all. Moreover, in such a world a subject might not even be rational. As such, epistemic secure realism would not only contradict many internalist views of justification, but some externalist views of justification as well. Any view that does not see good worldly justification and good rationality as intrinsically linked to metaphysical facts will be to a certain extent

incompatible.

1.3 HOW I INTEND TO ADVOCATE FOR EPISTEMIC SECURE REALISM

Chapters 2–4 concern themselves with the negative argument of the work. In here, I develop a transcendental argument that closely mirrors that of McDowell's in *Mind and World*. The target of this transcendental argument are two types of views: givenism (bad foundationalism) and coherentism. The rough structure of this transcendental style argument is roughly as follows:

1. I am forced to think of concept *C* as sensible.
2. If am forced to think of concept *C* as sensible I am forced to think of concept *C*'s instantiation as possible in some particular fashion *F*.
3. If view *X* is true, then concept *C*'s instantiation is impossible.
4. But I am forced to think of concept *C*'s instantiation as being possible in some particular fashion *F*, so I am forced to think of view *X* as being false.

This argument will then be played out against two views of the normativity of worldly thought: givenism (that is, defective foundationalism) and mere coherentism.

First, in chapter 2 we will set out the general method of the above argument. This involves clarifying concept *C*, sensibility, and the particular fashion in which *C*'s instantiation is possible. The particular concept *C* in this case will be the idea of 'judging that *p*', where *p* is some worldly proposition. Through phenomenological argument I will aim to establish that the idea of agentially succeeding to affirm correct things about the world is what I will call 'indubitably sensible.' Indubitable sensibility, in being indubitable (not doubtable), will designate some necessary feature of the phenomenology of affirmation. This notion I will then expand into a notion of 'indubitable possibility.' This notion of possibility will require the right type of modality for the next step of the argument.

The next part of the negative argument involves presenting and criticizing the givenist and coherentist types of views of epistemic normativity. What these hold in common—and therefore what I will deny so as to build my view in response—is a view of normativity that can function separate from the world itself. Seen as such, we can only derive normativity in two ways. Either we take certain experiential states as 'foundational', giving them epistemic preferential status over others, or we function using a 'web' model of justification where our goal is merely to make our system of beliefs 'cohere.'

The former is done by the givenist, which I critique in chapter 3. The givenist is a type of internal foundationalist that is very similar to the type of foundationalism that Wittgenstein targeted in his private language argument. The type of language interpreter Wittgenstein attacked was one that took the base pre-theoretic signs of language

as ‘data’ and built up from there a theory reinterpreted in terms of those signs. A analogical epistemic view would be a type of foundationalist internalism that takes internal seemings and tries to build a theory of justification by using them as the ‘axioms’ of the epistemic system. The analogy here however is not by accident; the argument I will launch against this view will be derived in part from textual analysis of passages from Wittgenstein’s philosophical investigations.

The latter argument is made by the coherentist. However, instead of critiquing a generic coherentist view, I will instead critique a specific and more complicated version as set out by Ernest Sosa in his most recent book *Judgment and Agency* (2015). I do this for two reasons: first, as an opportunity to show how even more complicated externalist views can have problematic internal conceptual parts. Second, to take the opportunity to build on Sosa’s system to make my own. Sosa’s view is an ambitious extension of his original epistemological views, connecting them to a more general metaphysically oriented analysis of the notions of human endeavoring and achievement. Yet, Sosa endorses a coherentist thesis for internal justification, and thus will fail against the paradox I intend to develop. Thus, I will first explain Sosa’s views in chapter 4 and then develop the coherentist paradox in chapter 5.

In chapter 6.1, I set out epistemic secure realism. My view is a modification of Sosa’s view, so—to give a car analogy—I must replace the defective conceptual machinery to get Sosa’s theory of judgment up and running again. As I will conclude, the problem will be that Sosa’s view of agentiality in justification possesses an iterative structure: judging well requires making sure that the process by which we affirm is reliable, which in turn requires that the process we use to check whether our affirmation is reliable, and so on. It is this notion of agency that generates his coherentist view, and thus is what has to go. In its place I will put a tweaked view of ‘will’, as presented by William James. This view of agentiality will not focus on iterative structure, but rather on conscious and deliberate will. I will argue that this view then summarily avoids the givenist and coherentist paradoxes.

Finally, in chapter 7, I sketch out how such a view could deal with more traditional problems of justification. First, there is the problem of various evil demon and virtual reality scenarios. Second, there is the problem of testimonial knowledge. Third, there is the problem determining the correct procedure for the global power of judgment. The first two will be answered in terms of the manifestation conditions of the judgment power. Evil demon scenarios could be seen as masking the environmental conditions for manifestation, while thorough-going unfriendly social environments could be seen as masking the *social* manifestation conditions of justification. Meanwhile, I will argue that in the good case of the global power of judgment the judging procedure would have to be within our reach. The world itself, in impressing itself on us and our social peers, would teach us how to judge it correctly. Thus, though judgment might be difficult, the worry that we could never succeed in it is implausible due to the many corrective mechanisms that embodied and social existence provides.

2 THE BEGINNING TRANSCENDENTAL AND PHENOMENOLOGICAL MATERIALS

In this chapter I will lay the groundwork for the transcendental argument against internal views of rationality. Firstly, I will set out the target view of rationality and the relevant notions of affirmation and judgment. For this, I will be referring to the relevant passages of McDowell and Sosa so as to set up the argument for later chapters. Secondly, I will explicate the particular phenomenological method I will be using and explain how that notion can provide the material for a transcendental argument. This will involve explicating the distinction between what James Conant refers to as Kantian and Cartesian skepticism. Lastly, I will set out the particular phenomenological theses we will be using to get the transcendental argument of the ground. Here I will be working with the notions of sense and conceivability that Edmund Husserl, Gottlob Frege, and David Chalmers provide.

2.1 RATIONALITY, AFFIRMATION, AND JUDGMENT

Epistemology as it is practiced today is mostly concerned with knowledge and therefore also concerned with beliefs, since they bear some sort of relation to knowledge, be it under a JTB model or otherwise. But what will be important for this paper is not knowledge or belief, but the ideal process which can be stored as knowledge. What we are concerned here then is what Sosa calls “endeavors”, which have “freely determined aim[s]” (2015, 192), but what equally might be called ‘tryings’, given the phenomenological arguments to come. The specific endeavor we are interested in is one where we are trying to, as Sosa says, “affirm with apt correctness” (2015, 66). Apt here for Sosa means “accurate because adroit” (2015, 15), which means that it must be true because it was skilled. Already now, the notion of judgment can sound very theory-laden, but I want to emphasize that this is not the case at all. If endeavors are simply ‘tryings’ and affirmations are simply the non-spoken inner analogues of assertions, then a judgmental endeavor is simply the process we go through when we try our best to make sure that we are going to affirm is true, and thus if we were to store it as a belief it would be *knowledge*. It would not be enough to require us to endeavor to affirm merely correctly rather than aptly correctly. Sosa has pointed out that such an act would be compatible with “guessing” (2015, 75), where we simply are trying to affirm correctly, but without a sense that the process we are going through to decide what we will affirm in some way *ensures* its truth. Since we are intentionally allowing this amount of risk into the process, whatever belief we store as a result would surely fall short of knowledge. Thus, we need to aim not just for correctness for our affirmations in our endeavor, but aim to affirm correctly through the process of the self-same endeavor. Which is to say, we are going to try to affirm in such a way that ensures truth of our affirmation, i.e., to intend to affirm in such a way our skill of affirming ensures the truth (aptly).

Now given that we are trying to affirm in such a way, we are engaging in an endeavor that is *normative*; after all we could succeed or fail in such an endeavor. Because it is an epistemic normative endeavor, it would make sense for us to be searching for something

like ‘justification’ or ‘reasons’ for forming a particular affirmation. The process one goes through in the endeavor is that of ensuring that one is justified in a certain way, or that one does have good reasons of a certain type. At the minimum, we are trying to assure that we are both *ex ante* justified and *ex post* justified. We are trying to be *ex ante* justified, because we are trying to affirm the correct right proposition for which we have the best evidence/reasons (whatever notion we might pick). We are trying to be *ex post* justified, because we are trying to ensure that the resulting affirmation, when so affirmed, is indeed justified based on what reasons/evidence there is. The distinction as such collapses in the case of the success of the epistemic act. Our endeavor then consists in something like ‘navigating’ the existing set of beliefs and their justificational interrelations in such a way that the resulting affirmation is justified/reasonable. When we do this job well, we are presumably *rational*.

Now we can formulate the principle that we will be critiquing as part of the negative argument. I take here as a basis for the principle Littlejohn’s formulation of “Standard Mentalism,” according to which “rationality supervenes upon an individual’s perspective” (2018, 2). Littlejohn gives this idea cash value in terms of “non-factive mental duplicates” but it will be important for the argument that we instead spell this idea out in terms of ‘phenomenological indistinguishability.’ We will have more to say about phenomenology in the upcoming sections, but for now, we can define the principle as follows:

COMMON KIND JUSTIFICATION THESIS

Indistinguishable epistemic states must have the same epistemic status.

If we assume this condition, we are constrained in the way that epistemic justification can be structured. Now, according to the traditional Agrippan trilemma, there are three general ways epistemic justification can be structured:

1. Foundationalist: Certain beliefs are foundational and need no further justification.
2. Coherentist: All beliefs need further justification, and justification goes around in a circular fashion.
3. Infnitist: All beliefs need further justification, but justification may not go around in a circular fashion.

Space constraints prevent me from doing full justice to infinitist-type views¹ (though it is not particularly popular in the literature). That leaves us with foundationalism and

¹For more on infinitism see Klein and Turri 2013 and Klein 1999. Though I have not expanded on it below, I find it plausible that at least some of the upcoming arguments against coherentism could be adapted to work against infinitism. I will argue that a coherentist structure just in being coherent need not be sensical. We might similarly say that an infinitist structure, just in being coherently infinitist (setting aside whether such a thing is even possible), need not be sensical either.

coherentism, twinned with the common kind justification assumption.

We now have judgmental endeavors, the system of epistemic justification that the judger must reflectively access so as to succeed in such an endeavor, and rationality, the norm that determines the correct process one must employ before deciding to judge. The task now is to provide an argument that shows how coherentism and foundationalism when twinned with the common kind justification assumption must fail, thus forcing us to give up said assumption. Before we can do that though, we need to develop the general method of the argument; that is, the phenomenological method.

2.2 KANTIAN SKEPTICISM AND THE PHENOMENOLOGICAL METHOD

Most philosophers—if not most of the general public—is aware of what is generally popularly known as *Cartesian* skepticism. What is less well known is the related notion of *Kantian*² skepticism. Kantian skepticism is focused on calling into question the very *contentfulness* of its target. An easy way to draw this distinction is to distinguish Kantian skepticism of perception from Cartesian skepticism of perception. While Cartesian skepticism about perception is focused on the question of whether perceptual experience is reliable, Kantian skepticism instead questions whether it even makes sense for our sense-perceptions (be they right or wrong) to be characterized in any sort of way, i.e., whether they *seem* to be like anything in particular (Conant 2012, 14–15). Kantian skepticism also distinguishes itself from Cartesian skepticism in being inherently theoretically unstable (35–36), in that it seems to lead us to deny something that intelligibly *cannot* be denied. In the case of perception, Kantian skepticism would force us to deny the seemingly unavoidable fact that our experience strikes us a certain way, for instance, as presenting a tree, a cup, or any other sort of thing. I claim therefore we can use a Kantian skeptical problem productively: if we show that a particular philosophical view would force us to deny a certain aspect of contentfulness belonging to a certain indubitable fact of our lives, we will be able to dismiss it out of hand. After all, while it could be possible to intelligibly deny that experience does not provide a certain warrant, it seems absolutely impossible to deny that experience strikes us a certain way. This would be akin to claiming not just that our experiences are illusory, but that even our very *having experiences* is illusory.

As it will become clear further on, I claim that we can get the kind of indubitable facts we need by practicing phenomenology. However, it is important to realize here that I do not intend to use the word ‘phenomenology’ as it is used casually in most of analytic philosophy. Typically, a phenomenological fact thought is to be some description of phenomenal state of affairs, usually involving ‘seems’ like ‘right now, it seems like I see a cup on the table’. These types of locutions have been traditionally thought of as non-committal and indubitable, though this has been disputed more recently.³ To remain

²See Conant 2012 for more on this distinction. N.B. that the terms ‘Cartesian’ and ‘Kantian’ skepticism are not meant to reflect any positive claims about Descartes’s and Kant’s *actual* skeptical views; they are merely shorthand labels.

³See Martin 2004 and Schwitzgebel 2008.

neutral on these disputes at this point in the dialectic, we will use the term to *definitionally* circumscribe the unarguable facts of our phenomenal way of being. This then provides the following (rough) taxonomy⁴:

1. Phenomenal ways of being: how things ‘appear’ to us from a first person perspective. (what is normally understood as ‘phenomenology’)
2. Phenomenal fact: a propositional description of how things ‘appear’ to us from a first person perspective, usually in the form of ‘it seems that *p*’.
3. Phenomenology: the a priori pursuit into the ‘logic’ or ‘λόγος’ of our phenomenal ways of being, i.e., a characterization of our phenomenal ways of being that is indubitable.
4. Phenomenological fact: An indubitable proposition that is the result of a successful phenomenological investigation.

Thus, according to these definitions, phenomenological investigation is always definitionally *possible*. What philosophers might disagree on is what phenomenological facts such investigation would yield (if any).

What this investigation looks like will become clearer as we proceed. So, let us continue by setting out the particular phenomenological claims which will be defended. Take the following (supposedly) unarguable phenomenological thesis, which involves some of the notions set out in § 2.1:

THE JUDGMENTAL ENDEAVOR PHENOMENOLOGICAL THESIS

I seem to be able to succeed in an endeavor to make apt affirmations (an endeavor to judge correctly) about my environment.

Given that this paper falls within the field of epistemology, there should, in some sense, be no need to justify this thesis. After all, Cartesian-style skepticism is almost paradigm-

⁴The reason why I go through this redefinition is in part to provide a more charitable reading of early phenomenologists like Husserl. Husserl defined phenomenology as an accounting of the concept of consciousness “in a manner which *cuts out all relation to empirically real existence*” ([1901] 2001, 82). This particular accounting of consciousness is one that Husserl thinks springs from our inner awareness of our experiences, which Husserl understands in terms of Descartes’s ‘Cogito ergo sum’ ([1901] 2001, 86–87). Husserl’s intention here seems to provide a scope for his investigation that, as he says, does not “depend on the knowledge and acceptance of ideas about the ego which have always remained questionable.” Instead of judgments of inner perception as his focus, he rather “ask[s] what could belong to this conceptually undemarcated and therefore unutterable kernel [of our empirical ego notion], what may constitute the self-evidently certain” and *then* arrives at these inner judgments. This would (charitably) seem to be what I define here, namely that phenomenological investigation is definitionally *a priori* and indubitable.

matically an epistemological subject, and the whole point of typical Cartesian-style skepticism is to show how that our ability to successfully endeavor to affirm is illusory, which in turn suggests that it has some sort of prior phenomenological plausibility to begin with. That said, we can say *something* to make this principle more plausible (though a full derivation would perhaps best be left to a different paper). We can be said to ‘derive’ this principle by weaving together four different yet closely related notions: the ‘environment’, ‘affirmations’, ‘endeavorings’, and ‘environmental epistemic practices’. First, we can see our notion of ‘environment’ as derivative of our appreciation of empirical experience as providing a contiguous synthetic whole. For instance, I do not just experience cups, trees, and colors separately; I experience them together in time and in space, structured together in a certain way. The notion of our ‘environment’ follows by taking this synthesized whole as a topic in and of itself for research. Second, we already find ourselves with a notion of affirmation simply by reflecting on our ability to *characterize* our experience. Even fairly non-committal locutions such as ‘I seem to see a tree’ already express an affirmational structure. Yet, an ability to make affirmations about our environment is not enough if we do not see it as something that *seems* to be under our control as a freely-made endeavors. So, third, we presuppose in the act of doing philosophy itself a notion of endeavoring. Just as Sosa does, we can draw an analogy to the notion of endeavoring in Descartes’s *Meditations*. Descartes’s description of himself questioning and justifying his beliefs is to be understood as Descartes describing himself taking up a certain project to underpin his beliefs in a certain fashion. The notion of ‘taking up a project’ here already contains the notion of endeavoring, because it involves us doing something with an aim that belongs to *ourselves*, that is, one that flows forth from an conception of our own particular desires and goals. Last, take the notion of our existing epistemic practices for getting the environment right. For instance, to determine whether a chair is ‘really’ there and that I am not under an illusion, I might go up to it, touch it, swing it back and forth, and so on. In other words, we find ourselves with a preexisting set of epistemic investigatory practices that we do to (seemingly) gain more certainty in our environmental facts, and in fact that *seem* to assure us of success.

Weaving these four notions together we get something very close to our phenomenological principle. Insofar as I understand myself as seemingly able to endeavor, I understand myself as seemingly able to take up certain projects. Insofar as I understand the environment as an imagined object I can aim at, I can take up the project of trying to get it right. Insofar as I understand myself as able to conduct my endeavor in such a way that it is expressible, I understand myself as attempting to characterize this imagined object of the environment through using affirmational acts. Finally, by treating our normal epistemic practices as if they were sufficient for establishing reasonable certainty, I can imagine my endeavors to affirm correctly about environmental facts as successful. Such a thing should be almost impossible to not imagine, since all it requires of me is that I merely do what it seems I do in my ‘normal’ pre-theoretic state. This is not to say that these environmental epistemic practices yield knowledge. Instead, the claim is that one cannot help but *seem* to see our preexisting epistemic practices as aptly successful. After all, our waking life is grounded in the *practice* of these principles, so it seems

unreasonable to claim that this illusion could be convincingly dispelled.⁵

2.3 FROM INDUBITABLE TO SENSE TO INDUBITABLE POSSIBILITY

Now that we have our main thesis, we need a way to make it usable for argumentation. For now, it has an ‘it seems’ clause put in front of it that seems to preclude us from making any conclusions from it. What we need is something like the following thesis:

THE PHENOMENOLOGICAL PRINCIPLE

If ‘it seems that p ’ is a phenomenological fact, then it is indubitably possible that p .

Using this thesis, we could take the above thesis about affirmational endeavoring and use in modal-style reasoning. But, we will need to justify this notion of indubitable impossibility. As a first rough shot, when p is indubitably possible this means that p is possible qua the *indubitable* basic phenomenal ways of being and the resulting fundamental epistemic scheme. In other words, p ’s possibility is an epistemic given that must hold when considered prior to any empirical investigation and thus has a very strong a priori status. The implication of this thesis would be that any phenomenological fact would sketch out possible ways the phenomenal world could be, and that any philosophical view that contradicts such possibilities for any and all ‘indubitably possible’ worlds would have to be false. This thesis, like the previous affirmational endeavor thesis, is hard to make sense of and justify precisely because it is supposed to be a primal notion within our fundamental first-person conceptual scheme. But we can make this thesis more plausible if we think of it as providing a bridge from what we call ‘indubitable sense’ to ‘indubitable epistemic possibility’. The argument — if it is to be called an argument — goes something like this: a phenomenological fact makes indubitable sense, and anything that makes indubitable sense we must consider indubitably possible, i.e., must indubitably structure our epistemic scheme. Before we can do this though, we need a firmer understanding of the two sides of the bridge: indubitable sense and indubitable possibility.

Let us first address the notion of *sense*. Here I mean the term in a way that is subtly different from Frege’s meanings. To start, ‘sense’ as a notion need not be as loaded as Frege’s has gotten over the years. When I use the word ‘sense’ I intend it to invoke the meaning of sense as that what things which are *nonsense* lack. As such, sense should not be a controversial notion at all; it is simply the background assumption of our cognitive and phenomenal life. To say then that something has an indubitable sense then is to say that it this thing must have sense for *me* and that no argument could lead me to doubt this fact; I cannot *help* but see it as making sense. We need not buy into the more potentially doubtful ontological implications that Frege’s view of sense implies. For Frege, things having sense means that they have an “Art des Gegebenseins” (1892),

⁵As said above, a fuller argument is not within the scope of this paper. Such an argument would likely be made by adapting late-Wittgensteinian material.

a way of being given, but Frege means this to say that they have an ‘objective’ sense, one that stands independent of our grasping it. Frege pushes this line so far that he feels the need to place senses in a ‘drittes Reich’ – a third realm, one that exists separately from the realm of our first person sense impressions and the realm of things. Frege seems to have thought that something being given *to me* was not good enough, and therefore he created a realm to where the ‘objective’ ways of being given could reside. But it would seem to me that by focusing on the notion of *indubitable* sense we avoid the stipulation of a third realm altogether. There is nothing unduly ‘unobjective’ to a sense that presents itself to me as long as its so presenting is a *necessary* fact of *my* experiential way of being, in other words, that it is a phenomenological fact. There is also nothing problematic about understanding phenomenological facts as having sense; formulating a fact about how my experience is structured in a certain way to say just to say this: that experience is structured – i.e. given – to me in certain way.

On the other side of the bridge we have indubitable possibility as the base structure of our epistemic scheme. Here, we can contrast this notion with Chalmers’s very similar notion of primary/epistemic possibility. Take the notion of an “epistemic space” (Chalmers 2011, 61), the space of epistemic scenarios (epistemically possible worlds) that a subject is justified to entertain as possible, given their knowledge right now. Here we should take the phrase ‘possible world’ as just being an investigative term; because my knowledge is incomplete, the *actual* world could be many *possible ways* with respect to my current knowledge base. Chalmers then introduces the notion of “deep epistemic possibility”, that is, “ways things might be, prior to what anyone knows.” (2011, 62) Yet here, just like Frege, Chalmers’s notion seems to involve an unnecessary imposition of some sort of ‘objectivity’; *prima facie*, speaking of ways things might be prior to what *anyone* knows would seem to speak to a kind of modality that is almost decoupled from any particular knower. This is perhaps why there has been some resistance⁶ to Chalmers’s Modal Rationalism, that (minimally) “ideal primary positive conceivability” (2002, § 6) entails deep epistemic possibility (also called “primary possibility” (2002, § 13)). The idea here is that as long as some state of affairs is both (a) imaginable and (b) able to withstand ideal rational reflection such a state of affairs must be deeply epistemically possible, i.e., it must be possible in a way prior to what anyone knows. This raises the usual question about this type of strong a priori knowledge: how come we as limited empirical beings are able to access these ‘ideally reflected facts’?

Though the particulars of Chalmers’s formulation might very well give a satisfactory answer to this problem, we can avoid the issue altogether. Since indubitable possibilities are (as I claim) the result of phenomenological investigation, we need not have them function prior to *any* particular set of beliefs; it is enough that they function prior to *our* particular sets of beliefs, that is, *my* belief set and *your* (that means *you*, the reader’s) set of beliefs. I propose then that we see indubitable possibility as the following:

INDUBITABLE POSSIBILITY

p is indubitably possible (for you) if and only if *p* is epistemically possible with respect to a successfully produced ‘Cartesian’/phenomenological reduction of

⁶See Vaidya 2017 for on this debate.

my current set of beliefs.

This is to say, if when I go through a good phenomenological investigation where I try to assume as little as possible about my beliefs, and I find I cannot rule out p 's possibility with respect to that reduced set of beliefs, then I call p is indubitably possible. This notion of possibility is wider than Chalmers's. Mathematical falsehoods like ' $2+2=5$ ' are deeply epistemically impossible because they would not withstand ideal rational reflection, yet they would be indubitably possible in part because of how closely indubitable possibility cleaves to indubitable sense. While it might be indubitably impossible to grasp a proof of ' $2+2=5$ ' (at least for how we normally construe the meaning of ' $2+2=5$ '), there certainly is nothing incoherent in the phenomenology of considering ' $2+2=5$ ' as true; our doing 'reductio' proofs is licensed by this phenomenological sensibility.

Given these two notions, we can finally start to construct our bridge from phenomenological facts to indubitable possibility. To state the 'argument' more specifically now, phenomenological facts of the type 'it seems that p ' are such that ' p ' must have indubitable sense, and that our making sense of p involves understanding it as presenting an indubitably possible scenario for the world to be in. A valid inferential argument is still likely impossible, but we certainly can make this bridge more plausible by providing close analogical examples.

Take the following example: suppose a person at the doctor describes a symptom they have as follows, 'when I lie down, it is as if there is a heavy weight on my chest.' Now, this description is perfectly sensible; it should certainly help the doctor come to some diagnostic conclusion. Moreover, it would seem that this description's sense is parasitic on the description that follows 'it is as if' having a sense of its own. In other words, the patient cannot rightfully be said able to describe themselves of having a sensation which is like having a weight on their chest if they did not, to some degree, understand what it would feel like to have a weight on their chest, either through being able to coherently imagine such an experience or by having such an experience before. That is, we understand the meaning of 'there being a heavy weight on my chest' by understanding what it would *actually* mean for there to be a weight on my chest, i.e., as an indubitably possible scenario, one that need not ever occur in your life but whose occurrence would be perfectly consistent with your indubitable way of being in the world.

By analogy, we can say that the 'it is as if' description functions similarly to 'it seems that p ' characterizations. Take an example from perception. Imagine for instance that I see (or minimally seem to see) a tree in my garden. In reporting this state I might say 'It seems to be that I see a tree in my garden', which is a general instance of 'it seems to be that p '. This characterization is almost always readily available and easy to make, yet it is one that relies on an implicit common-sense understanding of what it would mean for me to actually see a tree in my garden. In general, one could say that descriptions of the type 'it seems to be that p ' require that ' p ' have indubitable sense; otherwise, how is one giving a description of a certain state using the 'it seems' locution without understanding what the experience we have seems like? After all, if someone describes an experience they had on a hallucinogenic drug as 'it seemed like there was an eagle on my shoulder' it would make no sense for them to have succeeded

in that description and not have an understanding of what it means for there to be an eagle on their shoulder. Moreover, seemingly understanding a proposition p — its having indubitable sense — involves seeing p as possible in the sense that its p being the case would not radically break my phenomenal way of being in the world.

This idea should extend to *any* well-formed phenomenological presentation, in particular to our phenomenological thesis. Insofar as we view the synthesized whole given by experience presenting — as McDowell puts it — “a having in view” i.e. ‘having in display’ of the world, we should view our *characterizations* of this experience as possibly being affirmations about this display, and thus about the world (2009, 260). Again, we rely on an implicit common-sense understanding. I *seem* to live in a world that I can try — in fact, *succeed* — to say things about. This is not to say that our actual state of affairs might be different; we might — as the familiar example goes — be but a brain in a vat. But this does not mean that the idea of us succeeding in trying to say something about the world is deeply incoherent. I seem to be doing it all the time, and there is certainly nothing incoherent in my description of this seeming. Moreover, for me to rightfully make sense of my day-to-day life, I *must* see this seeming practice as being possible even when taking on the most stringent of epistemic constraints.

2.4 RECAPPING THE ARGUMENT

Given the above arguments, the bridge from indubitable sensibility to indubitable possibility should be plausible. Recall that in § 1.3 we set out the rough structure of the negative arguments I was going to set out against givenism and Sosa’s coherentism. We can now update this rough structure using the notions we have developed in this chapter:

1. I seem to be able to succeed in an endeavor to make apt affirmations about my environment. (phenomenological analysis)
2. If ‘it seems that p ’ is a phenomenological fact, then it is indubitably possible that p . (phenomenological principle)
3. It is indubitably possible to succeed in an endeavor to make apt affirmations about my environment. (1, 2)
4. If view X is true, it is not possible to succeed in an endeavor to make apt affirmations about my environment.
5. Therefore, view X is false (3, 4)

Our next tasks then will be setting out givenism and Sosa’s coherentism, and developing the arguments that show how they make successful judgmental endeavoring impossible. We will first tackle givenism, but before we do so, it is worth reflecting some more on what this phenomenological analysis amounts to.

To make judgmental endeavoring plausible, we appealed our *seeming* epistemic practices. Another way of thinking about this is in terms of Wittgenstein's work post the *Tractatus Logico-Philosophicus*. Throughout his later work, Wittgenstein describes our current epistemic practices in a way that seems uninterested in the typical 'veridicality' questions, analyzing judgment even irrespective of whether it is right or wrong, as merely simply looking at its features. (§ 132) Certain judgment procedures exist for things like calculations, and they seem to work reasonably well to a point where we feel like we can *stop* at some point, certain enough that we need not continue to check the result (§ 212). We seem to simply go about things in our lives in such a way that seems to inescapably seem to make sense, in fact, to such a degree that certain guiding beliefs form "hinges" for our thoughts (§ 343), and in fact almost certainly includes something like the procedures we do to settle our mind.

To be fair, I am ironing over Wittgenstein's tendency to interpret epistemic states in terms of purely nonepistemic action, but what I intend to convey is how Wittgenstein tries to make our basic naive way of being in world seem like it already makes *sense*. Wittgenstein of course pushes this to a quietist limit, but we need not go so far. What we can however do is incorporate elements of his thought into the transcendental argument against epistemic common kind justification. In fact, in the next chapter I intend to do just that.

3 THE MYTH OF THE GIVEN PARADOX

The goal in this chapter will be to set out the view of ‘givenism’—a particular internalist foundationalism—and show how this view makes judgmental endeavors impossible. We will first explain epistemic givenism within the context of judgmental endeavors and foundationalism. After that, we will provide the transcendental argument against it.

McDowell himself has tackled the historical context of the type of view we intend to critique in the afterword of *Mind and World*. First of all, givenism as I have been calling it takes its name from Wilfrid Sellars’ Myth of the Given, which he critiques in *Empiricism and the Philosophy of Mind* by specifically starting with a critique “sense-datum theories” ([1956] 1997, 13). It is also similar to the verificationist theories that Quine critiques in *Two Dogmas of Empiricism* (1951). And it also is related to the private language argument that Wittgenstein critiques in the *Philosophical Investigations* ([1953] 2001). Typically understood, what these critiqued views have in common is that they assume there is some set of beliefs, impressions, or what have you (the given), on which our epistemic system is founded such a way that justification originates or enters the epistemic system through them. We, however, are dealing the *internal* version of this view. Givenism as such is not a statement about the structure of ex ante justification, but rather about the decision procedures we go through when are endeavoring to aptly affirm. That is to say, we are not trying to figure out whether the structure of propositional justification originates in set objects, but rather whether the *internal* i.e. rational procedure we go through for committing ourselves to an affirmation has a foundationalist structure. On its own, the idea that our internal justification originates from certain parts of our epistemic system rather than others is not necessarily bad. However, when we combine it with the common kind justification thesis we get a view that is more problematic. Now, justification must both originate or enter solely through some part of the epistemic system, yet justification must be the same for indistinguishable states. This guarantees that the justification involved cannot enter into the system and instead must find origin in the system; if it did not, the justification could be different even if it was internally indistinguishable.

This then gives us a picture of what the process for a judgmental endeavor consists in. Say I want to settle my mind on whether p . To determine its ‘truth’, I must check its justificatory relations to the foundational bits of given. I go through this process until I am satisfied that the proposition is rationally justified, and then, ex post justifiably assert (say) p . Given that judgmental endeavoring is an internally phenomenologically present endeavor in which I affirm for *myself*, the bits of given must appear to me as experiences, e.g., visual experience. Since I am now constrained to my internal appearances, I cannot ‘step outside my body’ to check the reliability of my experiences, as it were. In fact, that would not be the point, since that would set further conditions for judging well that stood *outside* internally indistinguishable states. Thus, I can only talk about epistemic practices as relativized to a particular set of experiences that I use to build all my other beliefs. Let us call this type of judgment endeavors and respective judgments, ‘schjudgmental endeavor’ and ‘schjudgment’. Likewise, insofar as my beliefs are well built with respect to those experiences, I can get something like a best case for knowledge which we could call ‘schknowledge’. As such, schknowledge/schjudgments would

be something like ‘knowledge/judgments about the reality that my experiences seem to present’.

To put the dialectic more clearly: *givenism* is the position that the structure of internal epistemic justification is foundationalist *and* that the common kind justification thesis is true. Since the judgmental endeavor is an internalistically construed process, this requires to engage in *schjudgmental* endeavors to successfully ‘schjudge’ and hopefully store the resulting belief as ‘schknowledge.’ The question then is whether successful schjudgments about the environment amount to successful *judgments* about the environment.

But it turns out that this type of epistemic system for coming to ‘schjudgments’ makes the type of environmental judgmental endeavor that the we supposedly thought was phenomenologically required impossible. Here, I have chosen to use a version of Wittgenstein’s private language argument¹. Roughly put, the argument goes as follows:

1. I function within a schjudgment system, where all rules for characterizing experiences must be derived internally from foundational experiences themselves.
2. But there are no criteria for correct and incorrect characterizations of my foundational experiences besides the ones I arbitrarily hold myself.
3. So, there is no epistemic process I can go through to enable success in an endeavor to affirm.
4. So, it would be indubitably impossible to successfully endeavor to say something about my environment.
5. But the above contradicts the phenomenological thesis, so I must not function within a schjudgment system.

To start, when I am trying to figure out the foundational rules for judging a particular matter, I have to first of all figure out what the foundational rules *are*. Yet, to figure out what the foundational beliefs I have, I must avail myself of other rules for characterizing them, which makes them not foundational. My giving up on getting the world right and constraining myself to my senses, means that I have to give up any rules for the characterization of my foundational experiences.

We can illustrate with an example, adapted in part from passages by Wittgenstein (§§ 258–270). Suppose I have an experience *S*, and I try to characterize it in some fashion by saying ‘that is a Φ ’, where Φ could be a tiger, a tree, a chair, a building, etc. Now, who is to say that an experience presents an apple or a Pegasus? Well, if the experience functions as a foundational given within our epistemic system, then the answer is nothing and no one. Any appeal to a linguistic community or to a reality outside my

¹McDowell ([1994] 1996, 18–19) indicates that Wittgenstein’s private language argument is a good alternative way of construing his problems with ‘given’ based accounts of reasoning. My specific understanding of Wittgenstein’s argument though is mainly due to Kenny ([1973] 2006).

experiences would be to abandon the view taken. I could appeal to another experience, say a memory of how I identified the experience before, but then nothing justifies my characterization of that memory as an experience either, since in comparing my memory of an experience S to some experience S I would need rules to determine that I did that comparison correctly. Otherwise, there is nothing to say that the experience I am having now is not an S experience but a T experience. Even appealing to a *different* foundational bit of given Q would not work: either (a) that other bit of given justifies my characterization of S as a Φ , in which case S is not foundational, or (b) S justifies my characterization of Q and Q justifies my characterization of S , in which case the view is not foundational but coherentist.

Now someone might argue that these foundational experiences do not need to be characterized when I go through a judgment, and that in reality they are already characterized when I start. But it misses the point; the problem is not that I am confronted with a given experience which I subsequently characterize; the problem is that the experience of *being characterized* requires a framework to exist *logically* and not necessarily *temporally* prior to the experience itself. Take again the experience S , which now instead of having to characterize I find it already as being characterized as (say) an apple—that is, the experience comes *theory-laden*. Yet, since in the case where the experience does not come theory-laden and I would need a theory to recognize it, it stands to reason that the theory *itself* is the actual foundational given that *logically* justifies the experience as being an apple, which would mean that the apple experience itself is *not* a truly foundational given.

According to another objection, one might argue that the whole point of this sort of foundationalist view is that external circumstances determine whether I have characterized my experiences correctly, in which case the internal recognition of the experience can proceed as it may. If I characterize a particular experience as presenting a Pegasus when it was caused by an apple, then I am wrong not because of internal but external reasons. However, this objection shifts the correctness criteria. What makes my characterization of an experience now correct is whether I report it as representing an apple when it was appropriately caused, Why then should I require that I identify the experience *qua* experience correctly? As Wittgenstein says:

Imagine a person whose memory could not retain *what* the word ‘pain’ meant — but nevertheless used the word in a way fitting in with the usual symptoms and presuppositions of pain” — in short he uses it as well as we all do. Here I should like to say: a wheel that can be turned though nothing else moves with it, is not part of the mechanism. ([1953] 2001, PI, I, § 270)

Which is to say, what does it matter if my characterization of my experiences is internally inconsistent when my correctness criteria are external? If I constantly identified wildly differing experiences as being of an apple but only did so when appropriately caused by apples, then obviously my view does not care about identifying my experiences; all that mattered was my differential responsiveness to the appropriate causes, which is not a non-belief given. Again, we have without realizing it abandoned the internal founda-

tionalist view.²

Given that there is no procedure for assessing the foundational experiences, it should likewise mean there is no procedure *in general* for assessing environmental claims, since the justification for any environmental proposition would have to start internally from some foundational bit of given. And from here the Kantian skeptical conclusion is not far off: If there is no internal justification process we can go through to commit to a judgment, then we are precluded from any sort of successful endeavor. In fact, we are precluded not because of any particular feature of the world, but because of a piece of conceptual reasoning that supposedly functions prior to any particular state of affairs we might consider. The argument, I would claim, functions exactly at the same level of indubitable possibility.

As such, we are led into contradiction. We must either give up the phenomenological theses from § 2, or reject judgment systems. Given the two available options, it makes more sense to give up the latter. This is not to say that there is not some type of ‘given’, but the given is not that of piecemeal foundational experiences. Instead, the given is more of the type that Wittgenstein alludes to when he talks about hinges. We are not saddled with foundational beliefs, but rather with entire overlapping and interconnecting *systems* of evaluating claims that we operate from within. We operate on a coherentist “Neurathian” (McDowell [1994] 1996, 81) model of judgment, according to which we overhaul different parts of our judgment system, in the fashion “which a sailor overhauls his ship while it is afloat”, using one system to overhaul another. It is also this reason why that while it might seem that phenomenological seemings function as bits of given in the same way as this argument criticizes, this is not the case. The point here is not that phenomenological seemings are unarguable in the same way that we imagine experiences are in judgment systems; rather, that what makes these phenomenological seemings contentful in part is their imagined conditions of satisfaction, i.e., that the phenomenological seeming *already* contains the insight. If we had already understood the meaning of our phenomenological seemings we would have seen that we could not have understood it as phenomenological without its larger context within our epistemic system. Our basing ourselves on our phenomenology was not basing ourselves on a bit of non-empirical given; it was simply affirming our already existing epistemic framework.

However, as said above before, although the internal coherentist framework contains some valuable insights, this does not mean that it meshes well with the epistemic common kind thesis. Now that we have treated the givenist position, our next task will be to take internal epistemic common kind coherentism. However, since we are critiquing Sosa’s version of this position specifically, our next task will be explain it in detail.

²Modesto Gómez-Alonso has pointed out to me that the Wittgenstein passage might also be interpreted as such that the meanings of ‘wheel’ and ‘mechanism’ are reversed. In other words, if ‘sole judgments of external behavior’ (wheel) can be interpreted (turned) without the correct corresponding internal thoughts (movements), then these are not part of full judgments of thought (mechanism). Prima facie to me though, the passage seems compatible enough with both interpretations. So, since a more precise exegesis of Wittgenstein’s point goes beyond this paper, we will for convenience’s sake assume that the first interpretation is correct.

4 THE BASICS OF SOSA'S THEORY

The purpose of this chapter will be to give a quick rundown of Ernest Sosa's epistemic theory of judgment that he presents in his book *Judgment and Agency* (2015). In epistemology literature, Sosa is mainly thought of as an advocate for reliabilism and thus is identified more as a philosopher with externalist sensibilities. The view that Sosa presents in his book *Judgment and Agency* however is far more subtle than this typical roughshod characterization gives credit. For starters, Sosa connects his view to the Descartes' view in the *Mediations*, a book that one might think gives a paradigmatically *internal* view of justification. In fact, this is what Sosa's view offers: an *internalist* coherentist view of rational justification, coupled with an *externalist* and reliabilist view of ex ante justification. Though somewhat different, Sosa likely thinks that this position is for the most part compatible with the one he held in his celebrated paper *The Raft and the Pyramid* (1980). There, Sosa rejected a coherentism that sees the epistemic value of justification consisting *solely* in coherence (§ 9). Sosa's view of coherence here is that it has some small value, i.e., just enough to force us to act coherently in the first place, but that its *true* value comes into being when it enables us to improve our reliability at forming beliefs/making judgments. In fact, in *Knowing Full Well* (2011) Sosa sketches out a transcendental argument according to which even that 'small' internal justification finds its origin in our potential success if/when our causal mechanism of perception and memory are not radically mistaken. For the purposes of this paper though, these distinctions will matter little. All that matters is that the 'small' justification is there, even in a world where we are radically mistaken.

I will first sketch out Sosa's earlier views as they apply to belief and knowledge, since these are easier to understand. Then, I will transfer these views to judgment with the help of his work in *Judgment and Agency*. Finally, I will parse out Sosa's particular coherentist view, showing how it fits with the previously set out internalist 'process' rationality. With these materials in place, we will be able to successfully set out the Kantian skeptical argument in the next chapter.

4.1 SOSA ON BELIEF AND KNOWLEDGE

To understand Sosa's view in *Judgment and Agency*, it is helpful to first consider his simpler theory concerning belief and knowledge that he presents in *Knowing Full Well* (2011). There, Sosa sees knowledge as a particular type of belief that exhibits some 'correct' relation to one or more epistemic competences. To be knowledge, Sosa thinks that a belief must satisfy what he calls "AAA structure" (15): it must be (a) "accurate", i.e. true; (b) adroit (skilled), manifesting (an) epistemic competence(s); and (c) it must be apt, i.e. it must be accurate because it is adroit. This framework Sosa classifies as "a kind of performance normativity". To make this notion clearer, Sosa draws an analogy between these kinds of beliefs and the performance of an archer aiming at a target. Just as the archer can shoot at a target competently with an arrow and succeed to hit the target through their competence, so can an epistemic agent aim and hit the truth through a competently formed belief. The difference would merely be what particular "domain"

(6) we use to evaluate that particular performance. With the archer, the domain might be something like the rules of the Olympic Games by which the archery performance in which I compete is governed, and with belief-formation the domain would be that of “epistemic normativity” (Sosa 2011, 1), i.e. merely a different kind of performance normativity. On a first shot then, knowledge would simply be apt belief.

However, Sosa wants to distinguish different levels of aptness: a performance could be apt yet fail to be “meta-apt”, “because the agent handles risk poorly, either by taking too much or by taking too little” risk (7). Suppose instead of an archer we have a hunter, who not only has to shoot skillfully but also has to “pick [their] shots so as to secure a reasonable chance of success” (5). It would be good qua the domain of being a hunter to sometimes not shoot at a particular target — i.e., “forebear” — if the shot one would then take is particularly risky, for instance if the hunter has only one arrow left and no clean shot at their target. Even better is if when the hunter does successfully shoot an arrow, that the shot’s “success flows also from [their] target-selecting, shot-picking competences” (12). Thus, Sosa distinguishes three kinds of knowledge: “animal knowledge” (12–13), which is mere apt belief; “reflective knowledge”, which is an aptly formed belief about one’s aptness, and “human knowledge” or “fully apt belief” (92), which is when “reflective knowledge helps to guide the first-order belief so that it is apt” (13). It is when one has this human knowledge that Sosa says that one “knows full well”.

4.2 SOSA ON JUDGMENT AND AFFIRMATION

However, in his recent book *Judgment and Agency* (2015), Sosa has shifted his attention from the comparatively ‘passive’ notion of fully apt belief to the more active notion of apt judgment. Judgment for Sosa “is affirmation in the endeavor to affirm with apt correctness” (66). By ‘affirmation’ Sosa means something like the inner mental analogue of speaking the sentence out loud, i.e. assertion (66n2). So, when I say to myself ‘there is a tree in front of me’, then I have affirmed that there is a tree in front of me. Sosa then defines “judgmental belief” as “the disposition to so judge”. Sosa does this in part because he thinks that he can show more easily that judgment itself (unlike belief) inherently demands to be guided by reflective ‘meta-apt’ considerations (85). A judger does not only endeavor to affirm correctly, but to affirm aptly, because judgment as an act incorporates this demand in and of itself. To make sense of this demand, Sosa connects the act of judgment with his particular interpretation of Descartes: when Descartes attempted “to raise his first-order judgments, up to the scientia level” (250) (as he did with his clear and distinct ideas), he did so to avoid “error ... not just falsity”. Descartes demanded that when I affirmed something for myself that I make sure that my affirmation was safe, which for Sosa is another way of saying that I aimed to have the aptness of my affirmation come to pass through a “second-order awareness of my competence to so affirm” (80). That is to say, I affirmed because I was aware of the suitability of the conditions of my epistemic competence, and thus affirmed aptly (because I manifested my epistemic competence), meta-aptly (because I was aware of the suitability of the conditions of my epistemic competence), and fully aptly. Thus, taking Descartes’s demand for safeness of my affirmations seriously — which is the same as knowing that the manifestation of my

epistemic competence is suitable — is the same as recognizing the legitimacy of the act of judgment and its inherent demand for aptness.

Furthermore, Sosa expands his AAA framework. Before, all we had to go on for determining the aptness of some act was whether or not it was 'accurate because adroit', i.e. whether it in some way *manifested* some sort of competence, where what makes something a competence had to do with reliability within suitable conditions (2011, 7–10). In *Judgment and Agency*, however, he expands this notion to provide what he sees as a *metaphysical analysis* of competences. Now, when we talk about performances being apt, be they perceptions, actions, or affirmations, we mean to say that these actions were manifestation of a type of competence that has the same kind of metaphysical weight that a disposition like fragility or solubility might have (2015, 27). In fact, competences should be seen “as being a special case of dispositions” (2015, 24), and in fact are quite similar to them. Both are dispositions to behave in a certain way given that they are effected causally “in the right way” (29). Competences specifically though follow an SSS profile: an innermost competence (Skill) can manifest itself when one is in the right Situation and in the right personal Shape (awake, fit, etc.).

This theoretical framework has some interesting results for the classic barn facade example. The barn facade seems to provide two conflicting intuitions. On the one hand, something seems to go right: the person looking at the barn is seeing an actual barn. On the other hand, something is going wrong: their belief that there is a barn could have very easily been false, or perhaps the justification provided is not strong enough. But Sosa's hierarchical conception of knowledge/judgment is able to explain away these conflicting intuitions: On the one hand, if someone looking at a barn in barn facade county affirms for themselves that there is a barn in front of them, then that affirmation is apt insofar as they manifest their epistemic competence to visually distinguish barn-like things in an environment. On the other hand, their affirmation might have easily been false given their environment, and thus — insofar as they were endeavoring to judge whether there is a barn in front of them — their resulting judgment was inapt because they did not affirm with an awareness of the suitability of the (SSS) conditions of their epistemic competence. In other words, Sosa argues that the barn facade example was confusing because we had not properly cleaved the metaphysical concepts of knowledge and judgment at the joints. Once we understand the differences in requirements between aptness of judgment (second-order awareness) and aptness of affirmation, the confusion disappears. Yet, at the same time we can explain our discomfort with mere apt affirmation as springing from our desire to secure the aptness of our affirmation *for ourselves*, that is, as resulting from our desire to *judge* for ourselves.

4.3 SOSA'S COHERENTISM

Given, however, that Sosa's view sees judging as involving a (potentially implicit) second-order awareness, Sosa has to face the familiar threat of Cartesian-style vicious circularity, especially with respect to judgments of experience. For Sosa, experiential seemings follow the same SSS profile; instead of manifesting epistemic competences, they manifest “the relevant competence of the subject's perceptual systems.” (2015, 20) Here, Sosa

draws another metaphysical distinction, setting out “three sorts” of “states or events” that human lives have: (a) sufferings — pains or itches, for example — or mere doings, such as reflex actions; and performances of two sorts: (b) functionings (functionally assessable states); and (c) endeavors (with a freely determined aim). (Sosa 2015, 192)

Although Sosa allows various gradations in this distinction, he is fairly clear in specifying that perceptual seemings are “functional” (92) and thus are not subject to “direct agential control”; they merely provide “attractions to represent that *p*” (93). This in contrast to judgments, which are the result of endeavors with the freely determined aim to aptly affirm. But if that is the case, then it would seem impossible to aptly judge any proposition that originates from my perceptual system. How can I become second-order aware of the suitability of the conditions for the functioning of my visual epistemic competence if I have to use that very same visual competence to assure my suitability? It would seem that I would never be able to aptly judge any propositions that involve deliverances from my senses because I would never be able to aptly judge the reliability of my senses simpliciter. This in turn would seem to threaten our ability to aptly judge all a posteriori propositions.

Sosa’s response is as follows: reflective justification cannot possibly affirm the reliability of our senses simpliciter because we are not epistemically responsible for getting such justification; the problem simply disappears. Reflective justification, for Sosa, functions on a “web” model (2009, 239): we justify our beliefs and affirmations by resting them on other beliefs, not by stepping outside the web itself. As long as the web is “attached to the world through the causal mechanisms of perception and memory” (240), then our efforts to make our beliefs cohere and deductively follow will improve our knowledge of our environment further. But even *if* our web of belief is fundamentally disjointed from an outer environment, then improved coherence will *still* to some degree add an extra epistemic value (242). Thus, we neither need to nor are responsible for worrying about whether our web of beliefs hook up to some sort of outer reality. As long as we tend well to our particular web of beliefs — i.e., as long as we function as good reflective epistemic actors — we are justified, albeit in a small sense, to continue without having to worry whether our assertions and beliefs are grounded.

Moreover, Sosa does not just seem to endorse this coherentism out of practical considerations, but also for quasi-Kantian skeptical reasons. We can see this because while motivating his view of justification, Sosa references Davidson’s coherence theory of truth (2009, 109–132). Davidson does seem to work within a quasi-Kantian skeptical framework: Reflective justification is not limited to the system of beliefs because of tragic limitations; in fact, it would be “unintelligible” to “request for a ground or source of justification of another ilk” (Davidson 2001a, 141). In other words, Davidson dismisses this sort of view not because it is epistemically inconvenient, but because it is incoherent. As we saw in § 3, Sosa and Davidson are legitimized within the Kantian skeptical framework to draw this conclusion. The kind of internal justification that Sosa and Davidson argue against is indeed deeply incoherent because it leads us to deny deep undeniable phenomenal theses.

However, the question is whether the coherentism Sosa endorses will work either. Sosa is forced to endorse the epistemic common kind thesis, since (as was said above) the

functional seemings of senses provide only ‘attractions to represent’ that *p*. In fact, the justification that we get ‘in the small sense’ from tending to our beliefs *relies* on this common kind assumption. Yes, outer environment might be disjointed from our web of beliefs, but that situation is indistinguishable from it not being so. As such, our web of beliefs gain the small sense of justification because it *could be* going right.

Thus, our next task in the following chapter will be to develop the Kantian skeptical argument against Sosa’s view. In the end, while Sosa’s view has much going for it due to its systematic scope, it will not be able to deal with a different Kantian skeptical argument because of its fundamental disconnect from the ‘sources’ of epistemic normativity.

5 THE COHERENTIST PARADOX

In previous chapters we set out the various basic phenomenological theses involving judgmental endeavoring and had pointed to two views of internal rational justification which would govern this act: givenism vs. internal coherentism. We treated givenism in chapter 3; now our task is to develop an argument against internal coherentism, specifically against a particularly sophisticated version endorsed by Ernest Sosa that I set out in chapter 4. That is what will be the topic of this chapter.

Just as with the argument against givenism, the argument I will set out will be a type of rule-following argument. Wittgenstein however tends to draw on amorphous language games that have some sort of dislocated sense of authority. However, when we are trying to understand the *world*, we are not dealing with a dislocated community but rather a strict set of top down rules as set out by the way the world is. The analogy therefore will be rather that of playing a tournament game, in this case the *World Series of Poker* (WSOP). The idea here will be that intentionally internally functioning in a coherent fashion—playing a game that makes sense—cannot possibly amount to acting with the intention to get the *actual* world right—playing in the actual world series of poker.

The task in section § 5.1 will be to give a first pass of this argument and show how it might apply. However, this first pass will be not fully able to deal with Davidsonian triangulation arguments that might neuter the WSOP argument. Thus, in § 5.2 I will first give a sketch of the type triangulation argument views that might hurt the WSOP argument. Then, I will sufficiently modify the WSOP argument such that it is immune to these triangulation worries. Having done so, I will have shown persuasively that Sosa's coherentism—and any other internal coherentism for that matter—suffers Kantian skeptical problems and therefore must be rejected.

5.1 A FIRST PASS

Just as with the argument against givenism, we assume our phenomenal and phenomenological principles: the common kind justification thesis, the judgmental endeavor phenomenological thesis, and phenomenological principle, including any other phenomenological considerations we raised in chapter 2. Now, however, we endorse Sosa's framework instead of a schjudgment one. Thus, we take Sosa's notion of freely-made affirmation (of which judgment is a sub-type) and look at the cases where we use those endeavors to say something about our (seeming) environment. Since Sosa's notion of endeavors necessarily involves seeing them as different from functional competences, any freely made affirmational endeavors about our seeming environment must involve taking a (possibly implicit) second-order perspective on our first-order functional competences. It is this particular first/second-order structure that we will use to generate our paradox. To make the dialectic clear: Sosa's first/second-order structure is particular way to give cash-value to the process that successful judgmental endeavors must follow to be considered *apt*; the argument will aim to establish that this way of giving cash-value will violate the phenomenological theses and thus must be rejected.

The coherentist paradox goes as follows:

1. I function within Sosa's judgment system, so any system of beliefs I have is connected to my environment through the making of affirmational endeavors where one takes a (potentially implicit) second-order perspective on first-order functional competence. Specifically, one endeavors to determine (to some degree of certainty) whether the SSS conditions of the functional competence have been met, and if they are met affirming that p , where the 'attraction to represent that p ' is provided by the selfsame functional competence.
2. In the endeavors described above, one must either determine directly whether the manifestation conditions are met, or use principles that already exist within the epistemic framework.
3. The former is impossible, since e.g. determining whether my eyes are working *right now* without using memories of past perceptions would involve, as it were, 'stepping outside of my own body'.
4. Yet doing only the latter means that I am not aiming to affirm only when I have determined that the SSS conditions of first-order competences are met, but in actuality aim to affirm in such a way that my epistemic framework *seemingly* does not contradict itself, that is, irrespective of the actual SSS conditions.
5. But, since the first-order functional competences are what (possibly) undergird my system of beliefs to my environment, and I do not aim in my affirmational endeavors concerning those functional competences to fulfill the SSS conditions, then I do not aim to codify my actual environment in my belief system, but instead aim to codify a *seemingly coherent* notion of environment in my belief system.
6. Yet, if I aim to codify a *seemingly coherent* environment, then I am not aiming to codify my *actual* environment in my belief system, and thus cannot possibly succeed in an endeavor to affirm correctly about my environment.
7. Thus, either I try determine the SSS conditions of my functional competences directly, which is impossible, or I stick to a (seemingly) coherent understanding of my SSS conditions, in which case I cannot endeavor about the environment itself anymore. In either case, it is indubitably impossible to succeed in an endeavor to affirm correctly about my environment.
8. Thus, since that contradicts the phenomenological thesis, it must be that I do not function within Sosa's judgment system.

Notice here that the argument as such traffics on a disconnect between the norms of higher order judgment and the norms of the first order competences. The first order competences have clear norms: the SSS conditions. The higher order competences involved in judgment must surveil these norms, but to do so forces us into an uncomfortable

choice: either (3) one attempts to fulfill the norms directly, which leads to the familiar impossible problem of ‘stepping outside one’s own body’; or (6) one sticks to what one *thinks* are the SSS/norms of assertion, which then makes the endeavor not an endeavor to get the *actual* norms/SSS. It is this latter point that will require some explanation. It is not necessarily clear at this point in the dialectic why a seemingly coherent notion of environment should be so different from an actual environment. So, to make this point clear, let us develop a language game style argument, or, to be more precise, a *competition* game style argument.

Take the following example of playing a game:

‘WORLD SERIES OF POKER’ (WSOP)

A player is trying to participate in the ‘World Series Poker’ tournament. Unfortunately though, they do not know the rules of the game, nor do they (for some reason) have access to rule books or other reliable sources of information. Thus, they adopt a set of rules which they think are the rules of the poker tournament, and decide to play nonetheless, adjusting their rules as necessary by observing the reactions of the players and referees and by seeking coherence in their rules overall.

This example is similar to Sosa’s judgment system in several respects. Just as with the poker tournament, the rules of the ‘getting the environment right’ game — i.e., the SSS conditions of affirmations about our environment — are presumably determinate and not subject to change. If we are to succeed in our endeavor we must conform to them. Moreover, we also do not have access to the rules beforehand — i.e., that is, we cannot step outside our bodies — and thus must rely on the ‘reactions’ — sense perception — to correct our rules over time. *Prima facie* then, this example would seem to count in favor of Sosa’s view, because the position one finds oneself in WSOP is certainly not hopeless and it is clearly quite possible to learn a game even with this limited set of information. However, the example is disanalogous in ways that are unfavorable for Sosa’s judgment system, because a judge in that system is arguably far worse off than one in WSOP. For one, in the WSOP example the poker player presumably has a previous understanding of what games are and how they are played (they have winners and losers, they tend to involve taking turns, they can involve betting or stakes, etc.) and so they will likely have a fairly strong preexisting theory to use and revise when they sit down at the poker table. But with judgment there need not be any guarantee that our beliefs be mostly right about the environment because there is no previous experience with judging to draw from. Our career as judges is our first so we have no idea whether we are mostly right.

More importantly though, since we do not have reliable access to the secure norms, our intention to maintain a coherent notion of our game means that in our intent we are not playing the game of WSOP. Instead, we are playing the game of keeping our game coherent which would only be *incidentally* similar to WSOP. To make this clear, imagine the following situation: Suppose a player *did* have available to them, if they so chose, an avenue through which to get the rules of WSOP, either directly or through some

secure intermediate. Suppose now however that their beginning working theory of the game that they started off with was massively wrong. That is to say, this person theory of the game was so divergent that it caused them to always misunderstand the players' reactions in the game and draw incorrect conclusions from them; coherence on its own would not correct their understanding of the game and would instead steer them further away. Suppose now however that a referee ran up to the player and told them that they were *solely* mistaken in their understanding of the game, and that they had to completely start their working theory from scratch. But if now the player would *ignore* the referee and keep to their own theory, they would have revealed what their *actual* intentions were: they were never interested *at all* in playing WSOP, they were only interested in keeping their own game, whatever it might be, coherent. Moreover, it would seem that availability of the norms would not change the nature of the intention. If the 'working theory' game can so radically differ from WSOP, then an intention to keep the 'working theory' game coherent would have to not be an intention to learn and play WSOP.

Relating the analogy back to the issue of judgment, this would suggest the following:

If it is (indubitably) possible for a judger to have a belief system that is massively wrong, then it would not be possible to succeed in an endeavor to affirm correctly about my actual environment.

Because, just as with the WSOP example, any seemingly coherent notion of environment we have is only *incidentally* similar to any actual environment in which we live.

Prima facie, the antecedent of the conditional seems easily fulfilled, by, for instance, the supposed indubitable possibility of a brain in vat state of affairs. Yet, there is material in the 'semantic externalist' literature that might make one think otherwise. There are many variants of the argument¹, but the most interesting one for our purposes is the one that could be advanced along Davidsonian lines, and which is perhaps of the type that Sosa might endorse². According to Davidson, "belief is in its nature veridical" (2001a, 146), which is to say that in taking our beliefs as beliefs we must understand them as *necessarily* mostly massively right and thus not massively wrong. If Davidson is right, then the argument simply does not go through. Our next task, then, will be to analyze Davidson's argument.

5.2 DEFUSING THE TRIANGULATION ARGUMENT

The reason Davidson believes that belief is in its nature veridical has to do with how Davidson sees knowledge of the self (a), others (b), and the environment (c) stand in relation to each other. As Davidson sees it, we interpret others through the process he calls triangulation: I (a) know what you (b) are saying, because I correlate your actions with occurrences in the environment (c) and map those correlations back into my own lan-

¹See Brueckner 2016 for a treatment of Putnam's version of the argument. (though I personally find that variation unpersuasive for reasons that, for those who are familiar with the argument, will become clearer below.)

²See Sosa 2009 for more on this.

guage (a). The ‘triangulation’ here then is the use of a shared environment to ground the meaning of the other expressions. I effect this interpretation by using the Quinean principles of charity: correspondence and coherence. “The Principle of Coherence prompts the interpreter to discover a degree of logical consistency in the thought of the speaker; the Principle of Correspondence prompts the interpreter to take the speaker to be responding to the same features of the world that he (the interpreter) would be responding to under similar circumstances.” (2001b, 211) What distinguishes Davidson’s argument here is that he reverses the explanatory arrow: it is not just that I make sense of others by correlating their actions with meanings in their language, it is that I can only make sense of *myself* as having my actions correlated with meanings in other people’s language. All three varieties of knowledge go together hand in hand, and so you cannot have knowledge of self without knowledge of others and the environment.

Though a precise exegesis of Davidson’s point is beyond the scope of this paper, we can make sense of it by comparing it — as Davidson does himself (2001b, 209-210) — to Wittgenstein’s private language argument. According to a more traditional construal of Wittgenstein’s point, we cannot have infallibly known inner experience without thereby creating a ‘private language’, something that Wittgenstein claims is incoherent. Thus, any judgments about our inner experience can only be made with some background knowledge of the norms of outer language. Davidson simply extends this argument through triangulation; we can only form a theory of what another says through correlating their expressions with happenings in the environment. As such, knowledge of myself commits me to knowledge of others, and knowledge of others commits me to knowledge of my environment.

In sum, let us express Davidson’s claim as follows:

I am a believer because my utterances and actions could be made sense of by some other.

Or, to put the point more carefully:

THE TRIANGULATION THESIS

For any person *S*, if *S* is a believer, then an overwhelming number of *S*’s beliefs and utterances must be able to be rendered as sensible and true by some (possibly only hypothetical) interpreter existing in some environment.

Under this construal of the argument, what is still required before we can get that our environmental beliefs are largely correct (and thus rebut the WSOP argument from § 5) is that we have a belief system of some sort. *Prima facie*, this seems like an indubitably necessary³ fact; denying it would involve denying that I have beliefs, which seemingly is fact of my phenomenology. Yet, I want to claim that the only relevant fact that our phenomenology here can provide is that it *seems* that I have beliefs, and not necessarily that I actually have beliefs. According to our phenomenological principle, this would then indeed entail that having beliefs is indubitably possible, but not that it is indubitably

³Recall § 2.

necessary. The triangulation *thesis* can be true even though the overall triangulation *argument* for massive correctness does not work.

To make theoretical space for this position, let us postulate a notion that is phenomenologically similar to belief called *pseudo-belief*. Pseudo-beliefs would share with actual beliefs the quality of *seeming* to be a belief, yet would be incomprehensible by any other believers in any linguistic community. Correspondingly, there would be a distinction between actual believers and *pseudo-believers*. A believer can be comprehended by a community, while a pseudo-believer cannot, though both share the quality of internally seeming to have beliefs. Accordingly, we could say the following:

BELIEVERS AND PSEUDO-BELIEVERS

A believer mostly has beliefs. A pseudo-believer mostly (if not wholly) has pseudo-beliefs.

What I want to claim now is the following:

PHENOMENOLOGICAL UNDERDETERMINATION OF PSEUDO-BELIEVERS (UP)

There are states of affairs of being a pseudo-believer that are phenomenologically indistinguishable from the state of being an actual believer.⁴

This is not to say that the states are necessarily identical phenomenally speaking, in fact, it is likely that they are not. It is certainly plausible that a believer could distinguish between a belief and a pseudo-belief qua their phenomenal features. But it is impossible for us to determine qua phenomenal features whether we *actually* are believers instead of pseudo-believers. All we can be sure of is that we *seem* to be believers which at the most only guarantees the indubitable possibility of our being believers, not its indubitable *necessity*. Considered as such, phenomenology could not commit us to seeing ourselves *necessarily* as believers, and thus could not commit us to our beliefs about self, others, and the environment; radical Kantian skeptical doubt would still be possible.

As a start to advancing UP, let us make the notion of pseudo-belief more clear by providing closely analogous real world examples. As mentioned above, Davidson thinks

⁴This idea is similar to one Martin (2002, 2004) expresses, although his understanding of phenomenology is different than the one used in this paper. Martin claims that the only thing common between a veridical and a non-veridical sense-perception is the property of “being indiscriminable from a veridical perception” (69). However, what Martin calls the “phenomenological character of perceptual experience” (Martin 2002, 402) seems to line up far more closely to what one might call the ‘metaphysical nature’ of experience. Thus, the claim of indistinguishability is not one about the ‘phenomenology’ of non-veridical sense-perception in Martin’s terminology. It is neutral claim, because it does not pretend to give an account of non-veridical perception, giving only ‘epistemic’ criteria. Of course, according to this paper it is precisely these types of ‘epistemic’ criteria that phenomenology – definitionally speaking – is all about. The same point is made, only using different terminology.

that we interpret other people's behavior by mapping it to concepts and notions within our own language; we apply the principles of charity and coherence to do so, i.e., we *a priori* assume that the person we are interpreting *is* making sense and is trying to say something. Yet, we can still assume that interpretation is affected in such a fashion while allowing that sometimes such interpretation *fails*. In fact, sometimes an inability to understand a speaker is not due to the interpreter but due to a speaker, because the speaker simply is not making *any* sense. Students sometimes think they understand a concept when they do not, politicians think they sometimes understand a particular piece of policy when they do not; certainly, the world is replete with examples of people who *think* they are making sense when in fact they do not. Now of course one person's inability to make sense of someone's words or actions does not mean that a speaker's words or actions do not have any sense, and it is not a conclusion that anyone should draw particularly easily. If, for instance, I do not understand the explanation of some political situation that an economist provides me this not to say immediately that the economist is merely 'babbling'; I could simply be lacking the requisite expertise or conceptual vocabulary. But if I cannot make sense of them, and neither can my neighbor, nor my friends, nor any experts I consult, then at some point it does become empirically justified to draw the conclusion that there was nothing to make sense of to begin with.

Admittedly, just because someone fails to say what they mean does not mean that what they end up saying does not make *any* sense whatsoever. For instance, if a student does not explain a math concept well this is not to say that they have not understood anything, or that what they said could not on some level be made sense of at all. But this is not to detract from the possibility of someone failing to completely say *anything*. In fact, there seem to exist cases in the real world of people doing just that. One potential example might be found in cases of 'receptive' or 'Wernicke's aphasia. Those with this particular condition are able to vocalize, but frequently misspeak and misuse words and phrases, sometimes simultaneously and to such a degree that the sentences they make are completely unintelligible (Edwards 2005, 11). These errors do not seem intentional, nor easily predicted or corrected; the unintelligibility of 'non-words' in severe cases of aphasia is sometimes so severe that they "may be a some kind of random output from the phonological store" (sic) (44). What is more, patients with Wernicke's aphasia are "generally unaware of their production errors" (18). This lack of awareness encompasses also their ability to participate in conversation as well, as they frequently *seem* to think that they are answering questions when they in actuality are not. As such, Wernicke's aphasiacs would seem to provide a powerful example of pseudo-affirmation, that is, a behavior or action whose meaning cannot be made sense of using any particular language precisely because there is no pattern to be had, or at least none which the speaker could remotely be interpreted as intending. It would seem then that some behaviors cannot be made sense of because they do not make sense at all.

In turn, the Wernicke's aphasia example lends more legitimacy to the notion of being a full *pseudo-believer*. Given that we could potentially not distinguish at any particular point in time *prima facie* between an affirmation and a pseudo-affirmation (suppose we woke up tomorrow as a Wernicke's aphasiac), what is to say that we were making sense — i.e., could be made sense of — to begin with? Perhaps we had *always* been pseudo-

affirming, and had only thought we were affirming. If so, interpreters could not make sense of our actions and utterances and, as such, could not ascribe beliefs to us. What is being described then is the following modified ‘Brain in a Vat’-type scenario:

NONSENSE IN A VAT (NIV)

A mad scientist has hooked me (or just my brain) up to a machine that keeps me alive, but replaces the senses of all my thoughts with nonsense. As such, my thoughts are neither right nor wrong; they are simply meaningless.

Suppose I had always been an NIV, I would not have right or wrong thoughts; they would lack sense. Of course it seems *prima facie* that my thoughts have sense, yet, as we saw with the Wernicke’s aphasia example, this does not at all necessarily mean that those thoughts actually do have sense. The seeming perception of sense underdetermines thoughts actually having sense.

To finally get UP, we need to make plausible that being an NIV is phenomenologically indistinguishable from not being an NIV. We had already allowed that any particular pseudo-belief is *prima facie* phenomenally indistinguishable from an actual belief, but had left open the possibility that this need not mean that they are phenomenologically indistinguishable. Presumably, this would mean that though a particular pseudo-belief state φ might be phenomenally similar enough to the belief state χ such that these are indistinguishable on a first glance, further investigation qua their *phenomenal* properties would yield some sort of difference. If someone is a believer, i.e., mostly has belief states, then this would seem quite plausible. It is not uncommon for a person to have a confused or unclear thought that – though on first face seems to make sense – they later realize is confused, a conclusion sometimes drawn in part because of the ‘feel’ of a particular presentation. After all, dreaming experiences feel different than waking experiences, so why should the according experiences involving pseudo-beliefs not feel different than those of holding a regular belief? However, if we are pseudo-believers, then it seems far less likely that we could distinguish a pseudo-belief from a belief qua its phenomenal character. After all, if I only have pseudo-beliefs, I might not have any particular idea anymore of what an actual belief is. Comparing the phenomenal character of any particular pseudo-belief to any other pseudo-belief would not seem to show anything out of the ordinary. From the perspective of a pseudo-believer, the phenomenal characteristics unique to pseudo-states would be the *normal* character of phenomenal states. Thus, the various states of affairs of being a believer are not just *prima facie* phenomenally indistinguishable from corresponding NIV scenarios, but also *phenomenologically* indistinguishable; no investigation that constrains itself to the logic of the phenomenal presentations of states could reveal a difference between a believer state and a corresponding NIV state.

5.3 A SECOND PASS

Let us now recall the original thrust behind the coherentist paradox. A judger functioning under Sosa’s view must choose between two options: either one tries to determine

the SSS conditions of one's functional competences directly, or they stick to a coherent understanding of the SSS conditions. The former I said was impossible, since we could never get past the 'it seems' statements directly. The latter was supposed to be incoherent since trying to make my *seeming* environment coherent needed not bear any similarity to targeting my *actual* environment, which in this case was specifically supposed to become known through the low-level SSS functional competences. So, to be clear, what we needed to show was that these two acts (a) cohering the seeming environment and (b) judgmentally endeavoring about the actual environment, had sufficiently different goals and norms such that the success of (a) need never amount to the success of (b). The argument had been that the seeming environment could be massively wrong, so targeting the seeming environment was not akin to targeting the actual environment. However, the worry was that my actual environment would always cohere with my seeming environment due to semantic externalist considerations, and thus that our beliefs about a seeming environment could never be massively wrong. Put more carefully, the SSS conditions we were targeting could always be about whatever environment our beliefs were causally effective in; so, in the case of a brain in a vat, since my beliefs are causally effective about virtual trees, my SSS competences would include. e.g., a reliable ability to see trees.

But now, having motivated UP, we *can* doubt our commitment to the massive truth of our beliefs by simply doubting that we have belief systems in the first place. Davidson's triangulation argument does not go through, at least qua the resources that Sosa's position provides, because it now *does* seem to be indubitably possible for a judger to have a belief system that is massively wrong, albeit not wrong because it is incorrect, but wrong because it is nonsensical. Thus, step 6 of the rule-following argument set out in § 5 still goes through. When I aim to get my preexisting *seemingly coherent* notion environment, I am allowing that the notion of environment — in fact, my very notion of world, including self and others — could be massively incoherent.

It is important here to realize how massively large the disconnect between the seeming world and the actual world can be. Suppose someone was in an NIV situation, but they were still able to navigate the environment competently, albeit in an unconscious fashion somewhat in the fashion of blindsight, or perhaps in a prelinguistic fashion similar to how a child or very young non-verbal animal does. In such a case, one would still be given attractions to represent that *p* through e.g. ones visual powers, but one would not be able to make *sense* of the attractions to represent as such on any reflective level. We could properly be said not to be able to distinguish whether we ourselves could be in such a situation. Yet if we were, we would gain *no* benefit from coherence, in fact, our attempts at seeming coherence would only drive us further into the 'recesses' of our minds. In fact, even in the *good* case we have no idea whether our self-correction is helping us along. When I decide to reject the suitability for judging that there was a red light on, was I acting in such a way that would reinforce my meta-apt beliefs about my first-order competences, or was that the first step towards losing my mind? Perhaps there could be such a thing as *gradually* slipping into an NIV state, where a demon deceiver makes it *seem* like a particular belief does not fit in my belief system, and thus I think I act coherently and throw it out. Again, seeming coherence would do us no good, and

in fact might be leading us further down the incorrect path.

Thus, Sosa's coherentism fails. Sosa wanted to make an 'internal' coherentism that had an 'external' reliabilism in the background to secure access to what we judge coherently about. But the norm of coherentism need not be the same norm of the world at all; the WSOP example as thickened up with the NIV example showed that. Instead, the Sosian internal coherentists show themselves as completely unconcerned with their actual environment and only concerned, at the most, with appearances. The point can be put banally like this: if you try to get the *seeming* world right, you can never succeed in getting *actual* world right.

6 SETTING OUT EPISTEMIC SECURE REALISM

6.1 DIAGNOSING THE PROBLEM

In the previous chapters we developed two rule-following paradoxes. These two paradoxes were supposed to exhaust the space of plausible possible views that one could adopt when one held on to various phenomenological theses and the ‘Common Kind Justification Thesis.’ The ‘given’ rule-following paradox made it seem impossible to endorse an internal rational foundationalist view. Yet, the coherentist paradox made it seem impossible to do what looks like the opposite: to only derive justification from other beliefs. We might cynically conclude that the notion of judgmental endeavors is deeply incoherent. This would involve denying the phenomenological principles that were set out in § 2. Thus, I advocate that we let go of the common kind justification thesis. Once we do so there is new dialectical breathing room for a new position: epistemic secure realism.

What seems to cause the coherentist paradox for Sosa is the iterative structure of judgment. The paradox does not arise for first order competences because those competences are ‘ontologically integrated’, by which I mean that the success and SSS conditions for these competences are almost wholly determined by the ontological matters of fact themselves directly. However, the ‘surveying’ faculty involved in raising first order affirmations to second or higher orders is not ontologically integrated. The argument in § 5 exploits this fact by pointing out discrepancies between the critical domain of the world/first order competences and the critical epistemic domain of judgment. It was this epistemic domain that turned out to be incoherent, precisely because it was disconnected from the ontological matters of fact. The ‘obvious’ solution would be to modify our notion of judgmental endeavor such that it is similarly ontologically integrated. Instead of having multiple first order competences that one surveys and coheres together using higher order cognitive operations, we have only one *global* competence that functions in all contexts¹. As we shall see, exercising this competence correctly will mean coming into cognitive contact with the facts themselves, giving us something akin to an intellectual ‘seeing.’

That said, we will need to do some preliminary groundwork before we can fully spell out this new notion of judgment. First, we need to give a new notion of agency. Typical Sosian first order competences are ‘functional’ and as such do not require any sort of notion of full-blooded agentiality. We need to give a different cash-value to phenomenal considerations of conscious agentiality, albeit one that does not involve the aforementioned iterative Cartesian structure. This will be the task of § 6.2. Instead of a Cartesian notion agentially spelled out in terms of internal reflection, I will advocate for a ‘fully conscious’ view of agentiality à la William James. Second, in § 6.3 I will inject the Jamesian view of agentiality back into Sosa’s view of judgment. This will result in a new type of judgmental endeavor: the *fully conscious judgmental endeavor*. Third, I develop out the notion of critical domain that this new judgmental endeavor has when we apply it to the

¹We will address what this means for the ‘original’ first order competences (e.g. sight and memory) in § 6.5.

world, again, taking care to port over as much of Sosa's existing philosophical framework as I can. Last, I put all the results from the previous section to finally formulate epistemic secure realism.

6.2 CONSCIOUS ATTENTION AND THE WILL

Like Sosa, I do not wish to give up the notion of "conscious rationality" (2016), but if we are to keep this, it has to be spelled out in a different way. The most promising avenue to me seems to be the phenomenal notions of attention and will that James provides in *The Principles of Psychology* (1890a, 1890b).

To start, let us focus on attention. Attention, for James, radically breaks the notion of experience as merely just passive, but as active, and volitional. As James puts it, in Kantian terms, "Attention, implying a degree of reactive spontaneity, would seem to break through the circle of pure receptivity which constitutes experience" (1890a, 402). This is to say that while we are *given* an initial sphere of potential experience, we have some agency in picking out what in that experience we attend to, and in fact, that this attention in some ways distinctly patterns our experience. James himself says that "*my experience is what I agree to attend to*", which is perhaps somewhat too strong. At the very least though, we can agree that the phenomenon of attention patterns our experience of the world incredibly strongly.

James goes further, saying that "effort of attention is thus the essential phenomenon of the will" (1890b, 562). This is because James sees attention as the mechanism through which the will asserts itself in the world (564). By focusing on an idea and letting it fill one's mind, one is able to move one's body in a deliberate fashion. But even deliberating and deciding as acts consist of the exercise of the will on thought (528). The will, then, i.e. the exercise of one's *own* desires, seems to manifest itself as attention. When one deliberates, decides, and when one generally deliberately attends to parts of one's experience, these acts are all exercises of the will. (528–535)

Moreover, there would seem to be a connection between consciousness and attention. In his article "Attention is Rational-Access Consciousness" (2010), Declan Smithies argues for the idea that "attention is a distinctive mode of consciousness" which he calls "rational-access consciousness" (32). Like James, Smithies appeals to phenomenal considerations to begin the *prima facie* defense of attention, calling it "*sui generis*" (4), but he goes further by defending this phenomenological notion from more functionalist views on attention. Smithies allows that there might be other functional understandings of attention that might be seen as active in cases such as blindsight, but argues strongly that the conscious phenomenological notion of attention is the only kind that truly makes its contents "*rationaly accessible* in the sense that it is accessible to the subject as a reason that justifies the subject in forming a belief or performing an action" (22). As Smithies says, "blindsighted subjects do not spontaneously form beliefs or perform actions upon objects in the blind field, but claim to be merely guessing or acting randomly." (23). Moreover, while "Unconscious information in the belief system is rationally accessible to some degree for use in the control of action, reasoning and verbal report", "it becomes rationally accessible to a much higher degree when it is made accessible as the content

of an explicit judgment, which fully engages one's attention." (31–32)

This suggests that attention — that is, the most fully conscious and phenomenologically present version of attention — is the *essential* manifestation of our will and rational agency. It not only makes available to us the contents of our (sub)consciousness, but also allows us to steer our actions and thoughts. Moreover, insofar as reasoning and judging are volitional acts just as much as deliberating and deciding are, these too are exercises of the will as brought about through attention. This then implies a new division of phenomenal states that is similar to Sosa's endeavor/functional competence/sufferings-doings distinction. We have endeavors which are controlled by the will and executed through conscious attention; we have conscious experiences, whose contents are made available for consideration in endeavors and which are made available to the will through conscious attention; and we have all the other phenomenal states which do not fit this particular mold.

Taken all together, we get the following:

THE WILL AND CONSCIOUS ATTENTION PHENOMENAL SCHEMA

Rational agency (the will) is paradigmatically manifested in conscious attention. Conscious attention enables us to steer and execute conscious/willful endeavors, and enables us to make available for ourselves at least a certain set of passively occurring experiences. This then divides our phenomenal lives into three parts: (1) conscious/willful endeavors; (2) passive experience made consciously accessible; (3) all other phenomenal events that do not fit into the previous categories.

6.3 REASONING AND FULLY CONSCIOUS JUDGMENT

We can use this new structure to give an alternative account of the 'second-order awareness' that Sosa says is present in judgment. In Sosa's view, second-order awareness need not be "conscious nor temporally prior" (Sosa 2015, 79n20) for it to be active. But Sosa does allow for so-called "subconscious affirmations" (52n25) as a part of larger reasoning chains that lead to conscious affirmations, such as when one does a quick mental calculation. This kind of reasoning however can still be brought to the surface for use in "pondering or deliberation" (66n3); in fact, the retrievability of such subconscious affirmations (which are *steps* of affirmational reasoning) is a necessary feature of intentional 'endeavors' for Sosa. This then allows us to suggest that the *primary* notion of judgment is not its subconscious version but its *fully conscious and attentive* version.

A way to make sense of this suggestion is by seeing it as part of the desire for *scientia* judgments. In the Cartesian story, we examine our beliefs partly because we realize how much of our intellectual framework we hold unconsciously. The judgmental endeavor, as such, insofar as we try our best to get at the truth aptly, partly involves making conscious what was first *unconscious*, i.e. bringing it to our attention. Once we have brought a previously unconscious belief to our attention we can then revise it based on our own criteria, i.e. as in accordance with our own will. Moreover, if I want to make sure that my belief is correct, i.e. that it is apt, then surely I would want to make sure that my agency as a knower is fully manifest in my judgments and judgmental endeavors. Af-

ter all, wherever my agency is *not* manifested, i.e. where my judgment is less conscious, I could be in error.² Thus, judgmental endeavors and the judgments that result from them are ideally made willfully in the fullest sense: fully conscious and attentively.

But given this understanding of the Cartesian demand, we are also inclined to see *reasoning* as ideally realized in a fully conscious and attentional fashion. As deliberation is to deciding, so reasoning certainly is to judgment. Ideally, one would assume that the one flows naturally into the other; I reason so as to settle my mind whether *p*. Following Frege and Kant, it is tempting to see the process of reasoning as consisting in purely step-wise inferential reasoning. But such a view might bring vicious circularity back into the picture. After all, if reasoning can only avail itself of inferential processes, then it would be stuck within a coherentist framework. Instead, we can now appeal to a far more ‘primitive’ notion of reasoning by linking it to the willful attentive structuring of thought. As a result, judging that *p* could be seen as ideally involving *fully consciously and attentively* deciding whether *p* is true. Likewise, reasoning about whether *p* then is simply the conscious thoughts one went through so as to determine whether *p*. Whether this thought was inferentially structured or not need not preclude this string of conscious thought from being considered reasoning; just as long as it was *this* string of willful *conscious* thought that led one to the act of judgment. Put in Sosa’s terminology then, we could say that reasoning as such constitutes the “second-order awareness” (2015, 79) that is supposed to guide apt judgment. This ‘awareness’ then would simply be the conscious thoughts — i.e. the reasoning — that led one to affirm as one did.

Taken all together, we advance the following thesis:

THE J-ENDEAVOR THESIS

Call the fully conscious, attentive, and intentionally executed endeavor to affirm with apt correctness a ‘fully conscious and attentive judgmental endeavor’, or J-endeavor. Such an endeavor is evaluable normatively with respect to some critical domain. When an agent succeeds in such an endeavor in a fully conscious and attentive fashion (a K-judgment), then the success is attributable to them and their conscious reasoning processes.

We have not yet fully explicated J-endeavors in volitional terms. Specifically, we still have not given an account of how to understand the volitional endeavoring towards *aptness*. Before we give this account, however, let us examine how the notion of J-endeavoring interacts with Sosa’s notion of the critical domain.

6.4 J-ENDEAVORS AND THE CRITICAL DOMAIN

As Sosa says, any particular domain “will have its proper objective(s), and its correlated standards of success, and of proper competence” (2015, 126). Take for instance the domain of playing basketball: what partly constitutes that domain is what counts as a point, from where one can shoot, how many times one can dribble, etc. In other words, the do-

²Or, to put this demand in Kantian terms à la *What is Enlightenment?* ([1784] 2008), wherever I do not think for myself, I put myself in nonage of others.

main determines what does and does not count as success by imposing a framework for evaluating performances. These ‘norms’ that act on an endeavor or performance can find their normative source in various things: in “human convention” (102n9), in human “nature and needs” (99n8), in “evolutionary teleology”, and so forth. They can be rigid, in the case of a baseball game, or looser and communally determined, such as with the critical domain of ‘polite’ conversation (101–102). In the case of world-oriented judgment however, we had already decided in § 6.1 that the world itself would fix what makes such judgments both correct and apt.

Sosa seems to not say too much about the norms and their function in fully volitionally executed performances/endeavors, but we can quite easily draw these implications ourselves. *First*, notice that I must have some idea of what the norms are on my endeavor beforehand if I am even to endeavor (be it fully volitionally or otherwise) to succeed. Suppose that I were standing on a basketball court during a game, but that no one had told me the rules of basketball beforehand. If I had not even any idea of what the rules of basketball are, I could not even start to do a whole range of endeavors, such as scoring a point, dribble, defend my basket, and so forth. *Second*, notice that to actually succeed at some act (separate from my intentional endeavor to succeed) I must act in accordance with the actual rules governing that act. If I endeavor to score a point in basketball and actually succeed, I succeed *in part* because of how my physical movements sufficiently matched the rules for scoring a goal; I did not dribble too many times, I shot from within the boundaries of the court, and so forth. Success at an endeavor would then *necessarily* involve fulfilling the norms of acting on the endeavor. *Third*, if my conscious intention to succeed at some act is to successfully bring about that act (i.e. that my endeavor is apt), then not only do I intend to fulfill the norms of the critical domain of my endeavor and have my endeavor fulfill the actual norms, but I must have some awareness of the actual norms on my endeavor, either directly or possibly through some secure mediate.

To make these points clearer, take the following example: suppose that I am an Olympic gymnast intentionally endeavoring to get a gold medal in gymnastics. Insofar as I am engaged in this project, I want to ensure that I am actually fulfilling the Olympic committee’s criteria for a good performance. To make sure that my routine best fulfills the committee’s rules, I ask my trainer what they are. However, unbeknownst to me, I do not actually ask my trainer but an impostor posing as them. This person then told me — completely without any justification on their part — that the Olympic committee rewards extra points to routines with 720 degree spins. However, it turns out that these rules were recently added by the Olympic committee. So, when next I perform my gymnastics routine with a 720 degree spin I perform it aptly and receive the gold medal. Thus, I endeavored with an idea of fulfilling correct norms governing the act I am doing and I actually fulfilled the correct norms of my act, yet I did not actually fulfill the norms because of my endeavor to fulfill them. My first-level apt success was not creditable to me as an intentional agent, at least with respect to my efforts to consciously fulfill the Olympic committee’s rules.

To put this in terms that Sosa might endorse, I can have this second order awareness of my capabilities only if I have an awareness of the actual norms of my act that at the very least is given to me through some secure and safe mediate. If it is not safe (i.e. if

the intentional act is executed with respect to norms that could have easily been false), then I clearly do not act fully aptly and thus do not aptly succeed in my endeavor to act aptly. All together, we get the following:

THE 'NORMATIVE DOMAIN' (ND) THESIS

To *fully* consciously, intentionally, and aptly endeavor to perform an act with respect to some normative domain, we must (at some point)³ have a conscious awareness of the norms of the critical domain of my act (most likely through some safe mediate relating to the source of the norm).

6.5 THE WORLD-ORIENTED JUDGMENTAL COMPETENCE

We now have enough resources to give a fuller account of the nature of world-oriented judgment. As we saw in § 4, Sosa sees judgment as steered by a second-order awareness that considered whether the conditions were right for the application of a first-order awareness. Now, however, we have not a second-order awareness, but an attentive fully conscious awareness. This awareness is able to (seemingly) volitionally attend to what seems to be the contiguous experience of a world, and is also able to steer thoughts in deliberation in such a way that they terminate in decision. Given this (seeming) ability to volitionally structure thoughts, I propose that world-oriented J-endeavoring consists of deliberately structuring one's thoughts in such a way that the resulting judgment *necessarily* is true *because* of how one structured one's thoughts. In other words, world-oriented J-endeavoring consists in endeavoring to structure one's thoughts in such a way that they function as an ontologically integrated competence to get at the things themselves. Put more formally, we get the following:

WORLD-ORIENTED K-JUDGMENT THESIS

When I J-endeavor to settle my mind on a world-oriented matter and succeed in this endeavor (i.e., produce a K-judgment), then my reasoning and subsequent judgment functioned as the exercise of an ontologically integrated power to think true thoughts about the world. Moreover, if I possess such a power, then it should be considered subsumed under my general power to volitionally steer my thoughts.

This thesis, taken together with all the others from § 6.1-6.4, implies epistemological disjunctivism of perception as Pritchard (2015) defines it:

³Here, the clause 'at some point' is important to avoid an overly intellectualist point of view that would certainly not work for more physical acts like playing football, making a cabinet, and so forth. We need not hold the norms of playing football continually in our conscious awareness while we are actually playing football. In fact, it is doubtful whether one needs to have a fully conscious awareness *at all* to play football well. However, if one *does* endeavor aptly in a fully-conscious and volitional fashion — even if such an approach might be overly intellectualistic for the task at hand — then one must have an awareness of the norms acting on the act one is endeavoring to perform.

PRITCHARD'S EPISTEMOLOGICAL DISJUNCTIVISM

In paradigm cases of perceptual knowledge, the knowledge in question enjoys a rational support which is both factive and reflectively accessible.

We know that as consciously executed endeavors, the agent of a J-endeavor must have some sort of awareness of the norms that govern this endeavor. The norms that govern a world-oriented J-endeavor are supplied by the world itself. In fact, the world simply *is* the metaphysical embodiment of the norms. Given the phenomenal schema from § 6.2, it would mean that for J-endeavors to be possible we would need a (potentially mediated) awareness of these norms — that is to say, of the world itself — as provided through consciously accessible passive experience. Given that this awareness must be of the *actual* norms (albeit through some safe and secure mediate), this means that the awareness must necessarily be factive. Thus, for judgment to be possible, it must be possible for us to have perceptual states that provide the actual norms — that is, are factive — and which are accessible through the function of conscious attention.

However, to make sense of the epistemological disjunctivism embedded in this thesis while avoiding the paradoxes set out above, we must also endorse a ‘naive realist’ metaphysical disjunctive view of perception that directly contradicts the Common Kind Justification Thesis. If perception functions as an intermediary from our consciously accessible experience to the actual ontological matters of fact themselves, then we would have a similar second/first order structure that we originally had rejected in Sosa’s view. In the case of the trainer example in § 6.4, I can derive good reasons for trusting them to relay the rules of the Olympic committee faithfully; I could be well-situated to judge their expertise. With perceptual states though, I have none of this background information, and so the coherentist paradox plays out all over again. Necessarily, because if perception functions as an intermediary to the actual matters of fact, then there must be something do the mediating. This forces the mind into the surveilling role that it had for Sosa, which means that it cannot be ‘ontologically integrated’, since its functioning as a competence is now dependent on whether the intermediate perceptive faculty is functioning well. Thus, I argue that we endorse the following:

PHENOMENAL DIRECT REALISM

It is possible for a phenomenal experience to relate me to the world itself directly through its experiential content.⁴

⁴This position is roughly in line with the one that McDowell endorses in (2009) and (2013). At Charles Travis’s urging, McDowell has backed off the view that experience is propositional from top to bottom, because it would seem to mean that an educated perceiver would *literally* see more than an uneducated one, and, moreover, that the world itself exists in pure propositional form à la early Wittgenstein (the world is all that is the case). These would seem like overly paternalistic metaphysical assumptions to make. Thus, McDowell has modified his position to allow perception to be both contentful and of objects through the presentation of specific content. Travis (2016) has criticized this modified view as well, saying that this notion of content functions as an intervening intermediary for perception, but it seems unclear (to me) from Travis’s work why that is the case. McDowell’s

Moreover, to *completely* avoid any Kantian skeptical problems, we have to view the power of perception to see things directly as smoothly integrated into our fuller ontologically integrated power of the mind to think true thoughts about the world. Taken together as one fully holistic competence, we finally arrive at epistemic secure realism.

EPISTEMIC SECURE REALISM

Judgmental endeavoring is J-endeavoring, and when I successfully K-judge that *p*, I have gained secure cognitive contact with the fact that *p*.

Thus, though I do not get into *direct* cognitive contact with the fact that *p*, I have insured that whatever underlies the truth of *p* has impressed itself on my mind through secure mediates, and, in fact, that my endeavoring has helped ensure my confidence that these mediates were secure. Thus, while I might not know some fact of the past *directly*, full judgmental competence involves as part of it the use of my memory, and thus allows me to determine in such a way to use my memory such that I gain a secure connection to the facts of the past. (More will be said on this in chapter 7.)

Epistemic secure realism is a disjunctive thesis. To explain, take the following three phenomenally indistinguishable yet radically metaphysically different examples:

1. I am J-endeavoring to get the actual world right. However, a demon is radically deceiving me by feeding me false experiences, so I have no ontologically integrated competence to perceive. As such, I will never be in a position to judge that *p*.
2. I *seem* to think that I am endeavoring to judge whether *p*. However, a demon deceiver is subtly changing my thoughts without me noticing it. Thus, because I actually do not exhibit any control over my thoughts and resulting actions, I fail to actually *willfully* endeavor at all.
3. I am J-endeavoring to get the actual world right. The world is open to me through my senses. My thoughts succeed to be structured in such a way that the resulting affirmation must necessarily be true, in part because the way the world is *enables* me to possibly structure my thoughts just so. My J-endeavor succeeds and I produce a K-judgment about the world.

What is common to all these states is that it *seems* that I am J-endeavoring; however, in actuality, all these states are fundamentally metaphysically different. In state 1 I have no access to the things themselves through experience, and in state 2 I have no ability to think for myself. As such, in states 1 and 2 I could not be said to have the required ontologically integrated competence to judge successfully. Yet, in state 3, I do have the required ontologically integrated competence. State 3, besides the obvious phenome-

(2017) point (like so many of his) simply involves taking our existing linguistic practices at face value; of course we can see things, and of course our seeings have content. These two ideas need not be incompatible. But far more could be said about this issue than space constraints would allow.

nological similarity, is radically metaphysically different. My judgmental competence is not a competence that functions by coherently circulating my first-order competences; rather, it is the *full* volitional scope of my consciousness as integrated with the world itself. I am not reflecting on a competence; instead, my conscious volitional attentive structuring of my thoughts *is* my judgmental competence.

This view is immune to the coherentist and givenist skeptical paradoxes. Instead of viewing judgment as involving a reflective attitude towards a core epistemic skill, we see our attempt to judge aptly as involving our fully integrated cognitive capacities. As such, we simply do not take a reflective attitude towards our cognitive capacities as a whole, not just because this requirement would be onerous, but because it is incoherent. What I endeavor for in my judgments is to get at the world aptly through my own conscious thoughts, and not as a reflection on another ‘epistemic’ skill. Insofar as the judgment process involves reasoning so as to settle my mind, it would make no sense to take a meta-attitude to this process, since that meta-attitude in being part of conscious thought must necessarily already be part of our reasoning.

Despite the previous discussion, one might still be inclined to think that this metaphysical type of judgment act is simply impossible a priori, because it would never be possible to satisfactorily settle our mind without any sort of potential doubt. Given that we have allowed phenomenally indistinguishable yet radically metaphysically different situations, this would seem to legitimize this doubt. If two experiences are phenomenally indistinguishable, this would suggest that both these experiences have the same epistemic function and warrant. This would seem to revive the circularity problem. After all, if we, within the scope of our consciousness, attend to two phenomenally indistinguishable experiences, then it would seem that we are *still* confined within our own minds, and that the competence of judging is still a second-order one that does not have an inherent metaphysical seat in reality. But this move relies on confusing ‘seeming’ paradoxes of warrant with actual metaphysical problems with judgment. While there might be a ‘seeming’ reason for such a doubt, we should not confuse it with an ‘objective’ or even *epistemic* reason for doubt. Objective ‘doubt’ is only reasonable insofar as it is legitimized by the objective norms set out by the world itself, and ‘epistemic’ reasons for doubt are only legitimate insofar as they help us in the process of knowing, a process we have already stipulated in inherently metaphysical and thus objective terms. To illustrate, take as an example a world wherein we perceive the world directly and without failure, e.g. we have a perceptual power to see things as they are without failure. In such a world, there would be no illusions or hallucinations. So, if I see that there is a blue surface in front of me, then there must necessarily be a blue surface in front in me. Recall our discussion about fully conscious/attentive endeavors. From the perspective of the world, experience provides me a necessarily ‘safe’ contact with the world itself, plus, I have conscious access to this experience through the stream of experience itself. Thus, I am able to endeavor with an eye of what the norms are, and succeed in my endeavoring because the rules governing my endeavor are so simple; all I would need to do to aptly settle my mind on whether there is a blue surface in front of me is to trust my senses. Apt judgments would not only be possible, but also quite easy. Clearly, whatever ‘epistemic norms’ we would refer to so as to generate ‘doubt’ would not be justified in a world like

this; they would not help us get at the truth in this one, in fact they could only lead us further astray. But the *objective* norms as set out by the world itself are quite clear. We either meet or fail to meet them.

To link back to the phenomenological roots of the argument, my claim is that the act of world-oriented judgment only makes sense as the phenomenologically possible exercise of an ontologically integrated competence to think true thoughts about the world, one that must include our ontologically integrated perceptual competences. In such a phenomenologically possible world, this competence is simultaneously *interior* and *exterior*: it is an exercise of my conscious will, yet, because my will is integrated into the world itself, it is also an exercise of power that can make connections to exterior ontological states of affairs. Since the world, being the way it is, is the ultimate arbiter about our beliefs and judgments, I claimed that the only way we can make sense of our judgmental endeavors is as a competence that is deeply integrated into the world. Moreover, the only way to make sense of the success of our judgmental endeavor as justified by the world is to see such a success amounting to ensuring secure cognitive contact with the actual environment. In a way, epistemic secure realism amounts to just taking our phenomenologically prior epistemic practices seriously by seeing them as *contentful*, albeit in an indubitably sensible fashion; if the meaning of seeming to see that p is grounded in the idea of actually seeing that p , then our seeming to come to know that p should be grounded in the idea of *actually* coming to know that p . Finally, our idea of the seeming world is grounded in the idea of it being the *actual* world.

I have now set out epistemic secure realism. However, there are still some outstanding questions about the nature of intellectual seeing and lesser forms of justification. First, if consciously coming to know that p is like seeing, how can I then come to consciously cognize distant propositions like ‘there is an Eiffel Tower’? Second, if justification now consists in acting in such a way that our reasoning process functions as an exercise of an ontologically integrated power to think true thoughts about the world, how come people are justified for anything less than full cognitive contact with reality? As I will claim in chapter 7, the answers to these questions are related; the ontologically integrated competence is inherently *social* in nature.

7 ROUNDING OFF EPISTEMIC SECURE REALISM

In this chapter we will talk about various justificational factors that might plague the view of epistemic secure realism. In advocating for epistemic secure realism I have argued for what might seem an incredibly rigid view of justification, one that would seem to contradict all the various cases where justification and knowledge (for us, judgment) do not coincide. In fact, Gettier cases, the most famous type of thought experiment in epistemology, rely crucially on the idea of *insufficient* justification. Yet, if *successful* judgment is an exercise of a metaphysically integrated competence, what does that mean for unsuccessful judgment? We had said that in the case of an NIV as defined in chapter 5 we get no redeeming justification, but we remained silent as to whether that was the case in normal ‘Brain in a Vat’ cases. Moreover, there are other more mundane cases to consider. There is the question of *socially transmitted* judgment, where I judge whether the Eiffel Tower exists, for instance. Here, we could ask how a firing of my wordly judgmental competence could ever allow me to come to get into direct cognitive contact with something like that fact. Would it require me, for instance, to visit the Eiffel Tower myself if I want to get the correct kind of warrant? For that matter, there is the question as what to do when someone makes an ‘honest’ mistake while judging, one that one could not have possibly foreseen. Suppose for instance I judge that a friend of mine was in Paris because the New York Times (or your other favorite reliable periodical) reported them being there. However, it turns out that in this particular instance, the reporter fabricated the story for complex financial reasons. Surely, it would seem that I had judged in a way that was good *enough*. If not, why not, and if so, why did the metaphysical competence not guarantee my truth?

All in all, these questions amount to the following: if being rational is now an external matter, how do we account for what seem to be good instances of reasoning? Altogether, they form what we can call the *exculpatory* problem of justification. The task in this chapter to give a sketch of what an answer to this problem might be.

7.1 EVIL DEMONS AND VIRTUAL REALITIES

In Stewart Cohen’s paper “Justification and Truth” (1984), he develops what is now known as the new evil demon problem. Take two reasoners, *A* and *B*, who are both stuck in demon worlds which make their beliefs unreliable by fiat. *A* forms their beliefs based on the evidence in a way that we think a good knower might, but *B* uses guesswork, wishful thinking, and so forth to form their beliefs. Stewart here aims this thought experiment at reliabilism, since simple reliabilism would have to say that both *A* and *B* are as reliable, and therefore just as justified. Yet, the intuition is supposed to be that this cannot be the case. Surely, *A* is more reliable than *B*. If we accept the intuition, the challenge would seem to apply to epistemic secure realism as well. *Prima facie*, since both reasoners have no ability to come to know facts through their cognitive processes, both are equally as rational.

Yet, this is too quick. What such a response would miss is that rationality is not *world specific*. We have the intuition that *A* is more rational than *B* because we live in a world

where reasoning like *A* is the way we get to know facts through our own cognitive powers. Whether the reasoners now have those powers all depends on the metaphysical matters of fact of the demon world. The demon world could for instance be of a type where wishful thinking *is* reliable due to the very nature of the world, yet that that very competence was thwarted by a demon deceiver. Because the judgmental competence is worldly, a demon deceiver would be the perfect type of entity to *mask* it in the dispositional sense. One could in fact imagine a whole dizzying array of thought experiments, with demons, or unfriendly environments, masks, finks, and the like. Going through all the various options would go beyond the scope of this paper, and would require pulling out the full disposition and powers toolbox out of the metaphysics literature. What is important to realize for now though is that the intuition that *A* is more rational than *B* could be easily explained by the fact that *A* (presumably) is the type of rationality that one would have to have in a world that resembles *ours*. Seen as such, the intuition is just ‘metaphysical parochialism.’

This masking power approach also works for cases where we put someone into a particular disabling state in our *own* world. Take for example the case where we invent a virtual reality machine, say, tomorrow, and put someone in it. As we specified, the judgmental competence power has subsumed under it the sensible powers of sight, taste, touch, and so forth. As such, completely misleading one’s experiences in such a case could be seen again as ‘masking’ someone’s rationality and as such not taking it away. If, however, someone is now trapped in a virtual reality machine for a significant period of time, then we can perhaps avail ourselves of a semantic externalist response. Perhaps they do not anymore have normal sense powers but *virtual* sense powers, in which case they could endeavor to judge about *that* environment. Since the virtual environment was constructed in our environment, it stands to reason that that the particular procedures for judgment could not be radically different; it is very likely that we could attribute a *different* albeit very *related* rationality to this person.

Much more needs to be said on disentangling the various related threads. However, I hope to have firmly established that it is not at all clear that these radical skeptical views are a serious problem for epistemic secure realism. It is entirely plausible that the resources in semantic externalist and dispositional/power literature are enough to deal with these types of thought experiments.

7.2 PROBLEMS IN SOCIAL JUSTIFICATION

We just treated cases where we attributed justification to others in extremely bad skeptical scenarios. What we have not treated yet is attributing justification in mundane cases. Epistemic secure realism inherits the problems that perceptual naive realist theories have in explicating the mundane *bad* cases of justification. There is little trouble in attributing justification to judges in *good* cases; one is justified in knowing when one K-judges in virtue of this endeavor functioning as an expression of the global power to get securely at facts. But it is less clear what we are to say in cases where we are *seemingly* justified, yet are still incorrect in our beliefs. The worry here is that since the environment, being the way it is, sets out both the true and false facts and the global power of how to get

at the facts, there could be no such a thing as having a justified yet incorrect judgment, or—at the very least—no such a thing as being exculpated from being unable to secure the worldly justification.

To put the problem in greater relief, take an act of judgment such as trying to determine the truth of some political scandal, such as whether Saddam Hussein had weapons of mass destruction in 2003. In trying to determine whether that was the case, as a civilian, my likely only recourse is to read the newspapers, ask friends more knowledgeable of political matters, and so forth. Having done this investigation, I decide to conclude—not unreasonably—that since the New York Times is an estimable source of news and reported it was the case he had the WMDs and since most US officials were saying the same, that it is the case that Saddam Hussein has weapons of mass destruction. However, as it later turned out, the political administration had been misleading the public by leaking false information to the reporter Judith Miller and then later quoting the article in the NYT she wrote as justification for the claim. Yet, this normal method of justification would *usually* result in reliable beliefs. Two questions need to be answered: first, why, if *this* instance of judgment is not an expression of the global power, it is not such an expression; second, what my justification is in judging the way I did if it was not an expression of the global power.

A glib answer might be to deny that this case is a true instance of justified judgment failure. Perhaps as readers we *should* have been more suspicious of the reporting of the New York Times and others. If we had, we would have noticed the discrepancies. But, intuitively, it seems wrong to put all the blame on the sole judger in this case. To do so would leave the Bush administration officials off the hook for what seems a misuse of the credibility of the New York Times and more general print and television media. Certainly, readers of news media bear responsibility in how they parse a news source's reliability, but so do the *maintainers* and *participants* of media institutions as well. Intuitively, the more widespread the unreliability is in the systems of testimonial knowledge, the more forgivable mistakes are for a judger.

I argue that it is not impossible to fold these intuitions into the manifestation conditions of a global power. In § 7.1 we saw that an unfavorable environment could scupper the manifestation conditions of the global power. We also used a Davidsonian triangulation argument in chapter 5, albeit against coherentism endorsed with the common kind assumption. However, if we now take our epistemic scheme in non-demon, non-BIV/NIV situations as reflecting a real global power, then the triangulation argument actually provides a profile of that power. Thus, in endorsing the triangulation of environment, the individual, and others, we see that all these parts have a role to play in the manifestation condition story. Just as our perceptual faculties were seen as part of the global power, so too our *testimonial faculties* might be seen as part of the faculty as well. A thorough-going instance of testimonial deception might be seen as just as much a case of disposition masking: I still possess the global power, even if part of it is *masked* by a local thorough-going unreliability of testimony. The judger is responsible for navigating the social web, yes, but the fault need not lie with the judger if the social world contains some thorough-going unreliability. The global power of judgment then is not one that belongs solely to a single judger. Rather, it extends to many corners of the world through

the *shared* responsibility of many epistemic actors working together in an epistemic community. In exercising one's individual virtues one is able to become part of this global system, and thus gain secure knowledge through a global power of judgment.

This however does not mean that *any* local unreliability provides an excuse for a failed instance of the global power. K-Judging should still have far more stringent requirements than any particular piece of passively acquired simple knowledge. Thus, any Gettier cases of accidental true belief as formed from a J-endeavor will still not result in K-judgment. Judging that a certain barn is a barn facade merely by its appearing as a barn from a distance will never amount to good judgmental practice, in part because the epistemic endeavor was far too passive. We know what the procedures are for determining whether an object is what it seems to be, and they involve either getting close to the object to inspect it from multiple sides, or investigating the object as a social entity and thus gaining secure access to the factive nature of the object through the extended epistemic power of testimony (or any other methods not mentioned now). What distinguishes judgment, the active acquisition of knowledge, from mere passively acquired knowledgeable belief is the sense of *conscious investigation* embedded in the endeavor.

This now allows us to grade the mistakes a knower makes in terms of the manifestation conditions for social knowledge. Though I am not solely responsible for social knowledge, I do bear an important responsibility in that *I* am the one affirming the proposition, and thus must bear the environmental and social risks on my own. The full manifestation of the power though is not a matter that is ultimately solely up to me, and thus can be graded in terms of friendliness of the environment and social sphere. No community of course is perfect and will have its flaws, so it might not be always entirely clear how to draw the lines of responsibility. But this does not mean there are not clear cut cases where the responsibility falls on the community rather than on the individual.

7.3 THE CORRECT PROCEDURE FOR GLOBAL JUDGMENT

The last issue we are going to treat is that of the correct method or procedure for judgment. Though correct judgment is in many ways an externalist procedure, it still must be done internally by the judge, or at the very least *consciously*. The question is then how in a world where we have a global power to get secure knowledge of the world we go about judging.

The simple way to solve this problem is to give a Wittgensteinian answer: we already know how we go about judging or—to put it in a phenomenological fashion—we are already thrown into a seeming judgment practice. However, now that we have the notion of a global judgmental endeavor, we can express the same point using a metaphysical inflection. So, instead of us getting thrown into a judgment practice, the world itself (in the good case) impresses itself onto our consciousnesses. Just as a child learns that a stove is hot by touching it, so do we learn the right ways to go about the world through our experiences of the way the world is. If we make a mistake in our estimation, we risk bodily harm; but if we succeed, we tend to get rewarded. This is not to say of course all beliefs must have some immediate pragmatic effect, or in fact should be understood solely in terms of their pragmatic effect. Rather, that the world, in being the way it is,

constrains our beliefs through its effects on our bodily existence. This same idea would also go for a community of knowers. A community that had a method that matched the global power would be more successful in pursuing their aims than another that had not. Thus, in any good case world, the access to experience and the stability of the community would itself insure secure access to the global power through our existing epistemic scheme. Since the world exists, and since our experience provides this transparent access, the worry that procedures for global judgment are unreachable should have no purchase; in fact, they would be forced on us at every turn.

Of course, when we are judging we do have to take an active stance, since we are trying ensure the security of our judgments through our own thoughts. Once again, we can understand this type of act in terms of our existing epistemic scheme. When I settle whether p through a judgmental endeavor, I am trying to aim to get such a security in the proposition that its not being true would significantly alter my epistemic scheme. I do this in the case of the barn by, e.g., looking at p from multiple angles, since this would now mean that p not being true would significantly put my belief in the reliability of visual sight in doubt, a belief that might function akin to a hinge proposition. Yet, these internal moves in my epistemic scheme now match the *external* secure cognitive contact I gain: insofar as my sight is a dispositional power that is part of my global power of judgment, my seeing the barn from multiple angles effects a secure cognitive contact with the fact that the barn is there. Similarly so if we investigate, say, the red-faced barn with intricate carvings as a social object. My investigation of multiple newspaper articles referring to the specific barn with those specific carvings grants me secure cognitive contact with the fact that that barn exists. So, insofar as my epistemic scheme represents reality as it actually is, the procedures contained within the epistemic scheme allow me to gain a secure hold over what is and is not true. Consciously integrating my belief into my epistemic scheme simply amounts to doing the procedures that provide secure metaphysical cognitive contact.

The overall point then is this: the worry that our judging practices are unknowable is unfounded. We already have a good sense of what it is to judge correctly, and our connection to the world through our bodily existence and social reality is so involved that the world itself—be it through physical or even social consequences—teaches us how to approach it. Moreover, though we understand our epistemic scheme usually as a set of beliefs, the procedures we do to establish our beliefs in the good case would be *worldly* procedures. We do experiments, we ask others in the know, we test our underlying beliefs and commitments; all these procedures are prescribed by the epistemic scheme and thus, through secure metaphysical connections, by the global power.

7.4 CONCLUDING REMARKS

Many of the answers I gave to the problems above have been under-described or incomplete. Still, I hope to have given a sense that epistemic secure realism does not immediately fail on the merits of more traditional epistemic problems in the philosophical literature. It certainly is true that endorsing the position would require a reconception of the nature of justification, or in fact the nature of intellectual virtue; any and all sources

of epistemic goodness would have to be world specific and inextricably bound to the world itself. But such a reconception would not be impossible. The intuitions that motivate most of the ideas of independent justification can be wholly maintained, as long we contextualize them within the world we actually inhabit.

All in all though, the task of this work seems to have been accomplished. I developed a transcendental argument according to which views that saw epistemic justification as having a common kind across good and bad cases are incoherent. As a result I advocated for the view of epistemic secure realism which rejects the common kind assumption in favor of a secure link to the world in the good case. The details of epistemic secure realism still need to be worked out. But, I hope to have made plausible that epistemic secure realism is the correct view to take of the metaphysics of knowledge and judgment.

WORKS CITED

- Brueckner, Tony. 2016. "Skepticism and Content Externalism." In *The Stanford Encyclopedia of Philosophy*, Winter 2016, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2016/entries/skepticism-content-externalism/>.
- Chalmers, David. 2002. "Does Conceivability Entail Possibility?" <http://consc.net/papers/conceivability.html>.
- . 2011. "The Nature of Epistemic Space." In *Epistemic Modality*, edited by Andy Egan and Brian Weatherson, 60–106. Oxford University Publishing.
- Cohen, Stewart. 1984. "Justification and Truth." *Philosophical Studies* 46 (3): 279–295.
- Conant, James. 2012. "Two Varieties of Skepticism." In *Rethinking Epistemology*, edited by Günter Abel and James Conant, 2:1–73. Walter de Gruyter.
- Davidson, Donald. 2001a. "A Coherence Theory of Truth and Knowledge." In *Subjective, Intersubjective, Objective: Philosophical Essays*, 137–158. Oxford University Publishing.
- . 2001b. "Three Varieties of Knowledge." In *Subjective, Intersubjective, Objective: Philosophical Essays*, 205–220. Oxford University Publishing.
- Edwards, Susan. 2005. *Fluent Aphasia*. Cambridge University Press.
- "Human Nature?" 1996. *European Journal of Philosophy* 4 (2): 235–254.
- Frege, Gottlob. 1892. "Über Sinn und Bedeutung." *Zeitschrift für Philosophie und philosophische Kritik* 100:25–50.
- Haddock, Adrian, and Fiona MacPherson, eds. 2008. *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press.
- Husserl, Edmund. (1901) 2001. *Logical Investigations*. Vol. 2. Routledge.
- James, William. 1890a. *The Principles of Psychology*. Vol. 1. Henry Holt & Company.
- . 1890b. *The Principles of Psychology*. Vol. 2. Henry Holt & Company.
- Kant, Immanuel. (1784) 2008. "Beantwortung der Frage: Was ist Aufklärung?" In *Akademieausgabe von Immanuel Kants Gesammelten Werken*, 34–42. Universität Duisburg-Essen. <https://korpora.zim.uni-duisburg-essen.de/Kant/verzeichnisse-gesamt.html>.
- Kenny, Anthony. (1973) 2006. *Wittgenstein*. Blackwell Publishing.
- Klein, Peter. 1999. "Human Knowledge and the Infinite Regress of Reasons." *Philosophical Perspectives* 13:297–325.
- Klein, Peter, and John Turri. 2013. "Infinitism in Epistemology." In *Internet Encyclopedia of Philosophy*, edited by James Feiser and Bradley Dowden.
- Littlejohn, Clayton. 2018. "Being More Realistic About Reasons: On Rationality and Reasons Perspectivism." doi:10.1111/phpr.12518.
- Martin, Michael G. F. 2002. "The Transparency of Experience." *Mind & Language* 17 (4): 376–425.
- . 2004. "The Limits of Self-Awareness." *Philosophical Studies* 120 (1–3): 37–89.
- McDowell, John. (1994) 1996. *Mind and World*. Harvard University Press.
- . 2009. "Sellars and The Space of Reasons." Last modified September 14. <http://www.pitt.edu/~brandom/me-core/downloads/McD%20Cape%20Town%20talk--Sellars%20EPM.doc>.

- McDowell, John. 2013. "Perceptual Experience: Both Relational and Contentful." *European Journal of Philosophy* 21 (1): 144–157.
- . 2017. "Travis on Frege, Kant, and the Given: Comments on 'Unlocking the Outer World'." Preprint, forthcoming in *In the Light of Experience: Essays on Reasons and Perception*, edited by Johan Gersel, Rasmus Thybo Jensen, Morten Sørensen Thaning, and Søren Overgaard. Oxford University Press.
- Orman Quine, Willard van. 1951. "Two Dogmas of Empiricism." <http://http://www.ditext.com/quine/quine.html>.
- Pritchard, Duncan. 2015. *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*. Princeton University Press.
- Schwitzgebel, Eric. 2008. "The Unreliability of Naive Introspection." *The Philosophical Review* 117 (2): 245–273.
- Sellars, Wilfrid. (1956) 1997. *Empiricism and the Philosophy of Mind*. Harvard University Press.
- Smithies, Declan. 2010. "Attention is Rational-Access Consciousness." Last modified August 11. <http://philpapers.org/archive/SMIAIR>.
- Sosa, Ernest. 1980. "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge." *Midwest Studies in Philosophy* 5 (1): 3–26.
- . 2009. *Reflective Knowledge*. Oxford University Press.
- . 2011. *Knowing Full Well*. Princeton University Press.
- . 2015. *Judgment and Agency*. Oxford University Press.
- . 2016. "Replies to Comments on *Judgment and Agency*." doi:10.1007/s11098-016-0733-0.
- Travis, Charles. 2016. "Deliverances (Indirection)." doi:10.1007/s11245-015-9357-9.
- Vaidya, Anand. 2017. "The Epistemology of Modality." In *The Stanford Encyclopedia of Philosophy*, Summer 2017, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.
- Wittgenstein, Ludwig. (1953) 2001. *Philosophical Investigations*. 3rd ed. Translated by G. E. M. Anscombe. Blackwell Publishing.