Virtue, Evidence, and Epistemic Justification

A Dissertation Presented for the

Doctor of Philosophy

Degree

The University of Tennessee, Knoxville

Alexander Steven Hallam
August 2016

To My Loved Ones

and

To anyone who seeks answers to difficult but important questions.

ACKNOWLEDGEMENTS

First I would like to thank my dissertation committee members, E.J. Coffman, Adam Cureton, Markus Kohl, and Stephen Blackwell. Their time and service is immensely appreciated. Special thanks go to E.J. Coffman, my director and committee chair, for providing thoughtful, challenging, and substantive feedback throughout the dissertation process. His expertise, advice, and guidance were invaluable in shaping my dissertation. Secondly, I would like to thank my loved ones (they know who they are). Without their love, encouragement, and support, I never would have completed this dissertation.

ABSTRACT

Evidence is a central concept in epistemology and more narrowly, theories of epistemic justification. Evidence is commonly thought to be what justifies our beliefs. On this view, a belief is justified for a person if that belief fits that person's total body of evidence. But it is also commonly thought that evidence isn't the only thing that justifies a belief. Some epistemologists even think that evidence isn't what justifies a belief at all. Virtue epistemologists give epistemic or intellectual virtues an important and fundamental role in theories of epistemic justification. On such views, for a belief to be epistemically justified, the belief must be formed responsibly or formed as the result of a reliable belief-forming process, faculty, or agent. Thus, virtuous character, agency, and inquiry are thought to be central to epistemic justification. The pressing issue to be explored in this dissertation is whether theories of epistemic justification in which evidence is central are compatible with theories of epistemic justification in which virtue is central. The aim of this dissertation is to argue for a hybrid view of epistemic justification in which evidence and virtuous (reliable) inquiry both play a salient role in the epistemic justification of a belief.

My hybrid or 'two-component' view of epistemic justification holds that a belief is justified along two dimensions: the *fittingness-dimension* and the *reliability-dimension*. More precisely, a subject's belief *p* is *categorically* justified (justified 'period' or 'full stop') when it *fits* the subject's total body of evidence *E* (the fittingness-dimension), but that in addition, the *reliability* of one's evidence-gathering methods (virtuous inquiry) can also play a salient role in increasing the *degree* to which *p* is justified (other things being equal) (the reliability-dimension). Furthermore, I will argue that being justified along the fittingness-dimension is all that is necessary for a belief to be justified full-stop (this also allows for the strength of one's

evidence to partly—sometimes wholly—determine the degree to which a belief is justified). Being justified along the reliability-dimension can only increase the degree to which the belief is justified.

TABLE OF CONTENTS

CHAPTER ONE: Introduction……...1

I. Preliminaries……1

II. Statement of the Problem...……....4

III. Goal and Structure of this Dissertation…….……..13

CHAPTER TWO: Intellectual Virtues, Virtuous Inquiry, and Epistemic Justification.....19

I. Preliminaries – Moral vs. Intellectual Virtues……...21

II. Virtue Responsibilism vs. Virtue Reliabilism……..22

II.1. Virtue Responsibilism…….……..23

II.2. Virtue Reliabilism…….……30

III. Virtuous Inquiry as a Reliable Belief-Forming Process……...…..44

III.1. Virtuous Inquiry and the Problem of Strange and Fleeting Processes…....48

III.2. Virtuous Inquiry and Greco's 'Stability' Requirement…….……...……...52

III.3. Virtuous Inquiry and Cognitive Integration…….57

IV. Vicious Inquiry……...59

V. Virtuous Inquiry as a Reliable *Evidence-Gathering* Method……..……61

VI. Conclusion…...……...……...……....61

CHAPTER THREE: Epistemic Rationality……..63

I. Epistemic Rationality and Instrumental Rationality.....67

I.1. Thomas Kelly's Argument…….69

I.2. Adam Leite's Reply to Kelly……..72

I.3. Kelly's Reply to Leite…….75

I.4. Instrumental Rationality vs. Methodological Epistemic Rationality……..77

I.5. Two Components to Epistemic Rationality?.....78

II. Synchronic vs. Diachronic Accounts of Epistemic Justification…….83

II.1. Trent Dougherty's Defense of Feldman—Reducing Responsibility……84

II.2. Responses to Dougherty—Recovering Responsibility?.....90

II.3. Doxastic Justification vs. Personal Justification……...93

II.4. Against Evidentialism: Keith DeRose and 'Epistemic Ought'……..…107

III. Conclusion……116

CHAPTER FOUR: Explanationist Evidentialism.....118

I. Evidential Fit……...……122

II. Juan Comesana and 'Evidentialist Reliabilism'…….128

II.1. Against Evidentialism…...……...…...……128

II.2. Reliabilism ……...……..131

II.3. 'Fit' ……...……..…133

II.4. Proto-Evidentialist Reliabilism ……..……134

II.5. Evidentialist Reliabilism ……135

II.6 Objection to Evidentialist Reliabilism…….137

III. Externalist Evidentialsim……..140

IV. Explanationist Accounts……...142

IV.1. Kevin McCain's Account of Evidential Fit……...144

IV.2. McCain's Account of the Basing Relation……154

IV.3. Objection to EX-WF……..164

V. Conclusions……169

CHAPTER FIVE: Towards a Two-Component Theory of Epistemic Justification.……...171

I. Evidence of Flawed Evidence Gathering (EFEG) as an 'Undercutting Defeater'…..172

I.1. 'Higher-Order' Evidence and Defeaters……...174

I.2. Michael Bergmann—'Defeaters and Higher-Level Requirements'…….177

I.3. 'Defeat' In Explanatory Terms…….180

I.4. The New Evil Demon Problem……...…….184

II. The Fittingness-Dimension….……...188

III. The Reliability Dimension.…….……..191

III.1. Support Facts and J-Factors……...…192

IV. A Two-Component View of Epistemic Justification ……...203

IV.1. Truth-Conduciveness and Justification……..205

IV.2 Virtuous Inquiry as a J-Factor...……206

IV.3. An 'Awareness' Requirement for Justification?..…….208

IV.4. Applications…….……..218

V. Conclusion..……...231

REFERENCES……236

VITA……243

CHAPTER ONE

Introduction

"A wise man proportions his belief to the evidence."

-David Hume – *An Enquiry Concerning Human Understanding*

Evidence is a central concept in epistemology and more narrowly, theories of epistemic justification. Evidence is commonly thought to be what justifies our beliefs. On this view, a belief is justified for a person if that belief is supported by that person's total body of evidence. But it is also commonly thought that evidence isn't the only thing that justifies a belief. Some epistemologists even think that evidence isn't what justifies a belief at all. Virtue epistemologists give epistemic or intellectual virtues an important and fundamental role in theories of epistemic justification. On such views, for a belief to be epistemically justified, the belief must be formed responsibly or formed as the result of a reliable belief-forming process, faculty, or agent. Thus, virtuous character, agency, and inquiry are thought to be central to epistemic justification. The pressing issue to be explored in this dissertation is whether theories of epistemic justification in which evidence is central are compatible with theories of epistemic justification in which virtue is central. The aim of this dissertation is to argue for a hybrid view of epistemic justification in which evidence and virtue both play a salient role in the epistemic justification of a belief.

I. Preliminaries

In general, an epistemically justified belief is one whose subject holds it in the right way. Different theories of epistemically justified belief, or *doxastic justification*, provide different accounts of what it is for a subject to hold a belief in the epistemically correct way. This sort of justification is different from what is often referred to as *propositional justification*, in which

justification is a property of a proposition (rather than a belief) relative to a person.[1] Propositional justification holds that a proposition can be justified for a person even if that person doesn't believe that proposition or even if she believes it for the wrong reasons entirely. More formally:

> **Propositional Justification:** proposition *p* is propositionally justified for a person S if and only if S has reasons or evidence for *p* such that if S were to believe *p* and base her belief that *p* on those reasons or that evidence, her belief that *p* would be doxastically justified.

Propositional justification is not the sort of justification that I'm concerned with. The sort of justification I'm concerned with is doxastic justification.

Doxastic justification is a property of *beliefs* rather than of propositions. In this dissertation I'm concerned with when a *belief* is epistemically justified. For a belief to be justified, a person must have good reasons or sufficient evidence for it, but also the belief must be *based on* those good reasons or evidence.[2] This is otherwise known as the *basing requirement*. Here is doxastic justification stated more formally:

> **Doxastic Justification:** Belief *p* is doxastically justified for a person S if and only if S's belief *p* is propositionally justified for S and S bases her belief *p* on whatever it is that makes *p* propositionally justified (i.e.—good reasons or sufficient evidence).

Unless otherwise specified, when I refer to 'epistemic justification' or 'justification' I'm referring to doxastic justification.

One more crucial distinction needs to be made about the sort of doxastic justification I'm concerned with. Importantly, we want to make sure that we're not confusing different epistemic properties. When we are talking about epistemic justification we are not talking about epistemic *blameworthiness* or *blamelessness*. That is, it's not the case that a belief is justified simply

[1] I take this distinction between doxastic and propositional justification from Michael Bergmann (2006, p.4).
[2] This general claim is not necessarily equating 'good reasons' with 'evidence'. A 'good reason' could also refer to something that justifies a belief other than evidence.

because a person is somehow blameless or blameworthy for holding the belief. Bergmann calls this sort of doxastic justification "subjective deontological justification" (2006, p.5). But this is not the sort of justification I am concerned with. I'm concerned with what Bergmann calls "ordinary justification" or a sort of justification that is more objective than subjective deontological justification (2006, p.6). Bergmann characterizes this more objective sense of justification by saying that it is the type of justification such that Jones' belief that *p* could have it and Smith's belief that *p* could lack it even if both Jones and Smith are equally epistemically blameless in believing in *p* (2006, p.5). To understand this distinction, consider this example by Bruce Russell (2001):

> Someone who grows up in a religious society and is taught to listen to the deliverances of an oracle can be epistemically blameless in believing those deliverances even though her belief may not really be supported by the evidence and so is objectively unjustified (Russell, p.36; Bergmann 2006, p.5).

Earl Conee and Richard Feldman (1985) give another example:

> A paranoid man might believe without supporting evidence that he is being spied on. This belief might be a result of an uncontrollable desire to be a recipient of special attention. In such a case the belief is clearly epistemically unjustified even if the belief is involuntary and the person cannot alter the process leading to it. . . . The person who believes that he is being spied upon as a result of an uncontrollable desire does not deserve to be blamed for that belief. But there is a fact about the belief's epistemic merit. It is epistemically defective—it is held in the presence of insufficient evidence and is therefore unjustified (Conee and Feldman, p.17; Bergmann 2006, p.5).

In both of these examples, the person is epistemically blameless for holding his belief, but it is still *objectively* unjustified. Bergmann suggests that objective epistemic justification is of the type that evidentialists like Conee and Feldman and Russell are concerned with. That is, when evidentialists talk about epistemic justification, they are concerned with propositional justification, in which a proposition can be justified even if a person doesn't believe that proposition and even if that belief doesn't satisfy a basing requirement. But evidentialists also

recognize an objective property of doxastic justification which is what S's belief *p* has when its content is propositionally justified for S and S bases *p* on whatever it is that makes its content propositionally justified (Bergmann, 2006, p.6). Bergmann ultimately disagrees with the evidentialists as to what is required for objective epistemic justification, but he does agree that it is something of this sort (2006, p.5). I agree.

Another way to explain the type of objective epistemic justification I'm concerned with is that it's the kind of justification that makes a belief a much better candidate for qualifying as an instance of knowledge than it would be if it lacked such justification. That is, when we talk about epistemic justification we are talking about a positive normative status that makes a belief a much better candidate for amounting to *knowledge* than it would be otherwise. Blameless beliefs of the sorts in the cases above are not good candidates for knowledge.

II. Statement of the Problem

One prominent theory of doxastic justification is advanced by *evidentialism*. It holds that the doxastic attitude that a person is justified in holding toward a proposition is the one that fits the person's *evidence* bearing on that proposition at some given time.[3] Generally speaking, the idea behind evidentialist accounts of doxastic justification is that justification turns *entirely* on evidence. However, there are also other popular non-evidentialist theories of doxastic justification that hold that epistemic justification does not turn entirely on evidence and that *epistemic* or *intellectual virtue concepts* play a significant or fundamental role in justification. Such theories are known as *virtue theoretic* accounts of epistemic justification.

Earl Conee and Richard Feldman, two prominent defenders of evidentialism, defend the following evidentialist principle of *propositional* epistemic justification:

[3] This also includes the belief being *based on* the evidence in question.

EJ Doxastic attitude *D* toward proposition *p* is epistemically justified for *S* at *t* if and only if having *D* toward *p* fits the evidence *S* has at *t* (Conee and Feldman, 2004, p.83).[4]

According to EJ, epistemic justification turns solely on evidence. For example, when a physiologically normal person under normal conditions sees a plush green field in front of him, believing that he is seeing something green is the attitude toward this proposition that fits his evidence. It is having the visual perceptual evidence of a plush green field that determines justification.

Additionally, Conee and Feldman defend a view of doxastic justification, or what they call, *well-foundedness*. In order for a doxastic attitude to be well-founded, the person must possess evidence that on balance supports that attitude (have propositional justification) and come to hold that attitude on the basis of that evidence (and not have any defeating evidence). Here is Conee and Feldman's well-foundedness more formally:

WF: S's doxastic attitude D at *t* toward proposition *p* is well-founded if and only if:
- (i) having D toward *p* is justified for S at *t*; and
- (ii) S has D toward *p* on the basis of some body of evidence *e*, such that
  - a) S has *e* as evidence at *t*;
  - b) having D toward *p* fits *e*; and
  - c) there is no more inclusive body of evidence *e*' had by S at *t* such that having D toward *p* does not fit *e*' (Conee and Feldman, 2004, p.93)

It is well-founded beliefs that I will be concerned with in this dissertation since in order to have a fully fleshed-out and complete evidentialist theory of justification, that theory must have an account of what it takes for someone to properly base his belief on his evidence. Why? Because well-foundedness is important for the evidentialist analysis of knowledge. That is, according to the evidentialist, a belief can't count as knowledge if it isn't properly based on the evidence one

[4] There are three doxastic attitudes: belief, disbelief, and suspension of belief. The justification for each of these attitudes turns entirely on the evidence. As I mentioned earlier, I'm concerned with when a *belief* is epistemically justified, but my discussion would equally apply to the other doxastic attitudes.

has (at least, according to evidentialists of the Conee and Feldman type). Thus, well-foundedness is what makes a belief a much better candidate for knowledge, unlike propositional justification.

In contrast to Conee and Feldman's evidentialist account of doxastic justification, there are virtue theoretic accounts of doxastic justification which hold that justification does not turn entirely on the evidence and that virtue plays a significant role, or even entirely determines, whether a belief is epistemically justified.[5] Hilary Kornblith (1983) argues for a *virtue responsibilist* view in which an agent is justified in his belief that *p* at time *t* as from time *t'* (where *t'* is earlier than *t*) just in case all of the agent's *actions* between *t'* and *t* which affected the process accountable for the presence of the belief that *p* at *t* were epistemically *responsible*. Lorraine Code (1984 and 1987) also argues for the centrality of epistemic responsibility in epistemology. She argues that epistemic justification is best understood as attaching to stable dispositions to act in certain ways that have their source in virtue. James Montmarquet (1993) argues for the view that S is justified in believing *p* insofar as *S* is epistemically virtuous in believing *p*. John Greco (2000a) argues for the view that S is justified in believing *p* if and only if *S*'s believing *p* results from the dispositions that *S* manifests when *S* is motivated to believe the truth (when *S* is thinking conscientiously). Linda Zagzebski (1996) argues for the view that a justified belief is what a person who is motivated by intellectual virtue, and who has the understanding of his cognitive situation a virtuous person would have, might believe in like circumstances. All of these accounts accept some sort of view of epistemic justification in which concepts of intellectual virtue are central or primary. In none of these accounts of epistemic

[5] I am charitably interpreting each of these accounts as theories of doxastic epistemic justification.

justification does evidence have primary status. That is, they all reject well-foundedness as a proper account of doxastic justification.

Jason Baehr (2011) offers a hybrid account of epistemic justification that is sympathetic to both virtue accounts and evidentialism. Baehr challenges the sufficiency of EJ, but he does not argue in favor of abandoning evidentialism, rather he aims to *supplement* evidentialism with a particular constraint, according to which, under certain conditions, justification requires *intellectually virtuous agency* (Baehr, 68). Baehr's argument for this claim centers around two types of issues: 1) cases in which the believer either fails to inquire or 2) inquires in some defective manner, i.e.—there is a 'mishandling' of the evidence. Consider the following four cases. Cases 1 and 2 are cases of defective inquiry. Cases 3 and 4 are cases of defective 'doxastic handling' of evidence.

CASE 1: GEORGE

> George epitomizes the vices of intellectual laziness, apathy, and obliviousness. He goes about his daily routine focusing only on the most immediate and practical of concerns: feeding himself, getting to work on time, doing his job in a minimally satisfactory way, paying the bills, etc. He lacks any natural curiosity and is almost entirely tuned out to the news of the day. Unsurprisingly, George has many beliefs he should not and fails to believe many things he should. In the former category is George's belief that exposure to secondhand smoke poses no significant health risks. Given his extremely narrow and practical focus, George is oblivious to all of the well-publicized research indicating the hazards of secondhand smoke. In fact, George actually has positive evidence in support of his belief. He recalls having learned from a reliable source some years ago that a considerable amount of research had been conducted concerning the effects of exposure to secondhand smoke and that this research had failed to establish any correlation between such exposure and any serious health problems. And as far as George knows, the research on this topic has not changed. Nor, we may suppose, does he have any reason to think that it might have changed (Baehr, p.70).

Baehr argues that regardless of the type of epistemic justification George enjoys according to EJ, there is another important sense in which these beliefs are unjustified. Specifically, while George's belief is well-supported by his evidence, the evidence is clearly not what it should be.

George has good evidence for his belief because of his intellectual "tunnel vision" (Baehr, p.71). George *ought* to have considered important and widely-publicized research. Since George's belief is based on evidence that ought to be other than it is, he ought not to believe as he does. In an important sense, George's belief is *unjustified*. The evidence that George has is defective and this defectiveness stems from his 'epistemic malfeasance'. Had George been even minimally attentive and discriminating, his perspective would have been very different and much more accurate, i.e.—he would have had better evidence (Baehr, p.73).

Now consider a second case:

CASE 2: GERRY

> Gerry holds the same belief as George and on roughly the same grounds. Therefore he too has positive evidence for thinking that secondhand smoke is benign. Unlike George, however, Gerry is not oblivious to the news of the day; in fact he is a reasonably inquisitive person and enjoys checking things for himself. The problem is that his inquiries tend to be insufficiently demanding and discriminating. They are prone to gullibility, carelessness, and hasty generalization. Upon hearing news reports affirming the danger of secondhand smoke, Gerry decides to look further into the matter. The first item he comes across happens to be a report published by an organization with major financial ties to the tobacco industry. The report is aimed, not at a fair and balanced treatment of the issue, but rather at exposing any apparent weakness or grounds for doubt in the recently publicized research. To any reasonably intellectually rigorous or discriminating inquirer, the dubious nature of the report would be evident. But to Gerry it is not. And the result is that Gerry's total evidence (which again includes his initial evidence for thinking that secondhand smoke is not a health threat) supports his belief (Baehr, p.70-71).

Gerry's evidence also well supports his belief but only because of his "undemanding and undiscriminating habits of inquiry" (Baehr, p.71). Baehr says that it is clear that Gerry is also unjustified.

In both of the above cases, George's and Gerry's beliefs are well-supported by the evidence they possess. Thus, EJ is satisfied. However, the problem as Baehr sees it, is that the *reason* these beliefs are well-supported has a lot to do with certain *defects* on the part of George and Gerry, i.e.—tunnel vision and undemanding and undiscriminating habits of inquiry (Baehr,

p.71). Thus, George's and Gerry's beliefs are unjustified even though the evidence they have supports their respective beliefs.

Let's now consider cases 3 and 4. In these cases, the epistemic justificatory status of the beliefs are affected by an *occurrent* mistake or defect by the person. More precisely, the justificatory status of the person's "doxastic handling" of information or evidence threatens to defeat or undermine the person's justification because of how the person treats or regards this information or evidence *at the time of belief* (Baehr, p.75).

CASE 3: DAPHNE

> [Like George above.] Daphne believes that exposure to secondhand smoke poses no serious health risks; she also has some positive evidence in support of this belief. However, she is neither intellectually lazy nor undiscriminating. Upon hearing about the relevant research, she does some looking into the matter and nearly all the information she comes across indicates that in fact environmental smoke is hazardous. Daphne's problem is that she is a hypochondriac raised by two chain-smoking parents. Owing to extreme anxiety about her health, she cannot accept any of the relevant evidence; indeed, she quickly and conveniently (though genuinely) forgets about or suppresses it. The result is that, as far as she can tell, her evidence continues to support her belief (Baehr, p.75).

In this case Daphne is suppressing evidence she is aware of that would otherwise cast major doubt on her belief. However, the evidence well supports Daphne's belief. So, Baehr is pointing out that even though Daphne's evidence supports her belief, it is still in some important sense unjustified and, like George, this results from her 'epistemic malfeasance' (Baehr, p.76).

CASE 4: DORIS

> Doris also believes with some positive evidence that secondhand smoke is benign. Upon hearing news reports to the contrary, she too engages in reasonably careful and discriminating inquiry on the matter and in doing so encounters a host of data that threaten to refute her belief. Like Daphne, Doris is unable to accept this data. But in Doris's case, this is due to her own extremely strong attachment to smoking. Her habit of smoking wherever and whenever she wants represents one of few pleasures in her otherwise lonely and unhappy existence. Unlike Daphne, Doris' cognitive constitution is such that she cannot simply 'forget' or suppress the relevant evidence. Instead she distorts or misrepresents certain critical aspects of it. The result is that from her standpoint, the case

> for thinking that environmental smoke is hazardous is weak and her belief remains well-supported (Baehr, p.75-76).

Like Daphne, Doris' belief is well-supported from her standpoint. Unlike Daphne, Doris doesn't suppress or ignore her evidence, but she does distort or misrepresent important aspects of it and then forms her belief on the basis of the resulting perspective. Consequently, Doris' belief is not justified (Baehr, p.76).

Baehr argues that there is an important and necessary relationship between epistemic justification and a person's evidence. He spends a great deal of time arguing for the following modified evidentialist principle:

> **E*** *S* is justified in believing *p* at *t* if and only if *S's* evidence at *t* appears to *S* to support *p*, provided that, if *S's* agency makes a salient contribution to *S's* evidential situation with respect to *p*, *S* functions, qua agent and relative to that contribution, in a manner consistent with intellectual virtue (Baehr, p.82).[6]

Baehr thinks that E* is a plausible evidentialist account of epistemic justification. E* essentially states that evidentialism must be supplemented by the constraint which, when applicable, requires epistemic agents to function in a manner characteristic of *intellectual virtue*. That is, even though evidence is necessary for epistemic justification, when a person's agency makes a salient contribution to the evidence that person possesses in support of some proposition at a given time, then the evidence is not enough to justify the person's attitude. When gathering the evidence, the person must have functioned in a manner consistent with employing intellectual virtue.

In response to cases like GEORGE and GERRY above, Conee and Feldman (2011a) argue that while George and Gerry should have gotten different evidence, such an evaluation is

[6] In subsequent chapters I will be careful to refer to a *belief* being justified and not a *person* being justified. There is an important reason for this distinction and it is crucial for my project. I will save this discussion for later but will briefly say that the reason for this distinction is that only beliefs can be epistemically justified, not persons (or so I will argue). Thus, E* should read "*p* is justified for *S*" instead of "*S* is justified in believing *p*".

an evaluation of their *conduct* and may be 'unjustified' in some sense here. For example, maybe George was unjustified in *inquiring* into the health risks of secondhand smoke. But this 'justification' claim about his inquiry is independent of whether his belief in the effects of secondhand smoke, under his actual conditions, has the sort of justification that is necessary for knowledge. The same applies to Gerry. The evidence that George and Gerry have consists in mental states and it is those mental states under the particular conditions and at that particular time that constitutes the evidence that they have. It is this evidence, and only this evidence that determines whether their beliefs are justified (Conee and Feldman, 2011a, p.313).

Conee and Feldman evaluate DAPHNE and DORIS somewhat differently, but with the same result. If Daphne's suppressing the evidence is understood as refusing to pay attention to some defeating evidence, then her evidence does not support her belief on balance. If her suppressing the evidence is understood as her immediately and completely forgetting the defeating evidence, then her perspective is as if she had never had the defeating evidence. In which case, she only has in her possession the initial supporting evidence. Thus, the belief that is justified is the one Daphne's evidence supports at that time. In Doris' case, Conee and Feldman say that if Doris' 'distorting' of the defeating evidence results in her having different evidence that does not at all defeat the initial support for her belief, then it is also as though she never had the defeaters. Her belief about second hand smoke is the justified doxastic attitude she should have. Thus, Conee and Feldman disagree with Baehr that Doris' and Daphne's beliefs are unjustified. However, they do agree with Baehr that their respective inquiries are badly done as a result of exercises of intellectual vices (Conee and Feldman, 2011a, p.314).

Feldman points to what he thinks is a crucial difficulty for the view that virtuous agency somehow plays a role in *epistemic* justification. To understand this he distinguishes between two

senses of the epistemic term 'rational' (Feldman, 2004b, p.233). First there is *current-state epistemic rationality*, which is the sort of epistemic appraisal that concerns whether believing some proposition is rational for some person at some exact time and in some particular circumstance the person happens to be in. This is different from another sort of epistemic evaluation, what Feldman calls *methodological epistemic rationality*. A belief is methodologically rational if and only if it is formed as the result of good epistemic methods (consideration of all the evidence, careful reflection, etc.).

Evidentialists like Conee and Feldman would contend that George, Gerry, Daphne, and Doris are current-state rational, but not methodologically rational. In each of the above cases, it would be *irrational* to believe differently given that their evidence quite clearly supports their respective beliefs. Each person's beliefs are precisely what they should be given the evidence at that particular time and under those circumstances. "Questions about what I should do, or what I should have done, with regard to evidence for a particular belief are independent of questions about the relation that belief has to the evidence I have at a given time" (Feldman, 2004b, p.235). Theories about current-state rationality are *not* theories of methodological rationality. Theories about the conditions under which beliefs are current-state rational are theories concerned with the conditions under which beliefs are well-supported by the evidence, i.e.—the fit between the evidence one has and belief. Such theories have nothing to say about the methods used to gather the evidence, what one should think about, or the methods by which beliefs should be formed.

Feldman thinks it's unclear whether methodological epistemic rationality is even an epistemologically central notion whatsoever (Feldman, 2004b, p.235). Instead, methodological rationality seems to depend largely on *practical* matters. For example: are there more important matters for George and Daphne to spend their time and energy on, what are the negative

implications of forming such beliefs, etc.? These are questions that cannot be answered without information about the agent's goals and preferences. Feldman thinks that central epistemological questions do not concern these sorts of practical issues, but rather are questions about the relation of beliefs to evidence.

III. Goal and Structure of this Dissertation

Given the vast and thriving literature that favors some sort of endorsement of methodological epistemic rationality (extant in virtue epistemology), as well as the existence of many discontents about evidentialism as a theory of epistemic justification, it seems prudent to explore the possibility that evidentialists like Conee and Feldman are wrong about methodological rationality. I intend to explore this possibility. More precisely, I'll explore the relationship between the evidentialist principle of well-foundedness and epistemic virtue, and their respective relevance to epistemic justification. My focus will be to examine the role that *intellectual virtues* such as open-mindedness, fair-mindedness, intellectual consistency, intellectual courage, attentiveness, thoroughness, consistency, etc., play in the justificatory status of a belief. In exploring this debate, there are important issues to resolve. One important issue to resolve is whether well-foundedness is necessary and sufficient for doxastic justification. I will argue that it is. However, I will argue that even if well-foundedness is necessary and sufficient for justification, this doesn't exhaust all the ways that the *degree* to which a belief is justified is determined. That is, the conditions that evidentialists like Conee and Feldman propose as minimally sufficient (i.e.—sufficient as well as necessary) for a belief's being well-founded, are minimally sufficient for a belief's being justified *period* (justified 'full stop' or 'categorically' justified), but that there are other factors, namely factors related to epistemic or

intellectual virtues, that also play a salient role in determining the *degree* to which that belief is justified.

My position as stated in the last paragraph means that I agree with Conee and Feldman that if a person's belief is well-supported by her evidence and she properly bases her belief on that evidence, then that belief is justified for her *period*. Thus, evidentialism is a satisfactory theory of categorical justification. But, I don't think that doxastic justification is solely a function of one's evidence. So I reject Conee and Feldman's account of well-foundedness as being a *complete* account of doxastic justification. I think that doxastic justification is also *often* a function of whether a person gathered her evidence virtuously or viciously. I say 'often' because one doesn't always engage in virtuous or vicious evidence-gathering methods (virtuous or vicious inquiry). Given this fact, intellectual virtues such as open-mindedness, thoroughness, intellectual courage, etc., are not a necessary condition for justification. Both Baehr and Conee and Feldman agree with me on this point. Baehr says that the virtue requirement in E* is not a necessary condition as it doesn't have universal application. That is, it only applies when a person's evidential situation involves or implicates her agency in some way (Baehr, p.82). On this point, Baehr says, "…justified…beliefs sometimes arise from the brute or default functioning of a person's basic cognitive machinery or endowment (*not* from an exercise of any intellectual character virtues)" (Baehr, p.80). For example, "…while working late one night in my well-lit study, the electricity suddenly and unexpectedly shuts off, immediately causing the room to go dark. In response I automatically and without thinking form a belief to the effect that the room has grown dark. Intuitively, my belief is justified" (Ibid). Conee and Feldman agree: "We also think that virtuous inquiry is not necessary for having all of the justification that knowledge requires. This is illustrated by any example in which a conclusive case for some

proposition is foist upon someone so forcefully as to overwhelm the person's actively engaged intellectual vices" (Conee and Feldman, 2011a, p.313). But contrary to Conee and Feldman, I *do* think that in many cases of forming a justified belief, the fact that one engaged in virtuous or vicious inquiry *does* often play role in determining the *degree* to which that belief is justified.

What I just said above does *not* mean that I wholly agree with Baehr either. Specifically, I reject E*. I reject Baehr's claim that engaging in *vicious* inquiry can render completely unjustified, a belief well-supported by the evidence the person has at that time (such as in cases 1-4). That is, I reject E*'s constraint regarding when or under what conditions a belief being supported by good evidence generates justification. To be more precise, I will argue that having good evidence in support of a belief is sufficient to justify that belief for that person full-stop, regardless of whether this evidence is against the backdrop of vicious inquiry. However, I do accept a constraint in which engaging in vicious inquiry can play a salient role in *decreasing* the degree to which a well-founded belief is justified. This is still a fundamentally important difference between my view and Baehr's.

I therefore accept something like the following account of doxastic epistemic justification:

DEJ: A) S's doxastic attitude D at *t* toward proposition *p* is well-founded if and only if:
- (i) having D toward *p* is justified for S at *t*; and
- (ii) S has D toward *p* on the basis of some body of evidence *e*, such that
  - a) S has *e* as evidence at *t*;
  - b) having D toward *p* fits *e*; and there is no more inclusive body of evidence *e*' had by S at *t* such that having D toward *p* does not fit *e*'

B) If S's engaging in *virtuous* inquiry makes a salient contribution to S's evidential situation with respect to *p*, then, other things being equal, the degree to which *p* is justified for S will *increase*.

C) If S's engaging in *vicious* inquiry makes a salient contribution to S's evidential situation with respect to *p*, then, other things being equal, the degree to which *p* is justified for S will *decrease*.

I will only hint at this point as to how virtuous and vicious inquiry contribute to increasing or decreasing the justificatory status of an already well-founded belief. It has to do with the fact that, other things being equal, virtuous inquiry increases the probability that the evidence that was gathered is *not* misleading. Likewise, other things being equal, vicious inquiry increases the probability that the evidence that was gathered is misleading. The *ceteris peribus* (other things being equal) clause is important because virtuous inquiry will not always make a justificatory difference. For example, there are circumstances in which, *all things considered*, virtuous inquiry cannot increase the degree to which a belief is justified (e.g., local and global deception scenarios). Evidentialists will reject clauses B and C, but I think many virtue theorists will embrace them. The rest of this dissertation is an argument in support of a hybrid view of evidentialism and virtue epistemology in which evidence and often the nature of one's inquiry (whether it's virtuous or vicious) both contribute to the justificatory status of a person's belief.

My goal in this dissertation is to argue for the following claims: 1) There is a satisfactory evidentialist account of categorical doxastic justification (i.e.—the concept of a belief's being epistemically justified period or full stop) and 2) virtuous inquiry is a reliable belief-forming process that can increase a justified doxastic attitude's degree of justification and vicious inquiry is an *un*reliable belief-forming process that can decrease a justified doxastic attitude's degree of justification. In arguing for 2, I want to point out that I'm following Alvin Goldman's (1979) view that justification is not a purely categorical concept. That is, certain beliefs are more justified than others and one important component in determining the degree to which a belief is justified is the comparative reliability of the belief-forming process that led to that belief (p.10). Of course, evidentialists will agree that justification isn't *merely* categorical, but they think that the degree to which a belief is justified for a person is solely a function of the strength of the

evidence the person possesses. I do agree that the strength of one's evidence can increase the degree to which that belief is justified, but I will also argue that the degree of justification is *also* often a function of reliability. Thus, I will argue for a *two-component theory of epistemic justification* in which justification is a function of one's evidence but is also often a function of reliability.

Here's how I will proceed. Chapter Two distinguishes between two different kinds of intellectual virtue properties—roughly, properties of *inquirers* (e.g., being open-minded) vs. properties of *inquiries* (e.g., being conducted or inquiring open-mindedly). Here I will raise objections to theorizing justified belief in terms of "inquirer-focused" properties. I will show that these objections do not foreclose the possibility of theorizing justified belief at least partly in terms of "inquiry-focused" properties. Here I will also give and defend a rough, preliminary statement of an evidentialist view of justified belief in which the exemplification of "inquiry-focused" properties can increase or decrease the degree to which a doxastic attitude is justified.

Chapter Three discusses two major issues. First, I will discuss the distinctions between different types of rationality: *epistemic rationality*, *methodological epistemic rationality*, and *instrumental or practical rationality*. I defend evidentialism as an important source of epistemic rationality from an initial objection to the effect that it simply conflates epistemic rationality with instrumental rationality. I will also argue that methodological epistemic rationality is not *merely* instrumental or practical rationality as Feldman claims, but that it also often plays an important role in fixing the degree to which a belief is justified. I then discuss a second issue which is whether an account of epistemic justification should be construed as *synchronic* or *diachronic*. Roughly, a synchronic account of epistemic justification holds that a belief is justified if it is based on good grounds at the time it is held (current time-slice justification). In contrast, a

diachronic account of epistemic justification roughly holds that a belief is justified if it results from responsible and/or reliable methods, proper practice, etc., leading up to the belief in question (cross-temporal justification). I will defend a hybrid account. Also, throughout Chapter Three I will discuss various arguments that support the evidentialist notion that having good evidence in support of a belief is sufficient for that belief to be justified. This discussion will help to motivate and set up a more fully fleshed out defense of evidentialism in Chapter Four.

Chapter Four is an extension of Chapter Three. Chapter Four motivates, presents, and defends what I think are the best "evidentialist" necessary and sufficient conditions for justified belief (full stop). Thus, Chapter Four fleshes out the details of an evidentialist account of doxastic justification. In doing so I will appeal to Conee and Feldman's work as well as what I think is the most detailed account of evidentialism in the literature, i.e.—Kevin McCain's (2014) *explanationist evidentialism*.

Chapter Five further develops, argues for, and defends the aforementioned two-component evidentialist/reliabilist view of epistemic justification.

CHAPTER TWO

Intellectual Virtues, Virtuous Inquiry, and Epistemic Justification

A good 'theory-neutral' definition of intellectual virtues is that they are characteristics or traits that are conducive to intellectual, cognitive, or epistemic flourishing or excellence. This chapter will explore the types of virtues that are significant to epistemic evaluations of belief within the context of responsible and reliable inquiry. By 'responsible and reliable inquiry', I simply mean the active and intentional process of attempting to find the true answer to some question or problem in which a) intellectual virtues (conscientiousness, open-mindedness, careful observation, fair-mindedness, tenacity, etc.) play a salient and prominent role and b) involves the use of reliable belief-forming processes such as reason, vision, memory, and introspection (as opposed to unreliable belief-forming processes such as tarot card or tea leaf reading).

One of two primary goals in this chapter will be to examine what is included in the category of intellectual virtues. I take it as uncontroversial among epistemologists that the types of virtues involved in responsible and reliable inquiry are the intellectual virtues. However, what I take as controversial among epistemologists is what counts as an intellectual virtue and how they figure into the epistemic evaluation of beliefs, i.e.—whether a belief counts as knowledge or is epistemically justified. The divide in this controversy centers around two different conceptions of intellectual virtues: intellectual *character virtues* and *cognitive faculties*.

Generally, intellectual character virtues are conceived of as traits or dispositions of *character* manifested in someone's agency. Examples include being persistent and thorough in one's analysis, being open-minded about new evidence, maintaining intellectual integrity, being careful in one's observations, etc. *Cognitive faculties*, on the other hand, generally include vision, hearing, deductive and inductive reason, memory, introspection, etc. These faculties are

often characterized as virtues, skills, or both. It is common for some epistemologists to deny the relevancy of intellectual character virtues to epistemic evaluation. For example, Alvin Goldman (1992) thinks that intellectual character virtues are of limited value in examinations of knowledge and epistemic justification. It is also common among some epistemologists to reject cognitive faculties as intellectual virtues. For example, Linda Zagzebski (1996) rejects the view that cognitive faculties are intellectual virtues and argues that only intellectual character virtues are of primary importance in epistemic evaluations.

The second primary goal of this chapter is to begin to argue for a plausible view for how intellectual virtues such as open-mindedness, fair-mindedness, careful observation, patience, etc., play a crucial and salient role in determining whether a belief is epistemically justified. Crucial to my view is that it is not intellectual virtues of *character* that play a crucial and salient role in the epistemic evaluation of belief, but intellectual virtues of *inquiry*. That is, virtues construed as character traits or dispositions have no bearing on the epistemic status of a belief (i.e.—whether the belief is epistemically justified/unjustified or counts as knowledge), but such virtues construed as *features* or *properties of inquiry* do (e.g., inquiring open-mindedly being a feature or property of some instances of virtuous inquiry). I will argue that this latter construal of intellectual virtues like open-mindedness, careful observation, fair-mindedness, etc., is the only viable way to make sense of how they bear on the justificatory status of a belief. To accomplish this goal, I will argue that it is the reliability of what I call 'virtuous inquiry' that bears on the justificatory status of a belief. By 'virtuous inquiry', I simply mean inquiry that is conducted in an intellectually virtuous way or manner—e.g., open-mindedly, fair-mindedly, consistently,

carefully, patiently, etc.[7] On my view, virtuous inquiry is a particular type of inquiry performed by an epistemic agent, where such inquiry is a reliable belief-forming process.[8]

This chapter is structured into three parts. In part I, I will briefly explain the difference between moral and intellectual virtues. In part II, I will explain the two dominant views in virtue epistemology: virtue responsibilism and virtue reliabilism. I will also point out problems with each view. In part III, I will begin to defend the plausible view that virtuous inquiry, as a reliable belief-forming process, often bears directly on the degree to which a belief is epistemically justified.

I. Preliminaries – Moral vs. Intellectual Virtues

In exploring the virtues relevant to virtuous inquiry, it is important to distinguish between moral virtues and intellectual virtues.[9] Let me start with a working conception of character virtues. Moral and intellectual virtues are a subset of character virtues. Character virtues are essentially traits of character that someone possesses that we evaluate positively. Examples of character virtues are courage, persistence, honesty, tolerance, fair-mindedness, humility, kindness, cautiousness, determination, benevolence, and so on. Notice that many character virtues are both moral character virtues and intellectual character virtues. As one example, fair-mindedness involves consistency in evaluation. This 'consistency in evaluation' might apply to the concern that all get their due (including oneself) in cooperative arrangements of mutual

[7] Strictly speaking, only people (more generally, thinkers) can be open minded, intellectually virtuous, etc. Thus, it's a category mistake to call an inquiry "open-minded". An inquiry can be conducted "open-mindedly" (i.e.—as an open-minded person would conduct it). However, the adjectival expressions I use such as "open-minded inquiry" and "intellectually virtuous inquiry" are just convenient shorthand for the more fundamental adverbial ones. The same thing applies to expressions like "reasoning open-mindedly" and "open-minded reasoning". This is important to make this explicit from the beginning as I will slip back and forth between the different types of expressions.
[8] In this chapter, I will refer to virtuous inquiry as a 'reliable belief-forming process'. However, in later chapters I will begin to refer to virtuous inquiry as a 'reliable *evidence-gathering* process'. As will become apparent, this latter description fits better with my own account of justification.
[9] In making this important distinction in part I, I appeal to Jason Baehr's account of virtues (2011, pp.1-2).

benefit or it might apply to how one evaluates evidence. This distinction is important because it helps us see the difference between moral character virtues and intellectual character virtues. I turn to this important distinction now.

To understand this important distinction we have to understand the ends that different virtues are concerned with. Moral character virtues distinctively involve motivations related to ends such as social justice, the alleviation of pain and suffering, etc. These types of virtues are concerned with the moral domain. Intellectual character virtues distinctively involve motivations related to ends that have an epistemic or intellectual dimension such as truth, knowledge, evidence, rationality, understanding, etc. These intellectual virtues are significant to responsibly or reliably conducted inquiry. This is not to say that there isn't a relationship between moral character virtues and inquiry, but the relationship there has to do with whether one ought to engage in certain inquiries as they relate to social justice, the alleviation of suffering, etc. However, the type of virtues directly relevant to evaluations of responsible or reliable inquiry are the intellectual virtues.

II. Virtue Responsibilism vs. Virtue Reliabilism

Virtue epistemologists all agree that intellectual virtues are cognitive excellences, but there is an important disagreement about what sort of cognitive excellence they are and how they figure into the analyses of knowledge and epistemic justification. The two standard views among virtue epistemologists are *virtue responsibilism* and *virtue reliabilism*. The responsibilist approach is character-based, while the reliabilist approach is faculty-based. A brief survey of virtue epistemology will help map out these two approaches.

There are too many players to account for in the virtue epistemology literature, so I offer what many take to be some of the key figures. Guy Axtell is largely credited with the distinction

between ‘responsibilists’ and ‘reliabilists’ (1997). Key ‘virtue reliabilists’ include Alvin Goldman (1992; 2001), John Greco (2000a; 2010), Alvin Plantinga (1993a; 1993b), and Ernest Sosa (1991; 2007). Key ‘virtue responsibilists’ include Lorraine Code (1984; 1987), Jonathan Kvanvig (1992), James Montmarquet (1993), Linda Zagzebski (1996), and Christopher Hookway (2000; 2003).

II.1 - Virtue Responsibilism

Responsibilists generally conceive of intellectual virtues as states of character or deep qualities of a person. Linda Zagzebski most prominently represents the responsibilist camp.[10] Zagzebski proposes a motivation-based theory of virtue according to which the concept of a motivation is fundamental, and the concept of a virtue is constructed out of the concept of a good motivation (p.82). Zagzebski says that a virtue is an excellence as well as a ‘deep’ trait of a person (p.89). However, on her view not all excellences are virtues (Zagzebski, p.102). Zagzebski argues that virtues are acquired excellences and therefore they cannot be natural excellences. It is often claimed that natural faculties (excellences) are intellectual virtues. For example, Ernest Sosa (1985) and John Greco (1993) list such natural faculties as vision, hearing, memory, and other cognitive faculties as intellectual virtues. Zagzebski thinks that in one sense this is correct because we do sometimes call any human excellence a virtue. Nevertheless, she argues for a much narrower Aristotelian view of virtue, namely that virtue is an acquired human excellence.

For Zagzebski, the philosophical motivation for pointing out the category of acquired excellences is the interest in focusing on those excellences we are responsible for (p.103). It is necessary (but not sufficient) that a quality or trait be acquired if we are to be responsible for

[10] Jason Baehr (2011) also provides a nice discussion of intellectual character virtues and their role in epistemology.

having it. For example, she says, "An animal acting on instinct alone does not have virtue, even if the instinct is a good one, and the same goes for the human" (Zagzebski, p.103). However, Zagzebski does concede that we might be able to shape or change some of our natural tendencies through education. It seems that she has in mind here the fact that we can become better reasoners or enhance our memory through education. If so, she thinks it might be appropriate to call the resulting quality a virtue. However, Zagzebski argues, "…we do not praise the natural faculty or capacity itself and blame the natural defect, and one reason for this is that these qualities are wholly involuntary" (p.103). At least with respect to moral virtues, that which is involuntary is outside the moral realm. But this voluntary/involuntary distinction largely obscures a crucial distinction required in eliminating natural faculties from the category of moral virtues, namely that a virtue is a deep quality of a person (Zagzebski, p.104).

A virtue is a deep quality of a person, closely identified with that person's selfhood, while natural faculties are only the "raw materials of the self" (Ibid).

> Natural faculties, capacities, and talents may be praised in the same way we praise natural beauty or strength, but we do not blame the lack of them. Virtues are qualities that deserve praise for their presence and blame for their absence. Even greater blame is due to a person who has the contrary of a virtue, namely, a vice, but we do not blame a person for having the contrary of intelligence or good looks (Ibid).

A virtue derives its praiseworthiness from the fact that the person who possesses the virtue could have developed the corresponding vice rather than the virtue they actually developed. "It is the fact that the person could have gone either way that distinguishes virtue from certain other excellences, particularly all those that are natural or inborn" (Zagzebski, p.105). Thus, Zagzebski holds that virtues are states of excellence that a person develops over time, but who could have developed the contrary vicious state instead. According to Zagzebski then, virtues are not natural faculties, but traits of character.

Another important feature that virtues possess is that they are formed through habituation (Zagzebski, p.116).[11] It takes time to develop virtues and vices. This fact is connected with the responsibility persons have for having these traits of character. "The features of gradual acquisition and entrenchment suggest that a virtue is a kind of habit…" (Ibid).[12] This of course brings up the issue of whether skills are virtues in some sense. Zagzebski argues that they are not. She provides a number of differences between virtues and skills, but there is one difference in particular that best captures this distinction in her view: "Skills serve virtues by allowing a person who is virtuously motivated to be effective in action" (Zagzebski, p.113). Let us look at some of the examples she provides.

First, Zagzebski discusses moral virtues and moral skills (p.113). Commonly used examples of moral virtues are compassion, generosity, and courage. Examples of moral skills that accompany each of these virtues are:

Skills of compassion: knowing what to say to the bereaved

Skills of generosity: being effective in giving to others

Skills of courage: knowing how to stand up to a tormentor

This also applies to the intellectual virtues (Zagzebski, p.114). Common examples of intellectual virtues are:

- Open-mindedness in collecting and appraising evidence
- Fairness in evaluating the arguments of others
- Intellectual perseverance, diligence, care, and thoroughness
- Being able to recognize reliable authority

[11] I will not go into detail about habituation. For more on Zagzebski's analysis of habituation, see section 2.5 in *Virtues of the Mind*.
[12] Zagzebski does not think virtues are identical to habits. For more see p.116-117 in *Virtues of the Mind*.

Examples of intellectual skills are (keeping in mind that many intellectual skills are not closely associated with specific intellectual virtues):

- Verbal skills: skills of speaking and writing
- Mechanical skills: e.g., knowing how to operate and manipulate machines and other physical objects
- Mathematical skills and skills of quantitative reasoning
- Spatial reasoning skills: e.g., skills at working puzzles

Zagzebski concludes that virtues and skills have numerous connections but that virtues are 'psychically' prior to skills (p.115). Virtues also have a broader range of application, whereas skills are more context specific. To sum up her view, she concludes that virtues are different from skills. Many virtues, but not all, have skills that correlate with them and allow the virtuous person to be effective when acting, thus we would expect a person in possession of some virtue to develop relevant skills. However, it is possible for someone to have a virtue but lack the relevant skill. Essentially then, Zagzebski thinks that virtues precede skills and are strongly connected to motivation, while skills are more connected to effectiveness in acting (p.116).

Now that we have Zagzebski's account of virtue, we are prepared to understand her accounts of knowledge and epistemic justification. I will begin with the latter. "A justified belief is what a person who is motivated by intellectual virtue, and who has the understanding of his cognitive situation a virtuous person would have, might believe in like circumstances" (Zagzebski, p.241). "An unjustified belief is what a person who is motivated by intellectual virtue, and who has the understanding of his cognitive situation a virtuous person would have, would not believe in like circumstances (Ibid)." On Zagzebski's account then, in order for a

belief to be epistemically justified it is necessary that the belief is formed as the result of intellectually virtuous motivation.[13]

The first thing one might notice about Zagzebski's account of epistemic justification is that it seems to fail to account for putative cases of justified beliefs that do not meet her conditions. If her view were correct, then it would seem that intellectually vicious persons, young children, and many animals couldn't have knowledge or justified beliefs. For example, for many perceptual and introspective beliefs there is no requirement that they be virtuously motivated for the belief to be justified. It would seem that I'm justified in believing that I see a red object in front of me, that I don't currently have a headache, that I'm currently sitting down, or that I hear a high-pitched tone in my ear. In each of these cases, it would seem that I'm justified in holding the perceptual or introspective belief I have at that moment whether their formation was virtuously motivated or not. Thus, if there are cases of epistemically justified beliefs that were not virtuously motivated, then it would seem that some of the conditions in Zagzebski's account of epistemic justification aren't in fact necessary for justified belief.

Zagzebski might have a way to respond to this criticism. To get to this response we first have to understand a bit about how she defines knowledge. According to Zagzebski, "Knowledge is a state of cognitive contact with reality arising out of acts of intellectual virtue (p.270)."[14] She defines an "act of intellectual virtue" in the following way:

---

[13] At times, I will use phrases such as 'virtuous motivation' or 'virtuous motives'. Strictly speaking, it would seem that only *persons* can be virtuous, but considering that Zagzebski and others commonly use these phrases in the literature, they are difficult to avoid. What such phrases are referring to are the *motivational component* of a virtue. The idea is that a belief is justified if it is what a person who is motivated by intellectual virtue and has a proper understanding of the cognitive situation that person is in, might believe in that circumstance. This is what it means to say that a belief was formed as the result of virtuous motives or virtuous motivation. It is not so much that the motive or motivation is, strictly speaking, virtuous, but that these phrases refer to a belief's formation being virtuously motivated.

[14] Zagzebski prefers this definition of knowledge to defining knowledge as a state of true belief arising out of acts of intellectual virtue. The former definition has the virtue of being noncommittal on such questions as the object of knowledge, the nature of truth, and the existence of propositions. She also thinks it allows a broader interpretation

> An act of intellectual virtue A is an act that arises from the motivational component of A, is something a person with virtue A would (probably) do in the circumstances, is successful in achieving the end of the A motivation, and is such that the agent acquires a true belief (cognitive contact with reality) through these features of the act (Zagzebski, p.270).

The obvious difference between a justified belief and knowledge is that knowledge requires that the belief be true or more broadly speaking is in cognitive contact with reality. A justified belief is a belief whose formation was virtuously motivated, whether the belief is true (etc.) or not.

Zagzebski also raises the objection I just mentioned about perceptual beliefs. The difference here is that the objection is directed at her conception of knowledge. She says that it would seem that her definition of knowledge fails to account for "low-end" knowledge, or the sort of knowledge that includes perceptual beliefs and short-term memory beliefs (Zagzebski, p.277). That is, it seems that a person can have knowledge without satisfying Zagzebski's definition of knowledge, i.e.—there seem to be cases of knowledge where the formation of the belief that amounted to knowledge was not virtuously motivated. On Zagzebski's account of knowledge, it would seem that only virtuously motivated persons could attain knowledge. Zagzebski's account of knowledge would seem too restrictive as an account of knowledge because animals, children, and intellectually vicious persons could not have knowledge of any kind no matter how trivial or how easily obtained.

The first thing that Zagzebski says in response to this worry is that her definition of an act of virtue doesn't require that a person actually possess the relevant virtue. Nevertheless, "…she must act the way a virtuous person would characteristically act in that same circumstance, and she must be successful because of these features of her act. What she may lack is the entrenched habit that allows her to be generally reliable in bringing about the virtuous end" (Zagzebski,

---

of knowledge since knowledge may very well include "cognitive contact with structures of reality other than the propositional" (p.271).

p.279). Presumably, this would allow those who lack or do not fully possess a virtue, to mimic a virtuous person and perform the acts in question. However, as Jason Baehr points out, adopting such a view suggests that Zagzebski's view is not genuinely or sufficiently virtue-based (Baehr, p.41). Baehr argues that, so construed, virtue is not doing any real explanatory work. On Zagzebski's view, beliefs amount to knowledge because they are formed as an intellectually virtuous person would form them. However, in the cases of the perceptual beliefs in question, the formation of the beliefs is merely brute or mechanistic (especially in the case of animals). This means that young children and animals lack any sort of motivation component to their beliefs whatsoever. "The upshot is that if this alternative formulation is correct, nothing having to do with virtue per se explains why the beliefs in question amount to knowledge" (Baehr, p.41).

Zagzebski's response is that there actually are "low-end" motives in play where there is "low-end" knowledge. In such cases, the intellectually virtuous person is (presumably) sometimes skeptical of her own senses when there is contrary evidence and she sometimes doubts her own memory when it is weak (Zagzebski, 1996, p.279). In cases of pain, she may not question that she is experiencing pain, but she may question, from time to time, why her introspection is trustworthy. Zagzebski maintains that most of the time the intellectually virtuous person doesn't engage in doubt or reflection about her perceptual and memory beliefs. However, the reason for this is that the person maintains a "presumption of truth" in such cases unless given reason to believe otherwise. This attitude is a virtuously motivated intellectual attitude, the opposite of which is what Zagzebski calls "intellectual paranoia". Thus, on this view, in cases of "low-end" knowledge, beliefs are virtuously motivated and result from acts of intellectual virtue (Zagzebski, p.280).

This response doesn't seem to help Zagzebski however. It doesn't get around the objection raised two paragraphs ago about young children and animals. Zagzebski's account of knowledge and epistemic justification is based on the assumption that there is some motivational component of a virtue in operation. Again, this seems to be gratuitous when it comes to many perceptual or memory beliefs. Baehr recounts the case of sitting in a brightly lit room, when suddenly and unexpectedly, the lights go out and the room goes pitch black (Baehr, p.44). In a case like this, you do not manifest any intellectual motives or action. Instead, you are overcome with the belief that it is pitch black. The belief that it is now pitch black is immediate and automatic. At no level do you seek the truth or engage in any introspection about the reliability of your vision. There are numerous examples like this one that not only involve vision but the other senses as well, e.g., when you hear a sudden bang or feel some intense pain. The same goes for memory. There are memories that just "pop" into our mind, e.g., when we smell cornbread and are reminded automatically and immediately of grandma's house. In none of these examples does anything like a motivational component of virtue appear to be operative. These examples therefore undermine Zagzebski's virtue-based account for not only knowledge but for epistemic justification. That is, if we run the same arguments against Zagzebski's view of epistemic justification, we get the same results.

II.2 - Virtue Reliabilism

As we have seen, virtue responsibilists do not classify our various cognitive faculties or abilities as intellectual virtues. Like virtue responsibilists, virtue reliabilists also recognize intellectual character traits as intellectual virtues but also argue that reliable cognitive faculties or abilities are intellectual virtues. These cognitive faculties or abilities include sense perception, memory, reason, introspection, etc. On this view, anything with a function has virtues (natural or

artificial) (Sosa, 1991, p.271). This would include things like hammers, people, machines, etc. On this view, a virtue is a quality of a thing that enables it to perform its function well. The function of the intellect is to attain truth and the faculty virtues (whether natural or acquired) make this possible (Battaly, 2008, p.644). If the reliabilist conception of intellectual virtues were correct, then virtue would indeed play a crucial and salient role in determining the epistemic status of a belief (I will discuss this in detail below).

Two prominent virtue reliabilists are Ernest Sosa and John Greco. Ernest Sosa considers cognitive faculties intellectual virtues (1991, Chapters 8, 13, and 16). He uses "virtue" and "faculty" as synonyms (Sosa, 1991, p.138-9, 234-6, and 273-4). Finally, he refers to intellectual virtues as "input-output devices" (Sosa, 1991, p.224) and "truth-conducive belief-generating mechanisms" (p.271). On Sosa's view, virtues are reliable in that they attain more true than false beliefs.

Perhaps the most ardent critic of virtue responsibilism among the virtue reliabilists is John Greco, so I'll focus my discussion on his arguments. Greco argues for two claims: First, contrary to Zagzebski, he argues that the intellectual virtues need not be acquired, need not contain a "strong" motivational component, and need not be an "excellence" in the Aristotelian sense. Instead, Greco argues that the essential component of an intellectual virtue is its "success component" or its reliability (2000b, p.179). Second, Greco argues that even if Zagzebski is right that the way human knowers are reliable is through exercising intellectual character virtues, then this says something about the mechanics of human cognition rather than the actual conditions for knowledge (2000b, p.180). I will explain both of Greco's arguments in turn below. However, before doing so, I should make a brief comment about the fact that what follows focuses on knowledge rather than epistemic justification. The crucial point of what

follows is to draw out the distinction between intellectual character virtues and faculty virtues. Although this discussion focuses on how and what intellectual virtues are relevant to knowledge, in doing so it will reveal the important role that virtues play in epistemic evaluation more generally. This includes epistemic justification.[15] First, let's look at Greco's two arguments for the two claims I mentioned above.

Greco's first argument is that Zagzebski makes Aristotle's account of moral virtue definitional of the concept of virtue in general and that this is problematic (2000b, pp.180-1). That is, Zagzebski takes Aristotle's view of moral virtues to be the model for understanding the intellectual virtues, but there are other non-Aristotelian conceptions that differ in important ways. It's not important for this discussion to go into detail about these differences. What is important is to point out that there are substantive differences between different accounts of virtue and that we can't just rely on Aristotle's definition as the definition for virtue in general. Greco adds that if we don't make Aristotle's definition of moral virtue definitional of the concept of virtue more generally, then we can see that Zagzebski, Sosa, and Goldman (Greco too) can all appropriately be labeled "virtue epistemologists". As Greco points out, the defining characteristic of virtue epistemology is that it makes the normative properties of persons conceptually prior to the normative properties of beliefs. That is, justified belief and knowledge are defined in terms of virtuous character. Greco argues that Zagzebski, like Goldman and Sosa (Greco too), do exactly

[15] Ultimately my goal is to evaluate epistemic justification, but I think the role that virtues of inquiry play in knowledge is largely the same for epistemic justification. Consider what Greco says about knowledge and justified belief: "On the present view knowledge and justified belief are grounded in stable and reliable cognitive character. Such character may include both a person's natural cognitive faculties as well as her acquired habits of thought. Accordingly, innate vision gives rise to knowledge if it is reliably accurate. But so can acquired skills of perception and acquired methods of inquiry, including those involving highly specialized training or even advanced technology. So long as such habits are stable and successful, they make up the kind of character that gives rise to knowledge" (Greco, 1999, p.287).

this. Consequently, Greco argues that the real issue is what type of character is required for knowledge and justified belief.

Greco argues that virtue reliabilists are correct in rejecting Zagzebski's conception of virtue. First, he argues that Zagzebski-type acts of virtue are not necessary for knowledge. Greco asks us to consider a cognitive agent who never engages in Zagzebski-type acts of intellectual virtue and yet can still be said to have knowledge:

> He never manifests the characteristic motivations of these virtues, and is never successful at bringing about their characteristic ends. For example, he never engages in acts that would be considered fair-minded, open-minded, careful, or thorough. However, suppose that despite all this the person is highly reliable in making correct judgments in certain important domains; he is almost never wrong in these areas of his expertise, and in fact outperforms other, more open-minded, fair-minded, careful and thorough persons. It seems to me that such a person does not lack knowledge for lack of Zagzebski-type virtuous acts. Rather, he acquires knowledge in some other way (2000b, p.182)
> .

Greco has us imagine some non-human cognitive agent who is infallible in human domains of inquiry, but who never manifested Zagzebski-type intellectual virtues. He thinks it's obvious that such agents would certainly be able to acquire knowledge, it would just be that they do so in a different manner than we do. The idea also seems to be that an agent's belief could be epistemically justified even though it was not motivated by Zagzebski-type virtues. An agent's belief is epistemically justified if the belief is formed as the result of the agent's reliability (even if the belief is false). For example, vision is a normally reliable belief-forming process, so an agent's belief is justified if it is formed as the result of vision and there is no defeating evidence, etc.

The second important claim that Greco argues for is that Zagzebski-type acts of virtue are not sufficient for knowledge (2000b, p.183-4). As he points out, Zagzebski-type acts of virtue

don't require the agent to actually possess the virtue.[16] Instead, it is possible for a vicious cognizer or a child to essentially imitate an intellectually virtuous person. A virtuous act is something that a virtuous person would (probably) do in the circumstances. But Greco points out that this opens up such a view to an objection that is raised against early versions of process reliabilism. Early versions of process reliabilism defined knowledge in terms of reliable processes without distinguishing precisely what the relevant processes are. This brought up the objection of fleeting processes. That is, it is possible for an agent to utilize a highly reliable process for a very brief time, and possibly very accidentally. Yet surely, we wouldn't attribute knowledge to this person.

The example that Greco gives is of a poor math student who utilizes a correct algorithm for solving some equation. In general, using the algorithm itself is a highly reliable process for solving equations. However, this particular student rarely uses an algorithm correctly. In fact, he usually chooses the wrong one or simply guesses at the answer. If so, then it seems incorrect to say that he has knowledge when he uses this reliable process. The reason, according to Greco, is that knowledge requires more from the agent. Simply adopting some reliable process is not enough. The agent must be reliable. Such agent-reliability might be in the form of some sort of reliable disposition to utilize reliable processes. However, in this case it is the agent-reliability and not the process-reliability that is doing the work (Greco, 2000b, p.183).

Now, given that Zagzebski doesn't require a knower to actually possess any intellectual virtues, Greco thinks we can also characterize Zagzebski-type acts of virtue as fleeting-processes. It might be that acting in the relevant way is highly reliable, but it would not be the agent who is reliable. The agent would not have the disposition to act in a reliable manner, and

[16] A point I discussed in the last section.

consequently would not have knowledge when she did act in this manner. Zagzebski does require that an agent have a "motive disposition" (Fairweather and Zagzebski, 2001, p.5), but this simply means that the agent possess a disposition to have the characteristic motivation of the relevant virtue. The difficulty is that Zagzebski doesn't require an agent to have a disposition for *success* in achieving the characteristic end of the virtue. In the case of knowledge, this would be arriving at true belief (Greco, 2000b, p.183). Greco suggests that one way to modify Zagzebski's view to deal with this objection is to add to her definition of an act of virtue that the agent actually possesses the virtue. Since such virtues have a reliability component, this would ensure that the agent who acts virtuously is reliable. But this brings up an important question: "…what role do Zagzebski-type virtues as such play in the revised conditions for knowledge?" (Greco, 2000b, p.184). Greco argues that they don't make an important contribution whatsoever.

Greco distinguishes between two aspects of the success component of Zagzebski-type virtues. First, the success component is defined in terms of the characteristic end of the relevant virtue, e.g., the characteristic end of open-mindedness is reliable success at being open-minded. Second, virtues involve success at achieving the final end of truth. Greco quotes Zagzebski here: "The intellectual virtues are a subset of truth-conducive traits that are entrenched and whose entrenchment aids their truth conduciveness" (Greco, 2000b, p.184; Zagzebski, p.178-9). However, Greco points out that Zagzebski-type virtues are not necessarily truth-conducive. He suggests that perhaps what Zagzebski really means is that in order for a character trait to be an intellectual virtue, it must actually be truth-conducive.

Now, if the above is true, Greco argues that there is a problem. If we understand intellectual virtues in the first way, then an agent could act virtuously but be unreliable, i.e.—there is no agent-reliability. If we understand intellectual virtues the second way, it will be

impossible to perform acts of virtue and not be reliable, so the above problem is avoided. However, if this is the case, then the agent-reliability is doing the work and not the Zagzebski-type virtues as such. That is to say that the agent has knowledge because she is reliable and not because she displays any of the other components of Zagzebski-type virtues. This shows that Zagzebski-type virtues are not relevant to the definition of knowledge. (The same thing applies to epistemic justification. It is the reliability component, not the motivation component of virtue, which contributes to the justification of a belief.) If Zagzebski is correct about virtues, she has at best given us a correct account of some special features of human cognition, i.e.—that humans are agent-reliable by having and displaying these virtues. However, Greco thinks that if this view is correct, it simply says something about the "mechanics of human cognition rather than the conditions for knowledge" (2000b, p.184). Ultimately then, if Greco is correct about Zagzebski-type acts of intellectual virtue, then they are neither necessary nor sufficient for knowledge (the same issue applies to epistemic justification).

Jason Baehr (2011) provides a promising way to reply to virtue reliabilists such as Greco. In particular, he offers an argument to support the view that intellectual character virtues sometimes satisfy virtue reliabilist's formal conditions for an intellectual virtue that is relevant to the analysis of knowledge. In conjunction with this point, the fact that virtue reliabilists typically view knowledge as true belief arising from an exercise of intellectual virtues, helps explain that intellectual character virtues are indeed relevant to virtue reliabilist accounts of knowledge. Greco, for example, defines an intellectual virtue as a personal quality that plays a critical and salient role in a person obtaining the truth, i.e.—only if it best explains why a person reached the truth (Greco, 2003; Baehr, p.52).[17] Baehr argues that there are many cases where getting to the

[17] It is important to understand what Greco means by 'salient'. Generally, when Y occurs because X occurs, we identify X's occurring as a crucially important or salient part of the causal history behind Y's occurring. For

truth demands a great deal of agency and simply appealing to cognitive faculties to explain the acquisition of truth is inadequate. It seems that in cases of "high-end" knowledge (e.g., scientific knowledge), getting at the truth is largely explained in terms of the exercise of intellectual character traits such as carefulness, thoroughness, tenacity, adaptability, creativity, circumspection, attentiveness, patience, honesty, etc. (Baehr, p.53).

If we consider cases involving scientific investigation, it is clear that reaching the truth isn't just a matter of exercising our cognitive faculties such as vision, hearing, inductive reasoning, etc. Baehr has us consider the following case as an example:

> A field biologist is trying to explain a change in the migration patterns of a certain endangered bird species. Collecting and analyzing the relevant data is tedious work and requires a special eye for detail. The biologist is committed to discovering the truth, and so spends long hours in the field gathering data. He remains focused and determined in the face of various obstacles and distractions (e.g., conflicting evidence, bureaucratic roadblocks, inclement weather conditions, boredom, etc.). He picks out important details in environmental reports and makes keen discriminations regarding the composition and trajectory of several observed flocks. As a result of his determination and careful methods of inquiry, he discovers why the birds have altered their course (p.53-4).

In this case, the biologist reaches the truth largely due to manifesting certain intellectual character traits. That is, these intellectual character traits seem to account for or best explain why the biologist reached the truth. The biologist discovers the reason behind the alteration in migratory pattern because of his "…patient, focused inquiry and his refined powers of observation and discrimination" (Baehr, p.54). From this, Baehr concludes that intellectual character virtues do sometimes satisfy the virtue reliabilist's conditions for an intellectual virtue. That is, in certain cases, intellectual character virtues can and do play a critical and salient role in reaching the truth (Baehr, p.54).

---

example, to say that a fire was caused by S's negligence is not to say that it caused the fire by itself, but it is to say that S's negligence is a salient part (perhaps the most salient part) of the total set of causal factors (Greco, 2003, p.118).

Virtue reliabilists, according to Baehr, claim that the difference between intellectual character virtues and faculty virtues, or cognitive abilities and skills, has to do with the source or origin of a belief. A true belief counts as knowledge if and only if its source is reliable and cognitive faculties are the only reliable sources of belief. If true, this would provide a principled exclusion of intellectual character virtues as reliable sources of belief. However, as Baehr says, this is all dependent upon what precisely a "source" of belief is. Baehr suggests that on a broad construal, something is a source of a belief if and only if it is the cause or salient cause of a belief (p.56). On this broad construal, Baehr thinks that intellectual character virtues are obviously sources of belief (consider the biologist example). On a narrow conception of a "source" of belief, something is a source of belief if and only if it generates beliefs independently of other beliefs or generates them in an immediate and non-inferential way (Baehr, p.57). On this conception, things like memory, introspection, intuitive reason, and the sensory faculties are all sources of belief, while the intellectual character virtues are not. However, Baehr thinks there is good reason for virtue reliabilists to deny such a conception.

If the narrow conception of a source of belief were correct, then virtue reliabilists wouldn't be able to account for cognitive faculties that do not operate in an immediate and non-inferential way. This narrow conception would rule out things like explanatory reasoning, deduction, and induction. Luckily, the reliabilist does not rule out such faculties. They are exemplars of what reliabilists consider intellectual virtues. If they did not accept this broader conception of intellectual virtues, then it would severely limit the class of things we could know, or be justified in believing, to the immediate and non-inferential. Baehr also denies that there is some less restrictive conception of intellectual virtue that would include these inferential and non-immediate cognitive faculties, while excluding the intellectual character virtues. He says

(and this is important), “This is because an exercise of character virtues is sometimes (perhaps always) manifested in and partly constituted by an operation of faculty virtues” (Baehr, p.58).

To see this, Baehr thinks it’s important to notice that explanatory reasoning, deduction, and induction are better classified as activities rather than as “mere default modes of cognitive functioning” (p.58). That is, there is a more active dimension to the exercise of these former faculties than the latter. Again, consider a scientist (such as the biologist) who acts out of open-mindedness, thoroughness, tenacity, a genuine commitment to the truth, etc. Baehr asks how we should understand the connection between the scientist’s acts of reasoning and her exercise of various character virtues. One way to look at it is that traits like open-mindedness, tenacity, a commitment to the truth, etc., motivate her to continue her inquiry (or, I would add, engage in the inquiry to begin with). This seems obvious, but these traits might also lead the scientist to think through the data in a reasonable way or to draw valid conclusions from the data. The scientist’s open-mindedness might help her avoid making logical errors or perceive otherwise unnoticed logical connections. Baehr argues (and this is crucial) that we cannot draw sharp distinctions between the scientist’s reasoning on the one hand, and her exercise of open-mindedness on the other. It is not the case that she displays open-mindedness and then reasons in the relevant ways. Instead, her exercise of open-mindedness is partly constituted by her acts of reasoning (Baehr, p.59).

Baehr further argues that the relation between intellectual character virtues and standard reliabilist virtues also extends to the functioning of basic cognitive faculties. The biologist gets at the truth about a change in migration patterns as the result of exercising certain intellectual character virtues. “As he studies the birds’ new winter habitat, he notices or sees certain subtle but critical geographical details that would normally go unnoticed” (Baehr, p.59). Baehr asserts

that the exercise of the relevant character virtues in this case is partly constituted by the operation of his visual faculty, i.e.—his inquiring in a careful and attentive way just is (or mostly is) a matter of making certain visual observations.

Essentially, Baehr thinks that intellectual character virtues are personal qualities that often play a critical and salient role in a person obtaining the truth, i.e.—they often best explain why a person reached the truth. Thus, intellectual character virtues are often a reliable source of belief. However, this view is not without some significant problems.

Ernest Sosa (2015) finds Baehr's arguments unconvincing. According to Sosa, "Knowledge is analyzed as belief whose correctness manifests the believer's pertinent competence" (Sosa, 2015, p.40). Consequently, the relevant reliabilist intellectual virtue (competence) must be one whose exercise can *constitute* knowledge. Sosa claims that knowledge is belief that is correct, that thus succeeds, *through* the exercise of competence. But Sosa restricts the "through" part of this definition because a belief might attain correctness "through" competence merely because the exercise of some competence puts one *in a position to know*.[18] This is not sufficient for the correctness of that belief to constitute knowledge.

> For, that exercise of competence may not immediately take the form of the correctness of a belief. It may rather take the form of putting one in a position to exercise a competence, such as *sorting by eyesight*, whose exercise *does* amount to a correct belief, a correct sorting" (Sosa, 2015, p.41).

Sosa's claim is that certain competences might put one in a position to know something, which is different from the exercise of that competence taking the form of the correctness of belief. To make this point clearer, Sosa has us consider a very simple example:

> Suppose a mysterious box lies closed before us, and we wonder what it contains. How can we find out? We might of course just open the lid. In pursuit of this objective we will then exercise certain competences, perhaps even character traits (if the box is locked, or the lid

[18] Sosa says that to be "in a position to know" that *p* is to possess the complete competence whose manifestation in a true belief that *p* would constitute one's knowledge that *p* (Sosa, 2015, p.43).

> stuck), such as persistence and resourcefulness. And perhaps these qualities (in certain contexts, and in certain combinations) do lead us reliably to the truth. Nevertheless, the exercise of *such* intellectual virtues need not and normally will not *constitute* knowledge, not even when that exercise does indirectly lead us to the truth (Sosa, 2015, p.42).

This is different than what happens when we finally open the box and look inside of it. Once we look in the box then we have an immediate answer to the question of what is in the box—say, that there is a necklace in the box—*through* a perceptual belief which manifests certain cognitive competences for gaining visual experience and belief. A belief manifesting such a competence and, more importantly, one whose correctness manifests such a competence, *does* constitute knowledge (Ibid). It may be that certain intellectual virtues such as open-mindedness or intellectual courage may put one in a position to open the box in the first place, but on Sosa's view manifesting such virtues is not constitutive of knowledge. Thus, there are intellectual virtues whose manifestation helps to *put you in a position to know* and other intellectual virtues whose manifestation in the *correctness* of a belief thereby *constitutes* knowledge (Sosa, 2015, p.43). Intellectual character virtues such as open-mindedness, intellectual courage, etc., are not directly knowledge-constitutive virtues.

Sosa thinks that 'knowledge-constitutive competences' are of main interest to virtue reliabilism. On the other hand, epistemically important character traits such as open-mindedness, intellectual courage, persistence, etc., are of interest to epistemology more broadly and are worthy of serious study, but they aren't central in traditional epistemology. They are only "auxiliary" intellectual virtues (Sosa, 2015, p.42).

Marlin Sommers finds Baehr's arguments unconvincing for similar reasons.[19] He argues that Baehr's arguments do not adequately support the claim that virtue reliabilist accounts of knowledge should include intellectual character virtues, and there is reason to think these virtues

[19] Marlin Sommers. "Can Character Intellectual Virtues Be Reliabilist Virtues?" (Unpublished).

will not fit into reliabilist accounts. In response to the biologist case, Sommers asserts, "…the virtues involved seem to explain why the agent obtained specific beliefs about the subject at all rather than why he obtained true as opposed to false beliefs." The idea is that if the biologist had not manifested certain intellectual character virtues he never would have engaged in the inquiry in the manner he did and consequently wouldn't have obtained any new findings. Any true beliefs the biologist might have formed would have been vague, broad, and uncertain. Sommers argues that the true beliefs in the biologist case are *because* of intellectual character virtues but that is the not the kind of 'because' for reliabilist theories. Sommers points out that a virtue can bear on epistemic conduct without being epistemic in the more specific sense required by reliabilist theories. For example, it might require intellectual courage to engage in the type of research some biologists engage in, but this doesn't mean that this courage explains the biologist's successful inquiry, i.e.—the sort of excellence performed by the biologist to get to the truth. Character virtues call for specific actions but they are not themselves skills that an agent can employ during inquiry to get to the truth.

I think Sosa's and Sommer's arguments can also be run as objections against the role that intellectual character virtues play in the epistemic justification of beliefs. Suppose the biologist case is concerned with epistemic justification. The possession of certain intellectual character virtues such as intellectual courage, allows the biologist to improve his epistemic position with respect to his inquiry. That is, intellectual courage is crucial in getting the biologist to engage in the type of research that he needs to in order to form justified beliefs about the migratory patterns of the birds in question. However, the biologist's belief is not justified 'because' of intellectual courage, rather it is partly justified 'because' of the reliable belief-forming skills and processes that he engaged in to form his belief about the migratory patterns of the birds in question. In

other words, the biologist needed intellectual courage to engage in the relevant research in the first place, but what contributed to the epistemic justification for his belief was the reliable belief-forming skills and processes he used *during* his research.

I would also like to add that various intellectual character virtues are crucial *during* the biologist's research, i.e.—not just as a cause of the biologist engaging in the research in the first place. For example, it might require further intellectual courage to *continue* to engage in research or it might require a certain amount of open-mindedness to open up new directions of research, but these intellectual character virtues do not play a role in the actual justification of beliefs. However, these intellectual character virtues are crucial in the sense that without them, the biologist might never have engaged in the research in the first place or decided to engage in certain observations during his research. Essentially, intellectual character virtues create opportunities for the biologist to explore, but are not relevant to a belief's justification.

I think these criticisms of Baehr's view are important and accurate. However, these criticisms have no bearing on the view of intellectual virtues that I will argue for in part III. I agree that intellectual *character* virtues do not play a crucial role in determining the epistemic status of a belief. However, my view is fundamentally different from Baehr's view. I'm interested in the relationship between virtuous *inquiry* and a belief's epistemic justificatory status. This means that, strictly speaking, I'm not concerned with intellectual *character* virtues whatsoever. Rather, I'm concerned with intellectual virtues of *inquiry*, i.e.—intellectually virtuous inquiry, not intellectually virtuous *character*. As I see it, the intellectual virtues that are directly pertinent to a belief's justification are not *properties of a person*, but *properties of inquiry*. Hence, I'm not going to argue that what does the explanatory work in the formation of certain justified beliefs are intellectual *character* virtues, but rather reliable-belief forming

processes (in this case virtuous inquiry). This should provide at least a good idea for what I'll argue for in part III.

III. Virtuous Inquiry as a Reliable Belief-Forming Process

In part III I will begin to make sense of the role that 'intellectual virtues' like open-mindedness, fair-mindedness, careful observation, patience, etc., or what I call virtues of inquiry, play in a belief's justificatory status.[20] Again, I wish to reinforce a crucially important distinction between virtues of character and virtues of inquiry. As we have seen, virtues of character also include open-mindedness, fair-mindedness, courage, patience, etc., but they are virtues that make up an agent's character, i.e.—the agent possesses these virtues as traits of deep character. They are dispositions that motivate the agent to act in certain ways. For example, in the biologist case above, the biologist possesses the virtue of intellectual courage and this courage motivates him to engage in certain types of research. However, it is not this courage that explains why the biologist reaches the truth, i.e.—the resulting true belief is 'because' of the biologists courage but in the wrong sense of 'because'. Rather it is the biologist's *excellent performance* in using reliable belief-forming processes that we credit with reaching the truth. The biologist engages in the research *because* of intellectual courage, but he reaches the truth *because* of his excellent skill as a researcher. Thus, it is not character but reliable belief-forming processes that are relevant to the justificatory status of a belief. That is, characterizing intellectual virtues as virtues of inquiry rather than of character allows us to make sense of the 'because' relation that is important to a belief's epistemic status, i.e.—whether the belief amounts to knowledge or is justified.

[20] Importantly, however, what follows is not the full account of my view, but a rough sketch of some of the key pieces. A full account will be fleshed out in Chapter Five.

Virtues of inquiry are *properties* of inquiry. They are not character traits possessed by an agent. Instead, these virtues belong to certain types of inquiry, i.e.—inquiry that is conducted the way an intellectually virtuous person would conduct an inquiry (e.g., inquiring open-mindedly, inquiry involving careful observation, inquiring patiently, etc.).

On my view, an agent need not possess any of the intellectual character virtues in order to engage in virtuous inquiry. In other words, engaging in virtuous inquiry can be accidental or simply done on a whim. An agent who has an intellectually vicious character can *inquire* open-mindedly without possessing the *character* virtue of open-mindedness. For example, a person who is otherwise disposed to be close-minded can inquire open-mindedly and be successful in doing so, i.e.—reaching the truth because he inquired open-mindedly. Virtuous inquiry is a reliable belief-forming process that is part of the agent's excellent performance or activity in forming a belief. It is 'because' of the virtuous inquiry that the agent formed the relevant belief.

Rather than conceiving of the intellectual virtues that play a salient role in determining the justificatory status of a belief as character virtues, we should view them as virtues of inquiry. Virtues of inquiry are properties or qualities of inquiry that, under certain conditions and with respect to certain propositions, are reliable means to reaching the truth or avoiding error. In this sense, intellectual virtues, as virtues of inquiry, often play a critical or salient role in determining the justificatory status of a belief. To illustrate my point, let's look at Sosa's 'mysterious box' and Baehr's 'biologist' examples.

Take Sosa's example first. It is true that manifesting certain intellectual character virtues puts you in a position to know or form a justified belief with respect to what is in the 'mysterious box'. But the manifestation of such intellectual character virtues is not constitutive of knowledge or justified belief. It is when you open the box and look inside that the intellectual

virtues constitutive of knowledge and justified belief are properly manifested. However, once you open the box and your visual perception is engaged, it is not just any sort of visual perception that plays a crucial and salient role in forming a true belief about what's in the box, but *careful* and *attentive* observation. 'Carefulness' and 'attentiveness' are specific properties of the inquiry being used to directly answer the question of what is in the box, i.e.—*carefully* and *attentively* inspecting the contents of the box using one's vision. Someone could look into the box, but do so in a cursory or perfunctory manner and therefore *not* pick up the same visual information as someone who carefully and attentively observed what was in the box. A perfunctory visual examination may only yield visual perceptions that are vague and indefinite (perhaps *something* shiny and platinum colored), while *careful* and *attentive* visual observation will likely yield visual perceptions that are distinct and definite (a *necklace* that is shiny and platinum colored). Importantly, all of this takes place *after* the box has been opened and the person *looks* inside of the box. The carefulness and attentiveness in this case are part of the visual process of observing what's inside the box. It is not as if a person is careful and attentive and then visually observes what's inside the box. Instead, the visual process of observing what's in the box is partly constituted by its being done carefully and attentively.

Baehr would say that the exercise of the relevant *character* virtues in this 'mysterious box' case is partly constituted by the operation of his visual faculty, i.e.—his inquiring in a careful and attentive way just is (or mostly is) a matter of making certain visual observations. But this is incorrect. The 'carefulness' and 'attentiveness' are *not* character traits, rather they are properties of a type of visual observation, i.e.—careful and attentive visual observation. Part of what constitutes a successful inquiry into answering the question of what is in the box is visually observing what's inside the box in a careful and attentive manner. It is the visual observation *as*

a form of inquiry that is relevant to knowledge and justification, and carefulness and attentiveness are *properties* of this particular instance of visual observation. The same thing applies to the reasoning employed in the biologist case.

Again, consider the biologist who acts out of open-mindedness, thoroughness, tenacity, a genuine commitment to the truth, etc. Remember, Baehr asks how we should understand the connection between the biologist's acts of reasoning and her exercise of various *character* virtues. We saw that one way to look at it is that traits like open-mindedness, tenacity, a commitment to the truth, etc., motivate her to continue her inquiry once she has engaged in the inquiry. That is, it's not just that open-mindedness, tenacity, a commitment to the truth, etc., put her in a position to know or form justified beliefs because these traits might also lead the biologist to think through the data in a reasonable way or to draw valid conclusions from the data. The biologist's open-mindedness might help her avoid making logical errors or perceive otherwise unnoticed logical connections. Again, Baehr argues that we cannot draw sharp distinctions between the biologist's reasoning on the one hand, and her exercise of open-mindedness (for example) on the other. It is not the case that she displays open-mindedness and then reasons in the relevant ways. Instead, her exercise of open-mindedness is partly constituted by her acts of reasoning. Baehr's mistake is in thinking that what is relevant to knowledge and justification are intellectual *character* virtues. Sosa's and Sommer's objections hold if we think of the intellectual virtues in play as character virtues, but these objections are beside the point if we conceive of the intellectual virtues directly pertinent to knowledge and justification as properties of inquiry. Then it makes sense to say that the biologist's reasoning is partly constituted by its being conducted open-mindedly.

Reasoning can have different properties, e.g., properties such as 'hastiness' or 'close-mindedness' and properties such as 'impartiality' or 'open-mindedness'. The reliability of the reasoning a person engages in is determined (largely) by the properties of that reasoning. So, if the reasoning is conducted in a hasty and close-minded manner (vicious inquiry), then it will be less reliable (other things being equal) than if the reasoning is conducted in an impartial and open-minded manner (other things being equal). Hence, it is not the biologist's intellectual virtues as character traits that are directly relevant to her belief amounting to knowledge or being justified, rather it is intellectual virtues conceived as properties of inquiry that are directly relevant.

III.1 - Virtuous Inquiry and the Problem of Strange and Fleeting Processes

One concern about my view is that it seems to suffer from some of the same problems that plague *simple process reliabilism.* Simple process reliabilism roughly holds that a belief is justified or amounts to knowledge if it was formed as the result of a reliable belief forming process. Greco (1999 and 2010) does a nice job of highlighting this problem. The problem is that simple process reliabilism is too weak because not all reliable processes can give rise to knowledge or justified belief. One major worry is that simple process reliabilism can't solve "The problem of strange and fleeting processes" and that on my conception of virtuous inquiry as a reliable belief-forming process, it will often count as a strange and fleeting process. For example, a person may inquire virtuously, but do so out of some unexplained urge or completely by accident or on a whim. That is, an agent might accidentally or unwittingly engage in virtuous inquiry, and even if engaging in such inquiry is normally reliable, we wouldn't ascribe knowledge or justification to this person's subsequent belief. Take Laurence Bonjour's (1980) now famous case of Norman the Clairvoyant as an example of the problem:

> Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against his belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable (p.62).

A common intuition is that Norman does not know (nor is he justified in believing) that the President is in New York City even though his belief was the result of his reliable belief-forming clairvoyant power. There are no reasons Norman is aware of for his belief, rather he just finds himself with this belief. Upon reflection, he finds that his belief is without support. However, in this case a reliable process forms Norman's belief. Yet Norman is unaware that it was formed this way. Hence, we are reluctant to say that Norman has knowledge or even justification in this case. The same thing allegedly applies to my conception of virtuous inquiry. In the same way that Norman's belief isn't justified or doesn't count as knowledge even though his belief is the result of a reliable belief-forming process, an agent's belief that was formed accidently or unwittingly by engaging in virtuous inquiry is also not justified.

In cases where a strange and fleeting process is used, the belief is formed by a reliable process but in some sense the reliability of the process is accidental, lucky, or something else. This would seem to apply to cases in which someone accidentally or unwittingly engages in virtuous inquiry.

Greco requires that we amend simple process reliabilism and require a *disposition* to use the process in question. In doing so, Greco argues we should adopt the following principle:

> **Agent Reliabilism:** a belief *p* has positive epistemic status for a person S just in case S's believing *p* results from stable and reliable dispositions that make up S's cognitive character.

Greco's view is called agent reliabilism because the dispositions referred to are dispositions that make up the agent's character (1999, p.287-88). If a belief is the result of the agent's stable and reliable dispositions, then this solves the problem of fleeting processes. Furthermore, these stable and reliable dispositions are not strange because they are part of the agent's intellectual/cognitive character.

The key to avoiding the strange and fleeting process objection is to ground the resulting belief in an agent's stable and reliable cognitive character. To say that a belief results from an agent's stable and reliable cognitive character is to say that the belief can be credited to the agent. That is to say that the belief is not the result of luck, but the result of his own cognitive abilities. Greco's account of credit attribution is as follows:

> S deserves intellectual credit for believing the truth regarding *p* only if
>
> **a.** Believing the truth regarding *p* has intellectual value.
> **b.** Believing the truth regarding *p* can be ascribed to S, and
> **c.** Believing the truth regarding *p* reveals S's reliable cognitive character. Alternatively: S's reliable cognitive character is an important necessary part of the total set of causal factors that give rise to S's believing the truth regarding *p*.
>
> And hence:
>
> > S knows *p* only if believing the truth regarding *p* reveals S's reliable cognitive character. Alternatively: only if S's reliable cognitive character is an important necessary part of the total set of causal factors that give rise to S's believing the truth regarding *p* (Greco, 2003, p.123).

If the belief is not creditable to the agent's stable and reliable cognitive character, then there is a risk that the belief is the result of some strange or fleeting process (or possibly luck or something else). On Greco's view, it would seem that my account of virtuous inquiry is missing this important credit component because virtuous inquiry is not creditable to someone with a stable vicious cognitive character. More precisely, the issue is that if virtuous inquiry is not expressive of or grounded in an agent's *stable* and reliable cognitive character, then any belief that results

from such inquiry cannot count as knowledge or be justified. Thus, it is Greco's stability requirement that is problematic for my view. I will address this problem below, but first I would like to highlight another potential problem for my view related to the present one.

Greco argues for an additional requirement in order for his agent reliabilism to deal with the problem of strange and fleeting processes. He argues that cases like the Norman case can be underdescribed regarding the issue of what he calls *cognitive integration* (Greco, 2010, p.152). That is, cognitive processes associated with clairvoyance are not sufficiently integrated with other of the person's cognitive dispositions (faculty virtues) so as to count as being part of cognitive character. A disposition is part of cognitive character only if it is both (a) stable in the relevant sense, i.e.—stable across close possible worlds and (b) well integrated with the agent's other cognitive dispositions.[21] It might be that clairvoyance is a process that produces only a single belief and it is completely unrelated and insensitive to other dispositions governing the formation and evaluation of belief. On the other hand, the senses, memory, introspection, reason, etc., are typically well-integrated with and sensitive to the agent's other cognitive dispositions.

There are two important aspects to cognitive integration: 1) if the products of a disposition are small in number and rare, and have little relation to other beliefs in the cognitive system, then the disposition is not well-integrated and 2) the beliefs in the system must be sensitive to counterevidence, or defeating evidence, i.e.—if the beliefs are insensitive to reasons that count against them, then they are not well-integrated. Greco writes, "In general, it would seem, cognitive integration is a function of cooperation and interaction, or cooperative interaction, with other aspects of the cognitive system" (Greco, 2010, p.152). This would mean

[21] Importantly, Greco is speaking of *human* cognition. That is, there may be cases where clairvoyant knowledge exists that involve cognitive agents who are very different from ourselves (Greco, 2010, p. 155).

that many strange and fleeting processes, although reliable belief-forming processes, are not well integrated with other aspects of the agent's cognitive system, and therefore the beliefs that result from such processes do not count as knowledge nor would they be epistemically justified.

What is of concern is whether this conception of cognitive integration also weighs against my view of virtuous inquiry. However, from this point on I will focus on the issue of justification rather than both justification and knowledge. This is because I'm ultimately concerned with the role that virtuous inquiry plays in determining the justificatory status of a belief. Even though Greco focuses on knowledge, his arguments equally apply to justification. I will argue that Greco's stability requirement is not necessary for a belief to be justified (I remain neutral about its role in knowledge). That is, I will reject the view that stable and reliable *dispositions* are an important necessary part of the total set of causal factors that justifies S's belief that *p*. However, I do agree that for a reliable process such as virtuous inquiry to play a salient role in determining the justificatory status of S's belief that *p*, it must be well-integrated with the rest of S's cognitive system. I will argue that even if virtuous inquiry does not reflect the person's stable and reliable dispositions, it is nevertheless well-integrated with the rest of S's cognitive system.

The remainder of part III will proceed as follows. First I will address the stability requirement issue and then I will address the cognitive integration issue. Additionally, examining these worries will allow me to flesh out some of the details of my account of virtuous inquiry and the role it plays in epistemic justification. However, a full account of my view will have to wait until later chapters.

III.2 - Virtuous Inquiry and Greco's 'Stability' Requirement

The first issue I will examine is whether my view of virtuous inquiry must satisfy Greco's stability requirement. The idea is that the agent's cognitive faculties and habits are by definition stable and reliable dispositions (Greco, 1999, p.287). So, any reliable process that reflects or reveals the agent's stable and reliable cognitive dispositions is not strange or fleeting because the cognitive dispositions of an agent themselves are not strange or fleeting processes. They are not strange because they are part of the agent's cognitive character. They are not fleeting because they are stable dispositions and are not the kind of thing an agent can adopt on a whim or engage in an irregular fashion. This would seem to weigh against my view since I claim that a normally intellectually vicious agent could (accidentally or unwittingly) engage in virtuous inquiry even though such inquiry isn't grounded in or reflective of that agent's stable and reliable character *and* that the reliability of that inquiry can still contribute to the justificatory status of the agent's belief.

As an example of a strange process, Greco considers the case of the 'Serendipitous Brain Lesion' (2010, p.149). Suppose an agent has a rare brain lesion that reliably causes the true belief that one has a brain lesion. Like clairvoyance, the worry is that even if the process is perfectly reliable, it doesn't seem that one can have knowledge or be justified in believing that one has a brain lesion on this basis. An example of a fleeting process is best exemplified in the case of the 'Careless Math Student' (Ibid). Suppose that a student is taking a math test and adopts a correct algorithm for solving a math problem. However, suppose that the student has no idea that the algorithm is the correct one but simply makes a lucky guess when choosing it, i.e.—he could just as easily have picked the incorrect algorithm. Using the algorithm to solve the math problem counts as a reliable process, but it would seem incorrect to say that the student

knows or is justified in believing the answer to the problem because there are nearby possible worlds where the student does not use the right algorithm. The algorithm is perfectly reliable, but the *student* is not reliable.

The reason Greco thinks that processes like clairvoyance, serendipitous brain lesions, and luckily chosen algorithms aren't reliable is because they are not properly grounded in the knower's stable cognitive virtues (or abilities, or excellences) (Greco, 2010, p.150). It would seem that in order for my view to completely avoid the fleeting processes objection, virtuous inquiry must also be grounded in the agent's stable and reliable cognitive ability. The math student's belief is not justified precisely because he himself is not reliable. That is, his use of the algorithm is not grounded in a stable and reliable cognitive ability and so there will be close possible worlds where he chooses the wrong algorithm instead of the correct one. A cognitive virtue by definition is stable and for them to be reliable, they must be stable across close possible worlds.

Now, the problem in the brain lesion case is that there is a lack of cognitive integration. As Greco puts it, "That is, the cognitive processes associated with the brain lesion are not sufficiently integrated with other of the person's cognitive dispositions so as to count as being part of cognitive character" (Greco, 2010, p.152). Therefore, a disposition is part of cognitive character only if it is both (a) stable in the relevant sense, and (b) well-integrated with other of the person's cognitive dispositions. Remember that there are a couple of key aspects of cognitive integration. The first aspect concerns the range of outputs: if the products that result from a disposition or process are rare, and if they have little or no relation to other beliefs in the system, then the disposition or process is poorly integrated. The second aspect of cognitive integration is sensitivity to counterevidence or to defeating evidence. In other words, if the

relevant beliefs aren't sensitive to reasons that count against them, then they are also not well-integrated.

Greco claims that cognitive integration is a function of cooperation and interaction (cooperative interaction) with other aspects or parts of the cognitive system (Ibid). Like the student's use of the algorithm, the brain lesion is not well-integrated with the rest of the cognitive system. It is a process that produces only a single belief (that one has a brain lesion) and it has no relation to other beliefs in the system. Also, the brain lesion is insensitive to other dispositions or processes involved in the formation and evaluation of belief.[22]

So how is this a problem for my view of virtuous inquiry as a reliable belief forming process? Take the following example:

> **Intellectually Vicious Vic:** Vic has an intellectually vicious character. He exemplifies the following intellectual vices: close-mindedness (fails to give consideration to competing evidence), unfair-mindedness (engages in selection-bias), sloppy when making observations, imprecise in his analyses, inconsistent in applying standards of evidence, and lacks patience and thus rushes through his inquiries.

These intellectual vices are not reliable belief-forming processes. They are generally unreliable in helping a cognitive agent get to the truth. Indeed, they are hindrances to getting to the truth. Further, these vices reflect Vic's stable and unreliable cognitive character. So, engaging in *virtuous* inquiry would not reflect Vic's stable and reliable cognitive character because his cognitive character is stable and *unreliable*. It is completely out of character for him to engage in virtuous inquiry. Hence, Vic's stable and *reliable* cognitive character is not a part of the total set of causal factors that give rise to his believing the truth regarding something like, say, climate change. I believe this is correct, but it doesn't impugn my view.

[22] There are worries about the reliabilist's ability to explain how reliabilism can be sensitive to defeating evidence. However, this requires a separate discussion and goes beyond the scope of this dissertation. For a more thorough discussion of the problem of defeating evidence for reliabilism, see Greco (2010) Chapter 10.

Greco's stability requirement isn't necessary in order for a reliable process to contribute to a belief's justification. The key to my view is that the reliability component of virtuous inquiry is found in the inquiry's success rate at uncovering *evidence* that is not misleading. Evidence plays a *primary* role in my account of epistemic justification. That is, what justifies a belief full stop, is the evidence that bears on that belief. However, if an agent's well-founded belief is based on evidence that he gathered as the result of virtuous inquiry (and there are no defeaters), the reliability of that inquiry can increase the degree to which that belief is justified by increasing the probability that the evidence is not misleading. Whether virtuous inquiry reflects the agent's stable and reliable character doesn't matter. Virtuous inquiry as a *process* is the only thing that needs to be reliable. In all nearby possible worlds, the process of virtuous inquiry still increases the probability that the evidence gathered by the inquiry is not misleading (other things being equal). This is the case regardless of whether the agent is reliable. The problem of strange and fleeting processes is not relevant to my account because my account requires evidence to play a primary role in a belief's justification and virtuous inquiry's role in determining the justificatory status of a belief is connected to its reliability in gathering non-misleading evidence. More straightforwardly, virtuous (reliable) inquiry has a higher probability of gathering non-misleading evidence (other things being equal) regardless of whether the agent's cognitive character is stable and reliable. Thus, Greco's stability requirement is not necessary for a process to contribute to a belief's justification.

Furthermore, for virtuous inquiry to contribute to the justificatory status of a belief, it doesn't matter whether engaging in virtuous inquiry is accidental or unintentional. Vic could completely lack any awareness of the virtuous nature of his inquiry and still inquire virtuously. He could also inquire virtuously by accident or on a whim. Virtuous inquiry is either reliable or

not. It doesn't matter whether inquiring virtuously was accidental or done on a whim. If the inquiry was virtuous, and therefore reliable, then the probability that Vic's evidence is not misleading is still increased (other things being equal).

Since, there is no stability requirement, there is no reason that a process can't be 'strange' (i.e.—not part of nor revealing of an agent's enduring intellectual/cognitive character) and still contribute to the justificatory status of a belief. Virtuous inquiry does not necessarily reveal an agent's stable and reliable cognitive character and can still be a reliable belief-forming process that plays a salient role in determining the justificatory status of that agent's belief. Additionally, virtuous inquiry can also be a 'fleeting' process (i.e.—the kind of thing an agent can adopt on a whim or engage in an irregular fashion) and still play a salient role in determining a belief's justificatory status. However, in order for virtuous inquiry to play a salient role in determining the justificatory status of a belief, I do think that as a reliable belief-forming process, it must be well-integrated with a person's other cognitive faculties. I believe it is. I turn to that brief discussion now.

III.3 - Virtuous Inquiry and Cognitive Integration

I see no reason to think that inquiring virtuously (whether out of character or not) is not well-integrated with the person's other cognitive faculties. Remember that Greco argues that many fleeting processes, although reliable belief-forming processes, are not well-integrated with other aspects of the agent's cognitive system, and therefore the beliefs that result from such processes are not justified (or do not count as knowledge). I agree, but I will argue that virtuous inquiry is a well-integrated cognitive process and so cognitive integration is not a problem for my view. That is, virtuous inquiry as a reliable belief-forming process is (or can be) properly

integrated with an agent's other cognitive processes, but that this is required for a different reason than the one Greco offers.

Virtuous inquiry is a reliable belief-forming process. For example, reasoning is a quintessential reliabilist faculty virtue. But it isn't simply reasoning that plays a role in determining whether a belief is justified, but *good* reasoning, which in many cases might be reasoning open-mindedly or reasoning carefully, etc. Reasoning open-mindedly and carefully is a process. You can't separate the 'open-mindedness' and the 'carefulness' from the reasoning. The idea is that reasoning open-mindedly would raise the justificatory status of certain beliefs because engaging in such reasoning is generally more reliable in gathering non-misleading evidence, compared to reasoning close-mindedly which is not reliable. This also applies to faculty virtues like vision. In many instances it is not completely accurate to say that vision is the reliable process that gathered some body of evidence, but *good* vision, and often times good vision just is visually examining something carefully and attentively. If you strip vision of the carefulness and attentiveness, its reliability will likely decrease.

Additionally, the process of virtuous inquiry has a wide range of outputs. That is, the products resulting from virtuous inquiry are many and varied. Virtuous inquiry generates a high yield of beliefs. Further, virtuous inquiry usually involves many different processes working together, i.e.—the senses, memory, various types of reasoning, etc. All of which are well-integrated with one another. That is, these different processes are not isolated from one another and work together during the process of inquiry. For example, reasoning open-mindedly is often well-integrated with using one's vision carefully and attentively.

Furthermore, the beliefs that result from virtuous inquiry are sensitive to counterevidence, or to defeating evidence. When we engage in virtuous inquiry, we carefully

check and double check data, consult colleagues, check proofs, check arguments for validity, cogency, and fallacies, make careful observations, conduct complex experiments, raise new questions, etc. Engaging in such activities often requires the full range of our cognitive system. Furthermore, the various processes involved in virtuous inquiry are related to one another and sensitive to the outputs of other processes governing the formation and evaluation of belief. For example, good inductive reasoning is sensitive to new data gathered from visual observation, good memory is well-integrated with sound deductive reasoning (remembering logical fallacies properly), and so on.

Now, if this picture is correct, then I will have succeeded in providing a plausible account of the role that the intellectual virtues play in epistemic justification. More precisely, I have argued that what we should be concerned with when trying to understand the role of virtue in epistemic justification is to focus our attention on virtues of inquiry rather than virtues of character. Too many problems beset analyses of knowledge and epistemic justification that appeal to Zagzebski-type character virtues. On the other hand, an analysis of knowledge and justification that makes sense of virtues as virtues of inquiry has the benefit of comporting well with certain reliabilist accounts of knowledge and justification. That is, virtuous inquiry is a reliable belief-forming process that plays a key role in determining the degree to which a belief is justified. Of course, my discussion of virtuous inquiry in this chapter only provides some of the material needed to provide a fuller account of how virtuous inquiry plays a salient role in the justificatory status of a belief.

IV. Vicious Inquiry

Throughout this chapter my focus has been on intellectual virtues and virtuous inquiry and the role that virtuous inquiry as a reliable belief-forming process plays in determining the

degree to which a belief is justified. This is because my main concern is with the role that virtuous (reliable) inquiry plays in epistemic justification. A thorough discussion of the nature of intellectual vices and vicious inquiry therefore goes beyond the scope of this dissertation. Thus, I will not discuss intellectual vices and vicious inquiry to the same extent as I've discussed intellectual virtues and virtuous inquiry. However, my view of justification also entails that vicious (unreliable) inquiry often plays a salient role in determining the degree to which a belief is epistemically justified. So, I would like to briefly say something about intellectual vices and vicious inquiry. Importantly, many of the same arguments I made in this chapter about intellectual virtues and virtuous inquiry apply *mutatis mutandis* to intellectual vices and vicious inquiry (with certain exceptions of course).

The intellectual vices are counterparts (largely) to intellectual virtues. Intellectual vices include, for example, close-mindedness, self-deception, obtuseness, hastiness, unfair-mindedness, etc. Just like the intellectual virtues, I conceive of the intellectual vices relevant to epistemic justification as properties of inquiry rather than as traits of character. In the 'Intellectually Vicious Vic' example Vic inquires viciously. That is, his inquiry is conducted close-mindedly (fails to give consideration to competing evidence) and unfair-mindedly (engages in selection-bias), he is sloppy when making observations, imprecise during his analyses, inconsistent in applying standards of evidence, and lacks patience and thus rushes through his inquiries. Vicious inquiry is an *un*reliable belief-forming process that often plays a salient role in determining the degree to which a belief is justified. Again, the idea is that when a person inquires viciously, that person engages in flawed or unreliable evidence-gathering methods. Evidence gathered as the result of unreliable evidence-gathering methods is more likely to be misleading evidence. The greater the probability that a body of evidence is misleading, the less

justified is the belief that is supported by that evidence. Thus, vicious (unreliable) inquiry often plays a salient role in decreasing the degree to which a well-founded belief is justified. Importantly, however, vicious inquiry cannot render a well-founded belief unjustified. From here on out, most of my focus will tend to be on virtuous inquiry.

V. Virtuous Inquiry as a Reliable *Evidence-Gathering* Method

Thus far I have been discussing reliable *belief-forming* processes and have deemed virtuous inquiry such a process. However, as I indicated in the last section, it is more accurate to say that virtuous inquiry is a reliable *evidence-gathering* process or method. This is due to the fact that virtuous inquiry plays a very specific role in my account of justification. That is, inquiring virtuously increases the probability that the *evidence* gathered is not misleading and therefore the degree to which a well-founded belief is justified is increased. The opposite applies to vicious inquiry. I make this point because from now on I will refer to virtuous inquiry as a reliable evidence-gathering process or method and vicious inquiry as an unreliable evidence-gathering process or method.

VI. Conclusion

In Chapter Two I distinguished between two different kinds of intellectual virtue properties—roughly, properties of inquirers (e.g., being open-minded) vs. properties of inquiries (e.g., being conducted open-mindedly). I raised objections to theorizing justified belief in terms of "inquirer-focused" properties. I argued that these objections do not foreclose the possibility of theorizing justified belief at least partly in terms of "inquiry-focused" properties. Further, I gave a rough, preliminary statement of an evidentialist view of justified belief in which the exemplification of "inquiry-focused" properties can increase or decrease the degree to which a doxastic attitude is justified.

In Chapter Three I will discuss the relationship between epistemic rationality, instrumental or practical rationality, and methodological epistemic rationality. I will defend the view that epistemic rationality (justification) is distinct from instrumental rationality and that methodological epistemic rationality (virtuous inquiry being a type of methodological epistemic rationality) is not merely a form of instrumental rationality (as Feldman suggests) but is also an important source of epistemic rationality (justification). This is important for understanding how virtuous inquiry is relevant to epistemic justification. I will then argue that epistemic rationality is best understood by appealing to a diachronic approach to epistemic justification rather than a synchronic approach. This discussion will begin to motivate and defend evidentialist necessary and sufficient conditions for *categorical* epistemic justification (which I will continue to motivate and defend in much greater detail in Chapter Four), as well as present arguments that a conception of evidentialism as a theory of categorical justification is compatible with my two-component theory of epistemic justification. More precisely, I will begin to motivate my two-component view of justification which holds that an evidentialist theory of categorical justification is compatible with the fact that reliable evidence-gathering methods (e.g., virtuous inquiry) can play a salient role in increasing the degree to which a belief is justified.

.

CHAPTER THREE

Epistemic Rationality

I now wish to switch gears and begin to evaluate two issues of fundamental concern to my project. One fundamentally important issue to sort out is what kind of rationality we are talking about when we discuss *epistemic rationality*.[23] First, I will examine a view of epistemic rationality that maintains that *epistemic* rationality *just is* a species of *instrumental* or *practical* rationality. I will conclude that epistemic rationality is not a species of instrumental rationality. In doing so, we will see that there is a plausible evidentialist conception of epistemic rationality (what Feldman calls 'current-state epistemic rationality') that clearly demarcates itself from instrumentalist rationality. Second, I will discuss whether virtuous inquiry (what Feldman calls '*methodological epistemic rationality'*) is an important source of *epistemic rationality* or is merely a species of instrumental rationality. I will argue for the former view. Importantly, I will not be arguing that there are two types of epistemic rationality, but that there are (at least) two things that contribute to epistemic rationality, i.e.—the evidence supporting a belief and the reliability (or unreliability) of the inquiry used to gather that evidence.

The other fundamentally important issue I will discuss, which falls out from the first issue, is whether theoretical approaches to epistemic rationality are best conceived as *synchronic* or *diachronic.* I will begin to defend a diachronic approach to epistemic rationality that contains both evidentialist and reliabilist components. I will then further develop this two-component theory of epistemic rationality/justification in the next two chapters.

[23] Throughout this chapter I will use the phrase 'epistemic rationality' which I take to be equivalent to 'epistemic justification'.

Let me begin to motivate the discussion by considering once again the biologist example from the last chapter:

> A field biologist is trying to explain a change in the migration patterns of a certain endangered bird species. Collecting and analyzing the relevant data is tedious work and requires a special eye for detail. The biologist is committed to discovering the truth, and so spends long hours in the field gathering data. He remains focused and determined in the face of various obstacles and distractions (e.g., conflicting evidence, bureaucratic roadblocks, inclement weather conditions, boredom, etc.). He picks out important details in environmental reports and makes keen discriminations regarding the composition and trajectory of several observed flocks. As a result of his determination and careful methods of enquiry, he discovers why the birds have altered their course.

In this case, the biologist forms a belief about the migratory behavior of a certain bird species. This belief is no doubt formed on the basis of some very plausible evidence. However, it would seem that another important aspect of this case is that the biologist forms his belief as the result of good epistemic methods (virtuous inquiry), i.e.—perseverance, tenacity, careful observation, etc. One way to look at this is that the methods the biologist employs in his research actually *expand* the type and amount of evidence he has. It appears that without these good epistemic methods, the biologist would not have gathered the strong evidence he now has. It would seem that the good research methods he used contributed to the rationality of his belief. To see this better, imagine the biologist used *bad* epistemic methods, i.e.—laziness, irresoluteness, reckless observation, etc. It seems obvious that the data he collected as a result would very probably be defective and/or restricted. Consequently, we would say that his belief is *irrational* as a result of employing bad epistemic methods.

While the foregoing remarks seem plausible initially, Feldman would argue that this paints an incorrect picture of the nature of epistemic rationality. Feldman distinguishes between two senses of epistemic terms such as 'rational' (Feldman, 2004b, p.233). One type of epistemic appraisal concerns whether a person is rational in believing a proposition at an exact time and

under the exact conditions the person is in at that precise time. Feldman says that this is an assessment of the person's *current-state epistemic rationality* of believing the proposition. Another possible type of epistemic appraisal of a belief concerns the methods that led to the formation of that belief. Feldman calls this *methodological epistemic rationality*. Beliefs are methodologically rational if and only if their formation results from good epistemic methods (such as virtuous inquiry).

In the version of the biologist case where he employs bad epistemic methods, it might be that he is methodologically irrational. That is, one can contend that the biologist should have been more persistent, tenacious, and careful in his observations. If he would have done so, he would have gathered better evidence. Because the biologist did not do these things, he is *methodologically* irrational. On the other hand, Feldman would argue that it is reasonable to maintain that the biologist's belief is *current-state rational*. If the biologist has not *considered* whether the methods he is using are good or bad, it would be quite irrational for him to believe anything else. In other words, the evidence the biologist gathered clearly supports his belief (maybe that there is no conclusive explanation why the birds altered their course). Given the situation the biologist is in, the resulting belief is precisely what he should (epistemically) hold.

This brings up an important issue: if virtuous inquiry is just methodological epistemic rationality (as construed by Feldman), then how could it play a crucial role in determining the justificatory status of a belief? On Conee and Feldman's evidentialist view, the only thing that is relevant to whether a belief is justified is how the evidence bears on a belief, i.e.—whether the belief fits the evidence one has. 'Justification' claims about methods of inquiry are independent of whether the belief is justified under the actual conditions at a given time (Conee and Feldman, 2011, p.313).

There are two important issues related to this distinction that need further examination. The first issue concerns whether virtuous inquiry amounts to merely methodological rationality, or what is often called *practical* or *instrumental* rationality, and whether such rationality is different in kind from the type of *epistemic* rationality important for determining a belief's justificatory status. This brings up the second issue which is whether an account of epistemic justification should be construed as *synchronic* or *diachronic*. Roughly, a synchronic account of epistemic justification holds that a belief is justified if it is based on good grounds at the time it is held (current time-slice justification). In contrast, a diachronic account of epistemic justification roughly holds that a belief is justified if it results from responsible and/or reliable methods, proper practice, etc., leading up to the belief in question (cross-temporal justification).[24]

I will examine both of these issues in turn. I will begin by examining the relationship between *epistemic rationality* and *instrumental rationality*. In doing so, I will look at works by Thomas Kelly (2003; 2007) and Adam Leite (2007). At issue is whether epistemic rationality *just is* instrumental rationality or whether they represent two separate and distinct types of rationality. I will argue for the latter view. I will then argue that although virtuous inquiry has a strong instrumental component, i.e.—is largely a means-end or goal oriented process, its relevance to epistemic normativity (in this case whether a belief is epistemically justified) is due to its *reliability* as a reliable belief-forming process. As a result, I will propose that if belief formation involves virtuous inquiry, then the belief can derive its degree of epistemic justification from two sources: 1) the *strength* of the evidence that supports that belief and 2) the *reliability* of virtuous inquiry (vicious inquiry can decrease the degree to which a belief is justified). In making this argument I will examine works by Feldman (2000; 2004a), Conee and

[24] Later in this chapter, I will provide a more specific and formal definition of what constitutes synchronic and diachronic accounts of epistemic justification. For now I offer these rough definitions.

Feldman (2011), Trent Dougherty (2010; 2011; 2012), Guy Axtell (2011; 2012), Keith DeRose (2011), and Mylan Engel (1992).

I. Epistemic Rationality and Instrumental Rationality

Before digging into part I, let me say a bit more about what I just said in the last section. I take Feldman's *current-state epistemic rationality* to be equivalent to what Thomas Kelly calls *epistemic rationality.* Kelly defines epistemic rationality as the type of rationality a person exhibits when that person believes propositions strongly supported by that person's *evidence* and refrains from believing propositions that are improbable given that person's *evidence* (Kelly, 2003, p.612). Kelly's example is that if I have strong, undefeated evidence that the butler committed the crime, and my belief that the butler committed the crime is based on that evidence, then my belief is epistemically rational. Notice that this is an *evidentialist* conception of epistemic rationality (justification). It is different from what responsibilists and reliabilists consider epistemic rationality. The responsibilist defines epistemic rationality as (roughly) belief that is formed responsibly. The reliabilist defines epistemic rationality as (roughly) belief that is reliably formed. For the evidentialist, it is how evidence bears on a belief that *solely* determines its epistemic rationality. It is this conception of evidentialist epistemic rationality that is in question in this chapter.

I will also argue that *methodological epistemic rationality* is not equivalent to what is referred to as *instrumental rationality.* Instrumental rationality is the type of rationality one displays in taking the means to one's ends (Kelly, 2003; 2007; Leite 2007; Lockhard 2013). As Kelly puts it, "Thus, if I have the goal of asking the speaker a question, and I know that I will only be able to ask the speaker a question if I raise my hand, then (all else being equal) it is instrumentally rational for me to raise my hand" (Kelly 2007, p.612). This implies that it's

instrumentally rational for a person S to do X if and only if S *believes* that doing X will help S achieve some goal S has. However, I will argue that instrumental rationality doesn't exhaust the ways in which the *methods* or *processes* used to form a belief contribute to a belief's rationality. More precisely, I will argue that the methods or processes used to form a belief can contribute to both instrumental and epistemic rationality.

Thus, according to methodological epistemic rationality, a belief is rational if and only if its formation results from good epistemic methods, regardless of whether the person has any epistemic goal or believes that doing X will help S achieve some epistemic goal (in this case, the formation of true beliefs). Instrumental rationality is a practical or means-end rationality, whereas methodological epistemic rationality is akin to a reliable process like virtuous inquiry. This doesn't mean that a reliable process isn't utilized as a means to one's end. The key here is that a reliable process isn't *merely* a means to one's end. A reliable (or unreliable) process might also contribute to the justificatory status of a belief. Thus, the type of rationality relevant to merely meeting one's ends is practical or instrumental rationality, but this is not the sort of rationality relevant to epistemic rationality (or so I will argue). From this point on I hold firm to this important distinction between instrumental rationality and methodological epistemic rationality.

Given the preceding discussion, there are two main issues that need examination in part I. The first issue is whether epistemic rationality *just is* instrumental rationality. That is, I wish to examine whether epistemic rationality is just a species of instrumental rationality. In doing so I will defend Kelly's conception of epistemic rationality from instrumentalist conceptions of rationality. The second issue is whether epistemic rationality and methodological epistemic rationality constitute two *different* types of epistemic justification. I will argue that this is a false

distinction and that there is only one type of epistemic rationality. However, as I've said, I will argue for a different conception of epistemic rationality than the one defended by Kelly and Conee and Feldman. I will defend a conception of epistemic rationality in which evidence and virtuous inquiry both play a salient role in determining the justificatory status of a belief.

I.1 - Thomas Kelly's Argument

Kelly examines the relationship between epistemic and instrumental rationality (as defined above). He asks whether epistemic rationality is just a species of instrumental rationality, i.e.—instrumental rationality in the service of one's *cognitive* or *epistemic* goals (2003, p.613). This is known as the *instrumentalist conception of epistemic rationality*. Kelly asks what the relationship is between the rationality one displays in responding to evidence in the epistemically appropriate manner, and the rationality one displays in *acting so as to acquire* that evidence. For example, if I want to know the identity of the person who committed the crime (the butler) I will need to engage in the activity of *looking for evidence* which bears on the question. I have the goal of learning a certain truth (who committed the crime), which gives me an *instrumental reason* to act in a certain way, i.e.—other things being equal, it is rational for me to engage in the activity of looking for evidence. This is an example of instrumental rationality in the service of one's cognitive goals. Suppose further that I'm able to uncover strong evidence that the butler committed the crime. "The character of this evidence singles out a certain response on my part as the epistemically rational response: it is rational for me to believe that the butler committed the crime" (Ibid). The instrumentalist wants to assimilate the rationality of responding to the evidence in the epistemically appropriate way to the rationality of looking for that evidence in the first place. Kelly thinks it is a fundamental mistake to think about epistemic rationality in this manner.

Kelly thinks that the most serious reason to reject the instrumentalist conception of epistemic rationality is that what a person has reason to believe does not seem to depend upon the content of any goals he/she might possess. Having an instrumental reason is to have the relevant goal. I have a reason to raise my hand *because* I have the goal of being called on by a speaker. Without such a goal, there would be no reason to raise my hand. In this way, an instrumental reason is a *hypothetical* reason in the sense that it depends for its existence on the fact that the individual for whom it is a reason possesses a particular goal or goals. Kelly thinks this contrasts with the *categorical* character which *epistemic* reasons seem to have. By 'categorical' I take Kelly to mean that epistemic reasons do *not* depend for their existence on the fact that the individual for whom they are reasons possesses a particular goal or goals.[25] According to the instrumentalist conception of epistemic rationality, facts about what I have reason to believe are contingent on my having specific goals (Kelly, 2003, p.621).

Kelly rejects the instrumentalist conception. He argues that whether we have epistemic reasons to believe something has nothing to do with having particular goals. Kelly says, "However, from the fact that some subjects are matters of complete indifference to me, it does not follow that I will inevitably lack epistemic reasons for holding beliefs about those subjects" (Kelly, 2003, p.621). Kelly points out that there are many subject matters for which many people have no preference for having true beliefs, as well as subjects in which people might prefer having no beliefs whatsoever. Take the following example by Kelly:

> Thus, I tend to see newly-released movies after many of my friends. During the interval of time which is bounded on one side by my friends' viewing of the movie and bounded on the other side by my viewing the movie, I often make a conscious, deliberate effort to avoid finding out how the movie ends—since doing so might very well interfere with my

[25] I don't want the reader to confuse the way Kelly uses the term 'categorical' with how I have been using the term. On my view, 'categorical' refers to beliefs that are justified absolutely, full-stop, period, etc. However, on my view (like Kelly's view), whether a belief is supported by the evidence does not depend on the fact that the individual possesses a particular goal or goals.

> enjoyment when I do see it (When conversations about the movie begin in my presence, I either excuse myself or, reminding the discussants that I have yet to see the movie, implore them not to "give away" the ending, and so on.) That is, I quite deliberately take steps to avoid acquiring information about the movie. Sometimes these efforts are successful, sometimes they are not. When they are unsuccessful—as when someone inconsiderately blurts out the ending in my presence—it does not follow that I have no epistemic reasons to believe the propositions which he asserts (Kelly, 2003, p.626).

On Kelly's view, whether he has the goal of finding out how the movie ends makes no difference as to whether he is epistemically justified in believing how it ends once someone inconsiderately blurts out the ending. To think this way is to conflate (1) the reasons which one may or may not have to *seek out further evidence* which bear on the truth of some proposition, and (2) the reasons which one may or may not have to believe that proposition.[26]

In defense of this, Kelly says that in avoiding finding out how the movie ends, he wants to avoid acquiring reasons for believing the truth about how the movie ends. However, he says that if the possibility of acquiring reasons for believing the truth of some proposition is contingent upon having some goal that would be better promoted by believing the truth about that proposition, then this project is incoherent, i.e.—"…there is no need to deliberately avoid the acquisition of epistemic reasons to believe propositions about subjects with respect to which one has no desire to believe the truth, for one knows *a priori* that there are no such reasons. (Indeed that there *could not be* such reasons)" (Kelly, 2003, p.628). But, Kelly points out that this project is not incoherent.

Kelly says that he might have epistemic reasons to believe the truth about how the movie ends despite not having the relevant goal of finding out the ending. This is apparent when one acquires an unwanted belief—because a friend blurted out the end of the movie or by accidentally reading about it in a magazine headline. We would explain that forming a belief

[26] This is similar to the point made by Sosa and Sommers in Chapter Two.

about how the movie ends is not just a matter of psychological compulsion. That is, it makes sense to *explain why* the unwanted belief was formed by *citing* one's epistemic rationality, along with the fact that the person was presented with the relevant epistemic reasons. However, we can't make sense of this if the instrumentalist conception of rationality is correct. Kelly writes, "Put simply: one cannot *immunize* oneself against the possibility of acquiring reasons for belief by not caring about the relevant subject matter" (Ibid). The idea is that I come to have reasons to believe how the movie ended whether I wanted them or not. Thus, with respect to the issue of whether epistemic rationality *just is* a species of instrumental rationality, Kelly says:

> When it is instrumentally rational for me to Φ, this is because Φing promises to promote some goal or goals which I possess. The attempt to assimilate epistemic rationality to instrumental rationality founders on the fact that one can have epistemic reasons to believe propositions even in cases in which it is clear that one's believing those propositions holds no promise of advancing any goal which one actually possesses" (Kelly, 2003, p.630).

Being epistemically rational is not contingent on the goals that one possesses. Our goals are what motivate us to inquire into some subject matter in the first place, but we can acquire epistemic reasons to believe some proposition even if we lack the relevant goal to find out the truth of that proposition.

I.2 - Adam Leite's Reply to Kelly

Although Adam Leite is not a committed instrumentalist, he offers a response to Kelly's argument on behalf of the instrumentalist. The first thing Leite points out is that Kelly's evidentialist conception of epistemic rationality is one that the instrumentalist need not accept in the first place. He argues that Kelly's characterization is question-begging in one important respect: "The instrumentalist need not accept that this is a kind of rationality at all—except insofar as one has the goals (whatever they might be) which would make such behavior instrumentally rational" (Leite, p.457). Instead, Leite offers a more 'neutral' characterization

which involves the question, "What ought I believe?". He argues that this question can be answered only given a specification of one's goals. Leite agrees with Kelly that in order to provide an all things considered answer to the question "What ought I believe?", one would have to specify one's other goals and one's preference rankings for which goals one wants satisfied the most. And in order to provide a purely *epistemic* answer, one would have to specify one's cognitive goal (if any) (Leite, p.457; Kelly, 2003, p.619-620). With this in mind, Leite offers the following instrumentalist conception of rationality:

> …the instrumentalist can offer an explanatory story roughly along the following lines: given that one has such-and-such cognitive goals, the rational thing for one to do is to believe (only) propositions that are strongly supported by one's evidence—since that's the best way for one to achieve one's cognitive goals (Leite, p.657).

What Leite argues is that Kelly conflates questions about evidence with questions about what one ought to believe. In other words, we cannot explain epistemic rationality without appealing to cognitive goals. Kelly on the other hand, thinks that evidence can provide reasons for *anyone* to believe a proposition regardless of their goals.

In order to understand Leite's response to Kelly, we need to look at one of Kelly's examples:

> If both of us know that all of the many previously-observed emeralds have been green, then both of us have a strong reason to believe that the next emerald to be observed will be green, regardless of any differences which might exist in our respective goals … in arguing for my conclusions…, I think of myself as attempting to provide strong reasons for believing my conclusions, and *not* as attempting to provide strong reasons for believing my conclusions for those who happen to possess goals of the right sort (Leite, p.458; Kelly, 2003, p.621).

Leite says that if you and I both know that all the previously-observed emeralds were green, then we both know something that provides *strong evidence* that the next emerald we both observe will also be green. He says that the instrumentalist can hold this view regardless of differences that might exist between our respective goals. Leite goes on to say that "…the instrumentalist

can plausibly hold that this shared evidence provides me with *a reason to believe* that the next emerald will be green *only insofar as I have a goal which would be served by believing what the evidence supports about this matter*" (Leite, p.658). He has us imagine someone who simply has no project at all to which the color of the next emerald would be relevant. Leite asks whether this person is being irrational—not doing what he has most reason to do—if given the evidence before him he simply fails to form the belief that the next observed emerald will be green. Leite says this person is not being irrational.

Leite makes the following comparison:

> I'm standing next to a door at a convention center. I idly notice that all of the many people I've seen come out of the door have been accompanied by dogs. I am certainly not being irrational if I fail to form the belief that the next person to come out of the door is likely to be accompanied by a dog. Perhaps I'm busy thinking about things of greater interest or importance to me. More generally, as I go through my day, I gain all sorts of evidence supporting all sorts of beliefs. But I don't form most of them, and it hardly seems plausible that I have any reason to do so, given that they are about matters of complete indifference to me (Leite, p.458).

Leite agrees with Kelly that in arguing for our conclusions we regard ourselves as providing reasons for *anyone*. However, Leite claims that when we argue for a conclusion, we think the reasons we offer are good ones for anyone *concerned* to form a true belief about the subject in question. But he says he doesn't think that just *anyone* must be concerned with forming a belief about the subject in question. Even if that person is in possession of good evidence to form a true belief, we wouldn't think him irrational if he doesn't form a belief on the subject in question. The reason for this is because we don't think that at that moment the person has any particular *reason* to form a belief on the subject in question (Leite, p.459).

Take Kelly's example of someone trying to avoid finding out how some movie ends. Leite argues that an instrumentalist *can* coherently describe such cases. On an instrumentalist picture, what this person is trying to do is avoid acquiring *evidence about* how the movie ends.

Leite thinks that there is no incoherence in this project since the existence of evidence about this matter is not contingent upon the person's possessing some goal which would be better promoted by believing the truth about the movie. The instrumentalist can argue that in fact, even though the person has good evidence about how the movie ends, he does not have reason to believe how the movie ends, inasmuch as doing so would not serve his goals, cognitive or otherwise (Leite, p.659-70). Leite's argument hinges on the following claims:

> Having epistemic reasons to believe that *p* is having epistemic reasons to *do something*: to form a belief on a certain subject matter. But the mere possession of strong evidence doesn't give me a reason to form a belief. There are many propositions which it would be pointless for me to bother to take any attitude towards, even though I possess strong evidence in their favor. I can't justly be charged with irrationality—even with epistemic irrationality—for failing to do so. Even if I have evidence which could serve as an epistemic reason for believing *p, I* don't have any reason to believe that *p* (Leite, p.461).

In order for someone to have a reason to believe that *p,* it's not enough just to have (strong) evidence of *p.* One must have a reason to believe what the evidence supports and that requires a goal or interest in the truth about *p.* That is, even if the person possesses strong evidence that *p* is true and even recognizes he possesses such evidence, he still might not have any reason to believe that *p* is true.

I.3 - Kelly's Reply to Leite

Kelly thinks that the distinction between having evidence for *p* and having a reason to believe *p* would be a useful distinction for the instrumentalist. However, he is skeptical that such a distinction can be made. Instead, Kelly offers an alternative view of the nature of evidence in which *evidence* is itself a normative concept, just like *reason for belief* (Kelly, 2007, p.467). As an example, he considers a standard Bayesian explication of evidence, according to which to regard something as confirming evidence for some hypothesis is to regard it as a reason to increase one's confidence that the hypothesis is true, while to regard something as disconfirming

evidence is to regard it as a reason to decrease one's confidence. Given such an explication, Kelly says, "…there is an internal connection between recognizing something as evidence and recognizing it as a reason to change one's present view" (Kelly, 2007, p.467-68).

Kelly claims that on *any* view according to which evidence is a normative concept, there is no distinction between having evidence that *p* is true and having reasons to believe that *p* is true. When offering evidence for *p,* one is *ipso facto* providing *reasons* to believe that *p* is true. Kelly points out that on Leite's view the following should not seem paradoxical:

> *I have overwhelming evidence that p is true. But I have no reason to believe that p is true* (Kelly, 2007, p.468).

Kelly thinks that there is something very wrong with this.

Kelly suggests that Leite's examples all have something in common. They are constructed such that it seems natural that the person has not *actually* considered the relevant proposition (Ibid). For example, in the convention center case above, the person is described as being 'busy thinking about things of greater interest or importance'. Kelly says, "It is perhaps unsurprising that we hesitate to convict an individual of irrationality for not believing a proposition that is seemingly so far removed from his or her attention" (Ibid). However, as Kelly points out, the instrumentalist is not concerned with what the person is *attending to*, but whether the person possesses some relevant goal with respect to the subject matter. Instead, we should consider someone who *explicitly* attends to *p* in that instance, is fully aware of the strong evidence in his possession that *p* is true, yet does not take up the belief that *p.* Kelly thinks it is far more difficult to avoid the verdict that the person has failed to respond rationally, *even if it is stipulated that he has no goal whatsoever* which would be better achieved by believing the truth about whether *p.* Kelly says, that on Leite's view, it is mysterious why an explicit judgment to the effect that one's evidence strongly supports *p* normally results in a belief that *p* is true.

I.4 - Instrumental Rationality vs. Methodological Epistemic Rationality

We have seen that instrumental rationality is the type of rationality one displays in taking the means to one's ends. It involves goal-oriented or goal-directed activity on the part of the epistemic agent. This goal-directed activity is for the purpose of achieving some goal that a person has—in this case, an epistemic goal. However, as I mentioned earlier, instrumental rationality doesn't exhaust the ways in which the *methods* or *processes* used to form a belief contribute to the belief's rationality. According to methodological epistemic rationality, a belief is rational if and only if its formation results from good epistemic methods*, regardless of whether the person has any epistemic goal or believes that doing X will help S achieve some epistemic goal* (in this case, the formation of true beliefs or gathering non-misleading evidence). That is, the person need not have any goal in order for a belief to have methodological epistemic rationality/irrationality. A belief can be reliably formed and evidence can be reliably gathered, regardless of whether the person knew she was using methods that would (probably) achieve her goal of forming a true belief or gathering non-misleading evidence. That is, a subject who's not explicitly aiming to form true beliefs or gather non-misleading evidence could nevertheless do so. Of course, whether this picture of rationality is correct requires more discussion.

Furthermore, methodological epistemic rationality doesn't create a problem for the type of epistemic rationality that Kelly endorses *if* we take his view to hold that a belief is justified full-stop if there is good evidence to support it.[27] On this view, if a person's belief that *p* at a given time *t* fits the evidence the person has at *t*, then that person has *reason* to believe *p*, regardless of whether the methods she used to form that belief were epistemically good methods. This is precisely the view I'm defending. Additionally, on my view methodological epistemic

[27] Of course, it's doubtful that Kelly would be sympathetic to my hybrid view.

rationality, or what I call virtuous inquiry, is compatible with this evidentialist view of categorical epistemic justification. I will discuss this point in the section below.

I.5 - Two Components to Epistemic Rationality?

I take Kelly's argument against the instrumentalist conception of epistemic rationality to be convincing. However, I'm interested in exploring whether epistemic rationality can only be cashed out in evidentialist terms. That is, on an evidentialist view of epistemic rationality, it is solely the evidence bearing on a person's belief that makes believing *p* epistemically rational. I think that the evidentialist is correct insofar as evidence is a normative concept. My aim in this chapter is to determine whether evidence is the *only* normative component to epistemic rationality. I contend that it isn't.

I have already argued in Chapter Two why I think virtue responsibilist accounts of epistemic normativity are unsuccessful. Instead, I argued that a conception of intellectual virtues in which these virtues are also properties of inquiry rather than just traits of character, offers a far more promising account of how intellectual virtues (and intellectual vices) figure into determining the justificatory status of a belief. I argued that virtuous inquiry is a reliable evidence-gathering process, the reliability of which plays a prominent role in determining the justificatory status of beliefs. But in order for this argument to be successful, I have to show that all epistemic normativity isn't simply reducible to evidential fit. Doing so will go a long way in helping to provide a plausible account of the direct role that virtuous inquiry plays in epistemic justification.

Although I agree with Kelly's view about the normativity of evidence, it is crucial to point out that his view does *not* rule out the possibility that the *reliability* of virtuous inquiry can help determine the degree to which a belief is rational/justified. Specifically, although virtuous

inquiry is a means-end or goal directed process, the resulting rationality is not necessarily merely instrumental.

Someone who engages in virtuous inquiry is concerned with discovering some truth or fact of the matter. A person inquires to uncover information of interest to that person. Virtuous inquiry is a means-end process and a person uses virtuous inquiry as a tool to secure some end. Thus, virtuous inquiry has an important practical or instrumental dimension to it. For example, a scientist who engages in virtuous inquiry is engaged in a process that is an effective means to his/her end. The biologist engages in virtuous inquiry to understand the migratory pattern of a certain bird species. It is *instrumentally rational* for the biologist to engage in virtuous inquiry if he is interested in achieving his goal of understanding the migratory pattern of this bird species.

On Kelly's evidentialist view, however, what makes the biologist's belief *epistemically rational* is that the evidence he uncovered as the result of his virtuous inquiry gives him an *epistemic* reason to believe that the bird species has a particular migratory pattern. That is, the belief fits the evidence that *p,* in this case that the bird species has migratory pattern X. On this conception of epistemic rationality, there is no other source of epistemic rationality except for evidential fit. Virtuous inquiry may well have been an effective means to attain the evidence that the biologist needed, but this is merely instrumental rationality. At least on this evidentialist picture.

Again, I think that Kelly's argument (if sound) only establishes that epistemic rationality is not just a species of instrumental rationality. It establishes that evidence is a normative concept and provides a reason for belief *regardless of one's aims or goals*. But what Kelly's argument does not establish is that *evidential fit* is the *only* normative component to epistemic rationality.

Virtuous inquiry is certainly an instrumental tool for reaching one's goal of gathering good evidence, reaching the truth, etc., but its instrumentality is not its epistemic normative component. Again, its normative component is its reliability in uncovering good evidence and reaching the truth. Now, it may seem that I'm begging the question with respect to the reliability of virtuous inquiry, i.e.—that reliability is a robust normative concept in epistemology. However, that is not what I wish to argue for at this point. Instead, I simply wish to argue that *if* reliability can help determine whether and/or the degree to which a given belief is epistemically justified/rational and *if* virtuous inquiry possesses such reliability, then virtuous inquiry could obviously play a prominent role in determining to what extent a belief is epistemically rational/justified. Again, Kelly's view does not rule this out.

Key to this view is that epistemic rationality need not have anything to do with reliability. That is, I do not claim that if there is epistemic rationality, then there is reliability. However, I do think that in order to be epistemically rational/justified (whatsoever) one's belief must fit one's total evidence. More generally, in order for a belief to be epistemically justified it must have good reasons to support it. Practical reasons are not epistemic reasons. Epistemic reasons to believe some proposition are evidence. Such evidence is truth indicative. As Kelly says, "…one's evidence is *what one has to go on* in arriving at a view" (Kelly, 2008, p.942). McCain says that in order for any items of evidence to be something that one has to go on in arriving at a view, one has to be *aware* of that evidence in some sense (McCain, p.10).[28] He says that it would seem strange to say that "S knows that *p,* but she also has no evidence for believing that *p* is true" (McCain, p.1). The same thing applies to justification. On this view, evidence is a primary and necessary source of epistemic normativity, but it need not be the *only* source. This

[28] I will discuss what evidence is in Chapter Four. Very briefly, all evidence consists of mental states. Of course, not all mental states are evidence.

means that epistemic rationality/justification comes in *degrees*. That is, if evidential support is a primary and *necessary* source of epistemic rationality, then 1) a belief is categorically justified (justified period or full stop) if it is supported by the person's evidence, which means that 2) any additional sources of epistemic normativity will *increase* the belief's degree of epistemic rationality/justification.

Of course, the *strength* of one's evidence also largely contributes to the degree to which a belief is justified (assuming there is no defeating evidence). That is, two people can have different bodies of evidence in which one body of evidence provides strong support for belief *p* and the other body of evidence provides weaker support for belief *p*. The two people have belief's that are both supported by the evidence and therefore both beliefs are epistemically rational/justified full stop, but to different degrees because of the varying strength of each person's evidence. But this doesn't mean that the strength of one's evidence is the only thing that determines degrees of justification. For example, suppose there are two biologists. Biologist A has gathered evidence X, while Biologist B has gathered evidence Y. But, Biologist A engaged in virtuous inquiry while gathering his evidence and Biologist B engaged in vicious inquiry while gathering his evidence. Suppose that both bodies of evidence equally support *p.* That is, absent any other evidence and as far as anyone can tell, both bodies of evidence provide equally compelling reasons to believe *p.* My intuition is that Biologist B's belief *p* is epistemically rational/justified given his body of evidence, but that Biologist A's belief *p* has a higher degree of epistemic rationality/justification because he gathered his evidence by engaging in virtuous inquiry, i.e.—using a reliable evidence-gathering process. Biologist B used an unreliable (vicious) evidence-gathering process and therefore the degree to which his well-founded belief is justified is diminished.

What is crucial to my claim is that there can be strong supporting evidence for propositions which are in fact false. That is, evidence is fallible. We can have what we consider strong evidence that supports *p* that we later find out was misleading. Even strong evidence can support false beliefs. This is why reliability is important. The more reliable an evidence-gathering process is, the more likely it is that the evidence it has uncovered is not misleading. Thus, if a belief is well-supported by one's evidence *and* the evidence was gathered using a reliable evidence-gathering process, then other things being equal, there is a higher degree of epistemic rationality. However, this claim will require further defense on my part and doing so is one of the primary goals of this dissertation.

This of course brings up the important issue of whether a *synchronic* or a *diachronic* approach to epistemic rationality/justification is the correct approach. More formally and more precisely, a synchronic view of epistemic justification/rationality holds the following:

> Belief *B's* justificatory status at a time *t* is fixed entirely by the "non-historical" (or, "non-backward-looking") properties that *B* has at *t*.

Evidentialism is a synchronic approach to epistemic rationality/justification in which a belief *B's* justificatory status turns entirely on the total evidence *E* that subject *S* possesses at time *t*. In contrast, a diachronic view of doxastic justification/rationality holds the following:

> Belief *B's* justificatory status at *t* is fixed at least partly by some of *B's* "historical" (or "backward-looking", or "pre-*t*") properties.

Responsibilist and reliabilist accounts of epistemic justification are examples of diachronic approaches to epistemic justification/rationality. I turn to this important discussion in part II of this chapter below. I will argue that a diachronic approach does a much better job than a synchronic approach in accounting for everything that contributes to a belief's epistemic rationality/justification.

II. Synchronic vs. Diachronic Accounts of Epistemic Justification

In part II of this chapter, I will continue to defend my two-component theory of epistemic justification. I will examine whether Kelly and Conee and Feldman have given us sufficient reasons to believe that a synchronic account of epistemic justification is correct, i.e.—one that rules out diachronic conceptions of epistemic justification. Thus, I will examine both synchronic and diachronic approaches to epistemic justification. I will argue that both approaches are required in order to give a full account of epistemic justification. That is, I will argue for two claims. First, that the evidentialist conception of justification provides a satisfactory *synchronic* account of *categorical* justification. This simply means that if person S has evidence *e* at time *t, e* supports S's belief that *p* at time *t,* and S bases *p* on *e* at *t*, then *p* is justified (well-founded) for S at *t* (justified *period*). But the second claim that I will continue to defend is that this synchronic account of categorical justification fits within a broader *diachronic* account of justification in which the reliability of virtuous inquiry plays a salient role in fixing the *degree* to which a belief is justified.

Importantly, I'm not arguing for two different types of epistemic justification, i.e.—synchronic and diachronic justification. Rather I'm arguing that there are two different *approaches* to epistemic justification, i.e.—synchronic and diachronic approaches. My view simply entails that there are two different dimensions along which a belief can be justified: 1) the synchronic dimension in which a belief is justified by the evidence one has and 2) the diachronic dimension in which the degree to which the belief is justified is increased if the evidence that supports that belief was gathered virtuously (reliably).[29] But all of this is still referring to *one*

[29] I will refine this distinction in Chapter Five. Ultimately, I will argue for a two-component view of epistemic justification in which a belief can be justified along two dimensions: (1) the *fittingness-dimension* and (2) the *reliability-dimension*.

type of epistemic justification. Another way to say this is that while epistemically justified beliefs can result from non-virtuous inquiry, well-founded beliefs that result from virtuous inquiry are (other things equal) *more* epistemically justified than beliefs that result from non-virtuous inquiry.

First, let us look at Feldman's view of epistemic justification as it relates to synchronic and diachronic issues. Feldman argues that evidentialism is a synchronic approach to epistemic justification (Feldman, 2004a, pp.188-89). The rational thing to do at any one moment is to follow the evidence you have at that moment. Evidentialism is not concerned with how to conduct inquiry over periods of time. In other words, it is not concerned with how one should gather evidence, when one should seek out additional evidence, etc. These are diachronic questions. Feldman claims that diachronic questions are *moral* or *prudential* questions rather than epistemic questions. Feldman says, "If the fundamental epistemic goal is just to have reasonable beliefs, then nothing about evidence gathering techniques or the like follows as a means to that goal" (Feldman, 2004a, 188). He thinks that whether you should gather more evidence or how you should gather evidence depends upon what topics are of interest to you, and you should do so only if having a true belief about the subject matter of the proposition makes a moral or prudential difference and only if gathering more evidence or using certain methods to obtain evidence improves your chance of getting it right (Feldman, 2004a, p.189). Evidentialism provides no guidance for what to do about such questions.

II.1 - Trent Dougherty's Defense of Feldman—Reducing Responsibility

Trent Dougherty (2010) offers a defense of Feldman's view. He argues for the thesis that "….any normativity concerning belief that goes beyond fitting the evidence, and in particular epistemic responsibility, is either moral or instrumental" (Dougherty, 2010, p.534). In support of

this thesis, Dougherty offers a reductionist argument in which *all* instances of epistemic responsibility are in fact either 1) forms of *instrumental rationality* or 2) *moral responsibility* insofar as there is anything defective *beyond* one's beliefs not fitting the evidence one has at the time. Thus, he is making the following identity claim:

> **Identity** Each instance of epistemic irresponsibility is just an instance of purely non-epistemic irresponsibility/irrationality (either moral or instrumental).

This is to be understood as a conditional: If X is an instance of epistemic irresponsibility, then X is just an instance of either moral irresponsibility or instrumental irrationality (Dougherty, 2010, p.537).

Dougherty's argument is a reductionist argument and he argues that we should prefer reductionistic theories because of their simplicity (Dougherty, 2010, p.537). He says that the only thing necessary to make the reductionist thesis more choice worthy is to demonstrate that it can explain the same data as the non-reductionist theory. That is, if there are two theories that are both capable of explaining the data, then we epistemically ought to choose the more parsimonious of the two. Dougherty thinks it is obvious that it is much simpler to posit fewer basic normative categories. Hence, if his argument is successful, then he thinks it more likely that evidentialism is the more acceptable view of epistemic normativity. Here is Dougherty's core argument more formally:

Core Argument

(1) The responsibilist is already committed to the existence of synchronic justification, moral rectitude, and instrumental rationality.
(2) We ought not multiply types of norms without necessity.

---

(3) So if epistemic responsibility can be explained in terms of synchronic justification, moral rectitude, or instrumental rationality, then the reduction ought to be accepted.[30]

[30] It's important to point out that Dougherty uses the terms 'synchronic justification/rationality' and 'diachronic justification/rationality'. He defines synchronic rationality as rationality in which a belief fits the evidence at a particular time and diachronic rationality as a cross-temporal assessment of rationality (Dougherty 2010, p.538). He

Dougherty's aim is to defend the antecedent of the conclusion of this argument, which would establish his reduction (Dougherty, 2010, p.544).

To get his reductionist argument going, Dougherty offers an anecdote that seems to be problematic for evidentialism:

> I have a friend—let's call him 'Craig'—who is a 'special creationist'. That is, he thinks that the major species of animals were created separately—as opposed to their having a common ancestor—and that this occurred 6,000 to at most 8,000 years before present (YBP is a common metric in sciences that deal with the distant past). I was frustrated with the persistence of this unfortunate belief. The problem, though, didn't seem to be that his beliefs didn't fit his evidence—they did seem to fit his evidence, for he had read very narrowly on the subject and had been raised and schooled all his life in an apparently reliable community which sustained this belief in the usual social ways, and which had reasonable-sounding stories for why people deny their views. Rather, the problem seemed to be precisely that he only had the very limited evidence he had, since I'd often recommended books challenging his views (Dougherty, 2010, p.538).

According to Dougherty, it appears that Craig's belief is synchronically rational, i.e.—it seems to fit the evidence he had during that current time-slice.[31] However, his belief was not diachronically rational, i.e.—cross-temporally rational. Now in order to assess the diachronic justification of Craig's belief we have to look at it over some particular interval of time.

---

also uses the phrase 'synchronic (evidential) justification'. I believe the way he uses these terms is a mistake. It leads the reader to believe that there are two types of justification instead of one, i.e.—synchronic justification and diachronic justification. Throughout this chapter I have used the phrases 'synchronic accounts/approaches to epistemic justification/rationality' and 'diachronic accounts/approaches to epistemic justification/rationality'. This is crucial if we want to maintain that there is only one type of justification that we are trying to explain. I will be concerned with synchronic and diachronic *accounts/approaches,* which I believe is the real concern in this debate.

[31] Contrary to Dougherty's claim, I don't think that Craig's belief is clearly justified prior to getting additional evidence from another person. Craig seems to have done some reading on the subject of special creationism, albeit narrowly, which doesn't seem to constitute sufficiently strong evidence to support his belief. The rest of his evidence also seems pretty flimsy. Craig doesn't seem to have enough evidential support for his belief to be even weakly justified. Just because Craig has *some* evidence that supports his belief doesn't mean his belief is justified. That is, mere evidential support isn't enough for propositional justification. Feldman supports this same point with an example: If you learn from a reliable weather report that there's a 51% chance of rain tomorrow, then while your total evidence now supports the proposition that "It'll rain tomorrow", you're not yet justified in believing it'll rain tomorrow (Feldman, 2003, p.15). Evidential support must be *sufficiently* strong for propositional justification. However, even if there is a question about whether Craig's evidence supported his belief initially, I don't think it causes any major problems for Dougherty's example because we can just suppose for the sake of argument that Craig's belief fits his evidence sufficiently for propositional justification and then proceed with the rest of Dougherty's argument.

Consider the time at which it is suggested to Craig that he read some books that provide strong arguments against the young-Earth creationist view. Suppose Craig refuses to read these books. This appears to be a case of epistemic irresponsibility (suppose there is no good reason why he shouldn't read them). Furthermore, it *seems* that Craig has no new evidence since he has refused to read the books (I will address this below since this doesn't actually seem to be the case). However, evidentialism requires that a change in the *epistemic* status of a belief must issue from a change in *evidential* status. To see this challenge more clearly, Dougherty provides the Craig Argument more formally:

The Craig Argument

(1) In the case of Craig the creationist, there is a change (a drop) in epistemic status of his belief.
(2) However, there is no change in his evidential status.
(3) If 1 and 2, then Identity is false, unless the negative evaluation can be explained by instrumental irrationality or moral irresponsibility.
(4) The negative evaluation can't be explained in terms of practical irrationality.
(5) The negative evaluation can't be explained in terms of moral irresponsibility.
---
(6) So Identity is false. (Dougherty, 2010, p.539)

This argument is supposed to weigh in favor of the irreducibility thesis against the reductionist. This is a valid argument, but Dougherty argues that it is not sound. Since his identity thesis is disjunctive, his strategy is to argue (~4 v ~5).

Importantly, Dougherty acknowledges that premise (2) may also be false. It is controversial whether the disagreement of an epistemic peer constitutes evidence against one's belief. He thinks that it does constitute evidence against one's belief. If so, then (2) is false and Craig's refusal to read the suggested books means that he is ignoring important evidence against his belief (ignoring his reliable friend's suggestion). So there is a good evidentialist explanation for why Craig's justification drops, i.e.—because his belief no longer fits the evidence (or at

least, doesn't fit the expanded body of evidence as well as it fit the initial body of evidence). Craig's belief at time *t* (before his reliable friend's suggestion) is justified because his belief fits his evidence, but his belief at *t + 1* (after his reliable friend's suggestion) is no longer evidentially justified (or less justified) because he gets partially defeating evidence from his friend's testimony. Craig ought to suspend judgment. Thus, the change in the epistemic status of Craig's belief is a change in synchronic justification over two times (*t* and *t+1*). So the appearance of epistemic irresponsibility in this case can be partly explained by the fact that Craig is not being sensitive to the evidence he has. Dougherty argues that this can be explained by a couple of things: 1) it could be the result of some cognitive dysfunction, but then it hardly seems that we could blame Craig for holding the belief he does (unless maybe the dysfunction was self-inflicted) or 2) through some discrete act that is a case of instrumental irrationality or purely moral irresponsibility (Dougherty, 2010, p.539-40). Let's look at how Dougherty cashes this out.

Dougherty holds that if there is any sort of evaluation of a belief other than evidential fit, then the evaluation is either instrumental or moral. There is either something at stake in the matter or there is not. If there is nothing at stake, then there is nothing irresponsible in not being diligent and thorough. If there is something at stake, then it is related to one's own interests or the interests of others. If it is related to one's own concerns then it is easily explained in terms of instrumental irrationality. If it concerns the interests of others then you either have a duty to promote their interests or you don't. If you don't then there is no irresponsibility. If you do, then the irresponsibility is moral. All cases of epistemic irresponsibility (other than not being sensitive to one's evidence) can be explained in one of these two ways (Dougherty, 2010, p.539-40).

However, Dougherty points out that the Craig case can be retold so that his belief does fit the evidence and yet there is still irresponsibility. The question is whether this causes problems for the evidentialist. Dougherty thinks not. Suppose now that premise (2) is true (there is no change in Craig's evidential status), but this time Craig lowers his credence to the right degree of belief such that his belief now fits his evidence (which includes his friend's testimony). There is peer disagreement that serves as evidence against Craig's view, but a suspension of judgment is not required. In this case, Dougherty stipulates that Craig's credence is still very high. With (2) set to true, there is a potential counter-example to the following evidentialist thesis:

> S's belief B at *t* is subject to negative epistemic evaluation only if B fails to fit S's evidence at *t* (Dougherty, 2010, p.541).

An epistemically responsible agent is one who considers his reliable friend's testimony as evidence against his own belief. If Craig is aware of this fact, wants to have true beliefs about creationism, recognizes that he has the time to read the books, and has no moral obligation not to read them, then he is clearly instrumentally irrational for failing to read the books. That is, he is failing to attempt to secure his goal or end by what is—by his own lights—an effective means. Hence, premise (4) is false.

Now, if Craig is blamed for not caring enough about the truth as a reason for failing to read the books, then such blameworthiness might be construed as a type of moral blameworthiness. It might be that to lack a sufficient desire for the truth makes someone a morally bad person somehow. For example, someone might have a moral duty to his/her community to believe certain things that contribute to the well-being of that community. This might include the duty to provide an accurate account of the nature of the world to one's offspring. If this is what is happening in Craig's case, then the failing is due to moral irresponsibility. Hence, premise (5) above is false (Dougherty, 2010, p.542).

If Dougherty's arguments are sound, then he has successfully defended the antecedent in the conclusion of his Core Argument. That is, he has established that epistemic responsibility can be explained in terms of synchronic justification, moral rectitude, or instrumental rationality. Hence, his reduction goes through (or at least, he has successfully defended such a reduction from one interesting argument to the contrary, i.e.—The Craig Argument).

II.2 - Responses to Dougherty—Recovering Responsibility?

Guy Axtell (2011) provides a lengthy response to Dougherty in an attempt to make a case for diachronic approaches to epistemic justification. In particular he defends a responsibilist view of epistemic justification in which the methodology we use, intellectual virtues in play, etc., play a prominent role in a belief's epistemic justificatory status. Axtell provides many arguments against Dougherty which I do not have the space to address. However, many of his arguments can be handled by the arguments I made against responsibilism in Chapter Two. Having said that, there are important points that he makes that do support my two-component view of epistemic justification.

As we have seen, evidentialism holds that when a person's *belief* fits the evidence, then all other forms of negative evaluation concerning belief are either moral or instrumental. Among the criticisms Axtell has against this view is that it is belief-centric. That is, Axtell thinks that it is a mistake to apply conceptions of epistemic normativity solely to *beliefs*. Instead, he thinks there is an important sense in which epistemic normativity applies to *persons (agents)* and their habits and dispositions (Axtell, 2011, p.432 and 434). Axtell takes *personal* justification to be more basic than doxastic or propositional justification (or at least considers it to be an important source of justification). More precisely, he claims that there is more to epistemic normativity and to an appropriate understanding of responsibility-relevant norms than evidential fit.

Axtell believes that the evaluation of *persons (agents),* as opposed to the evaluation of only beliefs (cognitive states), is properly part of the theory of knowledge, and that the evaluation of persons is never wholly captured by a synchronic approach to justification abstractable from motivational factors and from habits of inquiry. Axtell thinks the latter is properly part of the theory of knowledge. Dougherty rejects such a view. Dougherty argues that persons can only be evaluated for their actions or their properties. Beliefs are not actions, but acts of inquiry are. For example, walking to the library, checking out a book and reading it, interviewing witnesses, and conducting experiments, are all types of *actions.* According to Dougherty, "….such acts are just as easily evaluable morally and rationally when they are aimed at forming true beliefs as they are when aimed at finding a girlfriend" (Dougherty, 2011, p.627). But what about properties of a person?

Dougherty argues that epistemically evaluable properties cannot be actions, but maybe they are habits or dispositions. However, he thinks there is a difficult *generality problem* that plagues such a view. He asks us to consider the following dispositions: "*to inquire diligently; to inquire diligently on matters non-religious; to inquire diligently on matters horticultural; to inquire diligently on Tuesdays; to inquire diligently after having had coffee; to inquire diligently on matters caffeinated"* (Ibid). How much narrower can these norms be, he asks? And what about other types of dispositions such as *to play games fairly*? Dougherty asks, "Is there a sui generis, emergent kind of laudatory normativity which evaluates such dispositions?" (Ibid). His conclusion is that whether it makes sense to evaluate someone's dispositional properties doesn't give us any reason to postulate new types of fundamental normative categories outside of the moral, evidential, and practical. Again, Dougherty holds that every proposed example of epistemic irresponsibility is either moral irresponsibility or instrumental irrationality.

Furthermore, Dougherty argues "….that when a morally irresponsible action has to do with a belief the irresponsibility doesn't take on some sui generis 'epistemic' nature any more than forgetting to feed the cat takes on a sui generis, emergent 'feline irresponsibility'" (Dougherty, 2011, p.626). This also applies to *intellectual dispositions and habits*, *mutatis mutandis*.

> The bad habit of over-feeding the cat does not imply the existence of some sui generis, emergent cat normativity. It's just a bad habit that involves a cat. The habit of under-feeding a mind doesn't imply the existence of some sui generis, emergent epistemic normativity. It's just a bad habit that involves a mind. Some people are quite generally prudent, others quite generally imprudent, and others spotty. But it would be profligate to posit an emergent form of normativity for every hole in someone's prudence (Dougherty, 2011, p.629).

Just because some practical aim is 'epistemic' in nature, doesn't mean that some new natural kind of normativity is created, any more than in a case where the practical aim is distinctively automotive (Dougherty 2011, p.630).

Dougherty argues that the same thing applies to moral irresponsibility such as forgetting to mail an important check, drinking too much, or spending too much on a new watch. All of these forms of irresponsibility fall under the same category of moral irresponsibility. It wouldn't make sense to break these items up into different categories of irresponsibility because they are too dissimilar from one another, i.e.—*memorial irresponsibility, bodily irresponsibility, and fiscal responsibility*. Dougherty argues that these are all cases of moral irresponsibility in different domains of life. The same line of thinking applies to cases of epistemic irresponsibility. There is no new, natural kind of normativity called *epistemic irresponsibility* but instead there is a form of moral or practical failure that has epistemic *consequences* (Dougherty, 2012, p.282).

But before we concede to Dougherty, perhaps we should further consider the view that there is some such thing as 'personal justification' in a bit more detail.

II.3 – Doxastic Justification vs. Personal Justification

Mylan Engel (1992) argues that there are two types of justification: *doxastic justification* and *personal justification*. He argues that the following equivalency thesis is false:

> **(ET)** *S* is epistemically justified in believing that *p iff* S's *belief* that *p* is epistemically justified (Engel, p.135).[32]

The left-hand side of this biconditional is evaluating *S* (the person), as being epistemically justified and therefore it is persons or (more broadly) cognizers that are being evaluated. The right-hand side of this biconditional is evaluating the justificatory status of S's *belief.* Engel argues that the left-hand side and the right-hand side of the biconditional are not equivalent and that both persons and beliefs can be epistemically evaluated. The idea is that people are often epistemically evaluated in terms of the reasoning that *led* to their beliefs, e.g., if they have reasoned well regardless of the belief they have adopted. People's beliefs are also positively epistemically evaluated even when their belief was arrived at in an epistemically irresponsible way. Engel has us consider the following example to illustrate his point:

> Consider Sally the misinformed logic student. Sally is told by her incompetent logic instructor that modus ponens is an invalid argument form. Not realizing her instructor's incompetence, she comes to regard modus ponens as an invalid form. One day, in a moment of wanton logical abandon, she comes to hold a belief on the basis of a modus ponens argument with obviously true premises. Here we have a situation where Sally's belief is perfectly reasonable (since it follows from obviously true premises), but Sally is unjustified in believing it (because, given her situation, she is being epistemically irresponsible in using modus ponens). Thus, her belief is evaluated positively, while she is evaluated negatively, in direct contradiction to the claim that beliefs and believers cannot be evaluated independently (Engel, p.135-36).

[32] I have been concerned with *doxastic* justification or what it means for a *belief* to be justified. I do think that the left-hand side and the right-hand side of ET are not equivalent. However, I disagree with Engel that each side of the biconditional represents a distinct type of justification, i.e.—doxastic vs. personal. It is a category mistake to say that a *person* is *epistemically* justified in believing some proposition. It is *doxastic attitudes* towards a proposition that are epistemically justified or unjustified *for* a person.

Thus, according to Engel, we have two different domains of epistemic evaluation: personal justification and doxastic justification.

Engel defines these two different domains of epistemic evaluation in the following general way (Engel, p.136):

(PJ) Personal justification is a normative notion in terms of which *persons* are evaluated from the epistemic point of view.

(DJ) Doxastic justification is a normative notion in terms of which *beliefs* are evaluated from the epistemic point of view.

Beginning with (DJ), Engel says that for a belief to be epistemically justified is for that belief to have positive epistemic status. But this is too trivial, so he says we need to ask the question differently. Instead, we need to ask when from the epistemic point of view, should a belief be evaluated positively. The epistemic point of view is defined by the goal of maximizing truth and minimizing falsity in a large body of beliefs, so we get:

(Al) S's belief B has positive epistemic status *iff* B is true (Engel, p.137).

He says that this is because from an epistemic point of view, true beliefs are better than false ones. But since this is simply a restatement of the epistemic goal itself, we need a way to evaluate beliefs that will tell us which beliefs promote this epistemic goal of gaining truth and avoiding error apart from their actual truth-value.

Engel thinks that it is objective probabilities that allow us to evaluate beliefs epistemically. From the epistemic point of view, beliefs that are objectively more probable are better than beliefs that are objectively less probable (Engel, p.137-38). This gives us the following:

(A2) S's belief B has positive epistemic status *iff* B has a sufficiently high objective probability of being true (Engel, p.138).

When we put (A2) together with the earlier point that a belief is doxastically justified if and only if it has positive epistemic status, then we get the following non-trivial definition:

> **(DJ')** S's belief B is doxastically justified *iff* B has a sufficiently high objective probability of being true (Ibid).

Engel says (DJ') correctly describes which *beliefs* we want to count as epistemically justified (beliefs best suited for maximizing truth and minimizing falsity). Thus, we have a *prima facie* working conception of doxastic justification. Engel next considers personal justification.

Engel is concerned with the reading of (PJ) in which S does in fact believe some proposition, as opposed to being personally justified if he were to so believe (Engel, p.138-39). Given this, Engel starts with what he thinks is obvious: when we evaluate person S as epistemically justified or unjustified, we are evaluating him positively or negatively from an epistemic point of view. When we do so we are praising and blaming S epistemically. Thus, Engel gives us a first approximation of personal justification (Engel, p.139):

> **(PJ_j)** S is personally justified in believing that *p iff* S is worthy of epistemic praise for believing that *p*.
>
> **(PJ_u)** S is personally unjustified in believing that *p iff* S deserves epistemic blame for believing that *p*.

A person merits epistemic praise or blame for believing a certain proposition depending on whether he has been epistemically responsible in coming to believe that proposition. Being epistemically responsible entails reasoning carefully, checking her work, considering defeaters, weighing the evidence, etc. Being epistemically irresponsible involves recklessly adopting beliefs, ignoring counter-evidence, etc. Thus, we get a revised definition of personal justification (Engel, p.140):

> **(PJ_j)** S is personally justified in believing that p *iff* S has come to believe that p in an epistemically responsible fashion.

(PJ$_u$) S is personally unjustified in believing that p *iff* S has been epistemically irresponsible in coming to believe that p.

Hence, Engel's conception of personal justification is connected to notions of epistemic praise, blame, responsibility, and irresponsibility.

Engel argues that most epistemologists have conflated personal with doxastic justification. On his account therefore, there are two types of epistemic evaluation. When we evaluate a *person* we are evaluating whether they were epistemically responsible. When we evaluate a *belief* we are evaluating whether it is doxastically justified, i.e.—whether it is sufficiently objectively likely to be true. Engel provides three examples that he thinks reveal these two types of epistemic justification. I will argue that his analyses of these examples is misguided and that we can better account for what is going on in these examples by appealing to the important distinctions that evidentialists such as Conee and Feldman, Kelly, and Dougherty make with respect to epistemic and instrumental rationality. This negative verdict on Engel will support the view that I am trying to establish, which is that there is only one type of epistemic justification. Furthermore, it will help establish the plausibility of evidentialism as a correct synchronic approach to *categorical* justification.[33]

Engel offers three cases that serve as internalist counter-examples to the externalist theory of process reliabilism.[34] He has us consider a simplified version of process reliabilism:

(PR) S's belief that *p* is justified in W (world) if it results from a belief-forming cognitive process [BCP] which is W-reliable (Engel, p.141).

---

[33] I'm not saying that I think that evidentialism captures everything that is relevant to evaluating a belief as epistemically justified. What I'm trying to emphasize here is that evidentialists of the sorts I have discussed and will discuss, provide a satisfactory account of when a belief is justified *period*. But as I argued earlier in this chapter, this doesn't rule out other ways to assess a belief's justification, i.e.—the degree to which the belief is justified given that the evidence justifying that belief was gathered through virtuous inquiry.

[34] The rough idea of internalism is that justification is solely determined by factors that are internal to a person. Externalists deny this, asserting that justification depends on additional factors that are external to a person. The evidentialist theories we have considered are internalist. Virtue responsibilism and virtue reliabilism are externalist.

A BCP is W-reliable *iff* the indefinite probability of beliefs produced by that BCP in W being true beliefs in W is high. *Definite* probabilities attach to propositions, while *indefinite* probabilities are 'dyadic' relations relating classes by specifying the probability of a member of one class being a member of a different class. Standardly, the indefinite probability of A being a B is identified with the actual relative frequency with which A's are B's. Engel claims that (PR) provides an important conceptual connection between justification and truth.

> Since (PR) asserts that S's belief that p is justified in W if it is produced by a W-reliable BCP, and since, by definition, the indefinite probability of beliefs produced by W-reliable BCPs being true beliefs in W is high, it follows that justified beliefs in W have a high indefinite probability of being true beliefs in W (Engel, p.141-42).

Engel maintains that a counter-example to (PR)'s necessity needs to include an intuitively justified belief even though it was produced by an unreliable BCP.

Engel: Case 1

Consider the following:

> Consider a possible world W where, unbeknownst to us, Descartes' evil demon hypothesis is true and where we are unfortunate enough to reside. In W, virtually all of our beliefs turn out to be false, owing, of course, to the clever manipulations of the demon. Consequently, the BCPs (e.g. perception, memory, and inference) which have produced our beliefs are unreliable in W. So, according to (PR), virtually all of our beliefs are unjustified in W, since they are produced by BCPs that are unreliable in W (Engel, p.142).

In response to Case 1, Stewart Cohen and Keith Lehrer (1983) argue that this result is untenable:

> The truth of the demon hypothesis also entails that our experiences and our reasonings are just what they would be if our cognitive processes were reliable, and, therefore, that we would be just as well justified in believing what we do if the demon hypothesis were true as if it were false (Cohen and Lehrer, p.192; Engel, p.142).

The idea is that our experiences and reasonings in W are phenomenologically indistinguishable from what they would be in a verific world W^* where our beliefs *would* be justified. But since our experiences and reasonings are what justify our beliefs, we are intuitively just as well justified in W as we are in W^*. Because we are justified in holding our beliefs in W despite the

fact that they were produced by unreliable BCP's, the result is that reliable production is not necessary for epistemic justification and therefore (PR) is false. Before getting to Engel's analysis of this case, let's look at Case 2.

Engel: Case 2

Cohen raises another purported counter-example to (PR)'s necessity:

> Consider another demon world W' whose sole inhabitants are A and B. Now imagine that A is a good reasoner, i.e. reasons in accordance with the canons of inductive inference, whereas B engages in confused reasoning, wishful thinking, reliance on emotional attachment, guesswork, etc.[35]

In world W', the demon makes BCP's like inductive inference just as unreliable as BCP's such as wishful thinking, etc. According to (PR), reliability is the defining feature of justification-conferring BCPs. Hence, in W' the unreliable BCP of inductive inference is no more justification-conferring than wishful thinking, etc. Cohen therefore notes:

> Since the beliefs of A and B are both produced by unreliable processes (the evil demon sees to this), a reliabilist theory of justification must render identical epistemic appraisals of both sets of beliefs.[36]

A reliabilist theory must maintain that neither A's nor B's beliefs are justified in W'. However, Cohen asserts:

> Plainly, this cannot be correct. A's beliefs are conditioned by the evidence whereas B's beliefs are not. A is a good reasoner whereas B is not. A's beliefs are reasonable whereas B's beliefs are not. There is a fundamental epistemic difference between the beliefs of A and the beliefs of B. But the Reliabilist does not have the theoretical means to display this difference.[37]

On Cohen's view, the fundamental epistemic difference between A's and B's beliefs is that A's beliefs are supposed to be justified while B's beliefs are unjustified. Thus, because A's

[35] Paraphrased from Stewart Cohen's doctoral dissertation "Justification and Truth", in manuscript, p. 10; Engel, p.143.
[36] Ibid, pp.10-1; Engel, p.143.
[37] Ibid, p.11; Engel, p.143.

unreliably produced beliefs are intuitively justified, reliable belief production is not necessary for epistemic justification and therefore (PR) is false (Engel, p.143).

Engel contends that neither of these purported counter-examples (Case 1 and Case 2) are sound since they conflate personal justification with doxastic justification. He argues that in Case 1, internalists find the result unacceptable because intuitively *we* are equally justified in W as we are in W* since our reasonings and experiences are the same in both worlds. But Engel says that *our* being justified in both worlds is irrelevant to whether our *beliefs* are justified in both worlds. He agrees that we are not personally unjustified in either world since we could hardly be blamed for believing as we do. However, he argues that the evaluation that our beliefs in W are unjustified reveals the right result according to (PR). That is, in an evil demon world all of our beliefs are produced by highly unreliable BCP's and therefore, the indefinite probability of these beliefs being true beliefs is extremely low (at or approaching zero). Engel asserts the following:

> Obviously, objectively improbable beliefs run counter to the epistemic goal of maximizing truth and minimizing error, since they virtually ensure error, and there is nothing epistemically positive about beliefs which virtually ensure error. This suggests that intuitively our demon-manipulated, objectively improbable beliefs are in fact unjustified. Consequently, in the first example, (PR) yields the right doxastic justification evaluations, after all (Engel, p.144).

What is crucial to Engel's claim is the fact that the determining factor in whether a belief is doxastically justified is whether the belief has a high probability of being true.

Engel contends that the second counter-example fails for much the same reason as the first. That is, Cohen also conflates personal justification with doxastic justification. Implicit in Cohen's argument is that the epistemic difference between A's beliefs and B's beliefs is that since A is a good reasoner and B is not, A's beliefs are justified but B's are unjustified. However, Engel argues that Cohen's analysis is mistaken. He says there is a definite epistemic

difference between person A and person B. Person A presumably reasons in an epistemically responsible manner, while person B does not. Therefore, A is personally justified in his beliefs in W', while B is personally unjustified in his beliefs in W'. However, being personally justified is not the same as a belief being doxastically justified. In this case, A's beliefs are not doxastically justified (Engel, p.145). I now turn to explaining why I think Engel's analyses fail.

Engel's analysis of Case 1 fails because he is not sensitive to evidentialist accounts of epistemic justification. A more plausible analysis of Case 1 is that in both W and W* our *beliefs* are justified *period* (categorically justified or justified full-stop) because we have the exact same *evidence* in each world. Given that we have no evidence whatsoever that our BCP's are unreliable, we have no good reason, i.e.—we have no evidence—to suspect that our reasonings are bad. Thus our total evidence supports our belief. That is, we have no defeating evidence to suggest that we have misleading evidence. Our beliefs are epistemically justified because of the evidence. In both worlds there is the same strong evidential fit between our evidence and our beliefs. Of course, this doesn't rule out that there are other factors other than evidential fit that can contribute to a belief's epistemic status (I will put this issue aside for now). This does suggest that evidential fit is necessary for epistemic justification. I agree.

Furthermore, Engel argues that in an evil demon world all of our beliefs are produced by highly unreliable BCP's and therefore, the indefinite probability of these beliefs being true beliefs is extremely low (at or approaching zero). He goes on to say that "objectively improbable beliefs run counter to the epistemic goal of maximizing truth and minimizing error, since they virtually ensure error, and there is nothing epistemically positive about beliefs which virtually ensure error". There are two problems with Engel's view here. First, his view is dialectically improper. That is, this claim is question-begging. It assumes that a belief's being

properly based on strong supporting evidence isn't a positive epistemic state of affairs. This is highly implausible, even if the belief is in fact false and couldn't have easily been true.[38] Engel's view ignores the fallible nature of evidence, i.e.—that even strong supporting evidence can be misleading. Second, as Kelly correctly argued, one's goals are irrelevant to whether one's belief is epistemically justified. The fact that using these unreliable BCP's involves failing to reach one's goals in reaching the truth, is relevant to whether one is *instrumentally* rational. Thus, knowingly using an unreliable process is an instrumentally irrational means to achieving one's end of forming true beliefs, but despite this fact, one's belief is epistemically justified in W and W* because one's belief is supported by the evidence. Of course, whether one is instrumentally rational in W or W* is another issue and one I won't explore here. I conclude that Engel's analysis of Case 1 fails to establish his thesis.

Though I would like to reinforce that even if a belief is well-founded, using an unreliable BCP can decrease the degree to which that belief is justified. But again, an unreliable BCP cannot render a well-founded belief completely unjustified. Thus, even though two people, one in W and one in W*, have precisely the same evidence and both of their beliefs are well-founded, the person's beliefs in W are justified to a lesser degree since BCP's in W are unreliable. Other things being equal, evidence gathered using unreliable BCP's is more likely to be misleading evidence, hence the degree to which that belief is justified decreases. Of course, all things considered, by their very nature all BCP's in W are unreliable so all evidence gathered by BCP's in W will be misleading. On the other hand, a person's well-founded belief in W* that was formed using BCP's that are reliable is justified to a higher degree since evidence gathered by reliable processes is more likely to be non-misleading evidence. Reliability isn't necessary for

[38] Thanks to E.J. Coffman for suggesting this point.

justification but it can play a salient role in fixing the degree to which a well-founded belief is justified.

Engel's analysis of Case 2 also fails for much the same reason as his analysis of Case 1. He says there is a definite epistemic difference between person A and person B, which is that person A presumably reasons in a responsible manner, while person B does not. Therefore, A is personally justified in his beliefs in W', while B is personally *un*justified in his beliefs in W'. However, being personally justified is not the same as a belief being doxastically justified. In this case, A's beliefs are not doxastically justified. I argue that if there is no epistemic difference between A's and B's belief it's because their beliefs are based upon the same evidence (assuming that they have the same evidence in this case). However, if their evidence differs, then there will be a difference in epistemic justification. But Cohen maintains that A's belief is conditioned by the evidence and B's is not. Thus, A has good epistemic reasons to hold his beliefs (assuming his beliefs fit the evidence), but B does not because his beliefs are *not* conditioned by the evidence. For a belief to be well-founded a person must properly base his belief on his total body of evidence.

Furthermore, if the evil demon has rendered both A's and B's BCP's unreliable then there's no sense in which these processes could contribute to the positive justificatory status of their beliefs. It was stipulated that they are unreliable processes, so there is no difference in reliability between inductive inferences on the one hand and confused reasoning, wishful thinking, reliance on emotional attachment, guesswork, etc., on the other (at least in the world they are in). Thus, A's and B's respective beliefs have a different epistemic status because A's beliefs are conditioned by the evidence while B's is not. That is, A has epistemic reasons to hold the beliefs he does while B does not. Hence, Engel's analysis of Case 2 fails to establish his

thesis. However, I would like to add that if there were some third person C in a verific world, who properly based his belief on precisely the same evidence as person A, and like A was a good reasoner, etc., person C's beliefs would be justified to a higher degree than A's beliefs. A's beliefs are conditioned by the same evidence as C's beliefs so both of their beliefs are well-founded, but the BCP's that A uses are unreliable and C's are reliable.

Engel: Case 3

Engel also addresses (PR)'s *sufficiency*. He considers the 'Norman the Clairvoyant' case as a case that purports to establish the insufficiency of (PR). I discussed this case in Chapter Two, but I think it's instructive to revisit it in order to see why Engel's analysis of it fails to establish his thesis that much of the confusion in epistemology can be attributed to a confusion between personal and doxastic justification.

Here's the Norman case again:

> Suppose that Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact, the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable.

Engel says that one might think that Norman is epistemically unjustified because Norman has no *evidence* of the President's location and he is *unaware* that he has this clairvoyant power. From Norman's perspective it would seem as if his belief just popped into his head out of nowhere. Thus, it would be epistemically *irresponsible* for him to hold such a "spontaneously occurring ungrounded belief" (Engel, p.146). But Engel thinks this only shows that Norman is *personally* unjustified in believing the President is in New York City. However, he says this is irrelevant to whether Norman is *doxastically justified,* i.e.—whether Norman's *belief* is justified. His *belief* is

the result of a highly reliable BCP. Beliefs that result from such a BCP have a very high indefinite probability of being true beliefs and are therefore justified from the epistemic point of view. Likewise, Norman's situation is very much like a child's situation with respect to that child's perceptual beliefs. The child has no rationale for her perceptual beliefs and yet she is most certainly justified in holding them. Thus, Norman's *belief* is justified, but relative to other things he believes, he is epistemically irresponsible in holding the belief and is therefore *personally* unjustified in doing so. In other words, the fault lies with *Norman* and not with his belief. Hence, if (PR) is viewed as an account of doxastic justification, i.e.—as a theory of justified *belief*, then the Norman case fails to impugn (PR) (Engel, p.145-46).

I think Engel gets this case wrong for two important reasons. First, recall the issue from Chapter Two that cases like Norman's can be underdescribed regarding the issue of cognitive integration. That is, cognitive processes associated with clairvoyance are not sufficiently integrated with the rest of the person's cognitive system. Clairvoyance is a process that is completely unrelated and insensitive to the rest of the person's cognitive system governing the formation and evaluation of belief. The senses, memory, introspection, reason, etc., are typically well-integrated with and sensitive to each other as well as the rest of the person's cognitive system. Hence, if we take cognitive integration into consideration, Norman's belief isn't justified.[39]

Furthermore, Engel's parity case of a child's perceptual beliefs doesn't help establish that Norman's belief is justified. Unlike Norman's clairvoyant power, a child's perceptual faculties are well-integrated with the rest of that child's cognitive system. Also, unlike Norman, the child

[39] This doesn't mean that some epistemic agent (say some alien from another planet) couldn't have a clairvoyant ability that is well-integrated. This is a purely contingent matter. However, Norman's case is not such a case. This is a point that Greco himself makes (Greco, 2010, p.155).

has perceptual *evidence* that supports the beliefs formed as a result of those perceptual processes. Of course, one might think that Norman might have evidence that his belief is justified because his belief is accompanied by some sort of feeling of 'correctness' or something. Much the same can be said about memorial beliefs. However, unlike Norman's clairvoyant belief, memorial beliefs are usually well-integrated with other beliefs formed as the result of the person's well-integrated cognitive system. But it's also important to note that no such feeling of 'correctness' occurs in this case. Norman just finds himself believing that the President is in New York City. That's all. Finally, because Norman's clairvoyance isn't well-integrated, his belief isn't sensitive to counter-evidence or defeating evidence. This is because Norman's belief isn't well-integrated with his other beliefs formed as the result of his other well-integrated cognitive faculties. Given this, it would be strange to grant that Norman's belief is justified.

The second reason Engel gets this case wrong is because he misses a crucial reason why Norman's belief is not *doxastically* justified, namely that he lacks any *evidential* basis for holding such a belief and that he confuses this with Norman being epistemically *irresponsible.* It is not *Norman* who is unjustified, but his *belief* that is unjustified. His belief is not epistemically justified *because Norman lacks sufficient evidence to support his belief*. From his perspective he lacks the necessary epistemic reasons to believe that the President is in New York City, i.e.—he lacks proper evidence. Therefore, the evidentialist is better able to explain why Norman's belief is unjustified.

Furthermore, the fact that Norman believes the President is in New York City without any evidence to support such a belief can best be explained by his being instrumentally irrational. Believing something because you are suddenly overcome by that belief is not a reliable means to achieving the end of forming true beliefs and avoiding false ones. It is instrumentally irrational

for someone to use such a means to achieve this end, *unless* they are aware that it is a successful means to that end. Suppose Norman is unaware that he is clairvoyant but finds that these beliefs that pop into his head are highly correlated with actual states of affairs that have obtained in the world. For example, he is unaware of his clairvoyance, but he finds that anytime he has had beliefs about the President's location, they have been correct. Suppose that his beliefs about the President's location have been correct 100% of the time over a thousand separate instances. If this were true then we would have to say that Norman is both epistemically and instrumentally rational. He is epistemically rational because now he has one thousand pieces of evidence to support the belief that anytime he has had a belief about the President's location in the past he has been correct and he has good inductive reason to believe that such beliefs will be correct in the future. Individually, these beliefs don't amount to any significant evidence, but together these beliefs provide a strong evidential case that such beliefs are correct. Likewise, once Norman discovers this correlation, he is also instrumentally rational when he uses these beliefs as a means to figuring out where the President is located. Using these beliefs is instrumentally rational because doing so is a reliable practical method for figuring out the President's whereabouts if one has the goal of doing so.

Therefore, rather than appealing to PJ' and DJ', Norman's case can be better explained by adopting an evidentialist account of epistemic justification.

In sum, Engel makes three mistakes in his analysis of the Norman case: 1) he fails to recognize that Norman's clairvoyance lacks cognitive integration; 2) he confuses epistemic irresponsibility with instrumental irrationality; and 3) he fails to recognize the role that evidence plays in doxastic justification.

Engel's analyses of the three cases above fail to adequately defend (PR) as a theory of doxastic justification. Furthermore, his analyses fails to establish that the wide acceptance of these counter-examples is due to a conflation of personal and doxastic justification. This is because Engel fails to be sensitive to evidentialist accounts of epistemic justification and the important distinction between epistemic rationality and instrumental rationality.

Now, I don't wish to present this response to Engel as an argument against all forms of reliabilism and the role reliability can play in epistemic justification. Instead, I've presented this analysis to cast doubt on the view that there are two types of epistemic justification, i.e.—doxastic and personal justification. I have sided with the evidentialists in response to Leite, Axtell, and Engel.

II.4 – Against Evidentialism: Keith DeRose and 'Epistemic Ought'

I now wish to consider an important argument that Keith DeRose makes against evidentialism. In contrast to Axtell's and Engel's arguments, DeRose's argument has nothing to do (as far as I can tell) with making a distinction between personal justification and doxastic justification. DeRose argues that whether a subject epistemically ought to believe something largely depends on issues relevant to evidence gathering and processing of evidence.

DeRose rejects Richard Feldman's evidentialist thesis:

> O2. For any person *S*, time *t*, and proposition *p*, if *S* has any doxastic attitude at all toward *p* at *t*, then *S* epistemically ought to have the attitude toward *p* supported by *S's* evidence at *t*.

Feldman provides a way to specify what attitudes we epistemically ought to have in different cases:

> If a person is going to adopt any attitude toward a proposition, then that person ought to believe it if his current evidence supports it, disbelieve it if his current evidence is against it, and suspend judgment about it if his evidence is neutral (or close to neutral) (DeRose 2011, p.137; Feldman, 2000, p.679; Feldman, 2004a, p.178).

DeRose offers what he thinks is a counter-example to O2:

> Suppose that Henry firmly believes that p—it doesn't matter much what p is—and has excellent evidence for *p* (evidence that's strong enough to adequately support the firm and confident attitude Henry has adopted toward *p*). Suppose further that Henry doesn't possess evidence against *p*, so the attitude toward *p* that fits all the evidence Henry possesses is the confident belief that *p* which Henry in fact holds. But suppose that Henry doesn't believe *p* on the basis of the excellent evidence for it that he possesses. Indeed, Henry hasn't even considered *p* in the light of this excellent evidence, and the fact that he possesses good evidence for *p* is no part of the explanation for why Henry believes that *p*. Rather, Henry believes that *p* on the basis of some other beliefs of his that he considers to be good evidence for *p*, but which in fact constitute absolutely lousy evidence for *p*. Henry has no business believing *p* at all on the basis of the evidence he is in fact using, much less believing it as firmly as he does. If you want some details—maybe your intuitions are helped out by filling in the story a bit— suppose that *p* is a proposition to the effect that someone that Henry hates very much has done some horrible thing. It's Henry's hatred and resulting desire to believe his opponent has done this horrible thing that causes him to think of his lousy evidence as strong evidence and to thereby believe that *p*. (DeRose, p.138)

Further:

> Suppose then that there is some evidence Henry very easily could have, and should have, gathered, but that he negligently never encountered. This would have been very strong evidence against *p*. So strong that, despite the excellent evidence Henry possesses in favor of *p*, this negative evidence that Henry should have gathered would have completely outweighed the positive evidence he actually possesses, such that disbelief of *p* would have been the attitude that would have best suited Henry's evidence, had he gathered this negative evidence. (We want it to remain the case that the evidence Henry actually possesses still favors *p*, so don't imagine this case such that Henry has good reason to think that the evidence he neglected to gather would have counted against *p*. This would raise the suspicion that, whatever these good reasons for thinking the evidence will be unfavorable to *p* are, they probably already also constitute some evidence against *p* itself.) (DeRose, p.139).

DeRose believes that the following proposition 'P' is true:

> **P** Henry ought not to believe *p* at all.

DeRose thinks P is true for two reasons. The first reason is because Henry's belief is based on lousy evidence. The second reason (and the one I wish to focus on) is that even if Henry had brought all the evidence he had to bear on the issue, he still would have ended up not believing in

p if he had also been responsible in his gathering of evidence. There is relevant evidence that Henry easily could have and should have had. The fact that there is evidence Henry easily could have and should have gathered seems to ground a legitimate, and important sense, in which, because he should have already gathered that evidence, Henry *right now* ought not to believe that *p*. That is, P is true *right now* (DeRose, p.140-41).

In support of P, DeRose cites an important sense of 'epistemic ought' that Feldman endorses:

> [F]orming beliefs is something people do. That is, we form beliefs in response to our experiences of the world. Anyone engaged in this activity ought to do it right. In my view, what they ought to do is follow their evidence (rather than wishes or fears). I suggest that epistemic oughts are of this sort—they describe the right way to play a certain role....They are based on what's good performance. (DeRose, p.142; Feldman, 2000, p.676; Feldman, 2004a, p.175)

This sense of epistemic ought is connected to a certain 'role' we play as believers. To believe what you ought to believe and to the degree you ought to believe, is to perform well in filling the role of a believer. But as DeRose points out, if Feldman is correct about the tie between 'epistemic ought' and the role of a believer, then the evidentialist thesis is surely false. DeRose writes:

> For Henry's performance is poor, in almost every way I can think of. He's doing it (believing), but not doing it right. He's doing a lousy job of gathering evidence, and a lousy job of processing the evidence he does have. His belief seems irrational and unjustified, and even if it happens to be true, it certainly isn't a piece of knowledge. He's not doing a good job of filling the role of a believer. Yet, at least as I understand it, he satisfies Feldman's evidentialist thesis for believing what he ought (Ibid).

In order to do a good job as a believer or to fill the role of a believer well, one should both gather and process evidence well. If Henry had both gathered and processed evidence well, he would not believe *p*. So Henry did not perform well in filling his role as a believer, nor did he believe

what he would have believed if he had performed well in that role, despite the fact that he satisfies Feldman's account of believing what one ought.

DeRose further claims that "It's not the (more general) role of a believer that Feldman's account seems to answer to, but the more specialized role of an *evidence-processor* [my italics]" (DeRose, p.143). The evidence that Henry used didn't actually support his belief, and it was only through bad processing of the evidence that he used that he ended up believing that *p*. However, Henry did end up believing what he would have believed if he had taken all of the evidence he had and processed it well. But this is a very specialized sense of ought and if this is what is meant by 'epistemic ought', then O2 comes out true but in a very trivial way, i.e.—it's no surprise that O2 gets this very specialized notion correct (Ibid).

Conee and Feldman reject DeRose's counterexample to O2. They believe that O2 correctly implies that P is false. Henry has excellent evidence, E1, for the proposition B (that Jones did something terrible). Additionally, Henry has no reason to doubt B. He also has some lousy evidence, E2, for B. Henry mistakenly and for no good reason thinks that E2 is good evidence for B. He never considers B in light of E1. Henry only considers B in light of E2 and believes B on that basis. In spite of DeRose's claim that P is true, accepting B is precisely the attitude Henry ought to take. It is made reasonable by having E1 and Henry not having any reason to doubt B. Thus, he ought to believe B. Thus, P is false (Conee and Feldman, 2011, p.287).

Conee and Feldman think that we must not be misled by the credibility of P. There are other 'nearby facts' about things that Henry ought not do and these 'ought not's' may lead some to mistakenly believe P. Conee and Feldman have us consider P2.

P2 Henry ought not to 'believe as he does' concerning B.

They concede that P2 is true because Henry bases his belief in B on E2 which is bad evidence. Basing his belief on E2 is not well-founded. However, denying P is compatible with the claim that Henry ought not to believe anything, including B, on an ill-founded basis (Ibid).

Additionally, Conee and Feldman have us consider P3-P5:

P3 Henry ought not to infer B as he does.

P4 Henry ought not to base his belief on E2.

P5 Henry ought not to come to believe B as he does.

Feldman and Conee believe that P3-P5 are all true.

> They refer to things that Henry does, but he ought not do. The reasons he ought not do these things are evidentialist and fully compatible with evidentialism's verdict on P. He ought not make bad inferences such as the one he actually makes in accepting B, he ought not base his belief on bad evidence such as the evidence on which he actually bases his belief in B, and he ought not come to believe in the way he did. Nevertheless, B is a proposition that he ought to believe (Conee and Feldman, 2011, p.288).

All of this is compatible with the evidentialist thesis that Henry ought to believe B because the evidence he has *right now* supports B.

Furthermore, Conee and Feldman reject DeRose's suggestion that Henry neglecting or failing to gather evidence about B that he could have easily gotten, and should have gotten, somehow enhances the credibility of P. They argue that this additional stipulation does not make P more credible when P is properly understood. P is a claim about what attitude Henry ought to take toward B, given how things actually are. Henry's failure to gather more evidence does not support P. This produces the following distracting 'ought not':

P6 Henry ought not to have only the evidence he has about B.

Conee and Feldman point out that this is a practical evaluation of Henry's evidence gathering, but this is not what P is about. P is about a doxastic attitude that Henry has toward B and P is mistaken because having that attitude is the correct one, given the evidence that Henry *actually*

has. Feldman and Conee believe that we must consider the total actual evidence that Henry has. If 'ought to believe' is closely connected to epistemic justification, one ought not ignore any relevant evidence that one possesses and one is in no position to consider any evidence that one does not even have. "This leaves all of the evidence one does have as a non-arbitrary determinant of what one ought to believe" (Ibid). Hence, O2 is not arbitrary or trivial in any way.

Conee and Feldman's reply to DeRose is pretty much in line with the various defenses of evidentialism I've already discussed. Conee and Feldman provide a very plausible defense of evidentialism, insofar as evidentialism provides a plausible account of when a belief is categorically justified (i.e.—justified 'period' or 'full stop'). Crucial to Conee and Feldman's view is the following:

> No matter what the answers are to questions about how I ought to conduct my inquiry, where I ought to look for evidence, and so on, there always remain the questions, "What should I believe in the meantime?" "What should I believe until I have a chance (or the courage) to look at that new evidence?" It's that natural and central question to which evidentialism provides a good answer (Feldman, 2000, p.188).

The answer to the questions in the above passage is that a person ought to believe what the person's evidence supports at that time. If your belief *p* is well-supported by the evidence you possess and you don't possess any other evidence against your belief *p*, then there's no *epistemic* reason that you ought to believe anything other than *p*. Even if you should have and/or could have gathered different evidence, the fact is you didn't. Consequently, it wouldn't make sense to believe anything other than what your current body of evidence supports, in this case, *p*. Hence, when you believe *p* based on evidence that supports *p* over believing some other proposition not supported by that evidence, then you are believing precisely what you ought to believe. This evidentialist view holds that the doxastic attitude a person epistemically ought to adopt toward a

proposition is the one that is supported by the evidence the person has at that time. Importantly, Feldman says that the sort of 'ought' in play in this evidentialist view is a) purely epistemic and b) it is an 'ought' about what doxastic attitude is epistemically *permissible* at any given time (2000, p.179). What is epistemically permissible to believe at any given moment is a function of what your total body of evidence supports. You ought to believe what your evidence permits.

Furthermore, Feldman says that how a person *actually* came to possess the evidence the person has does not change what the person ought to believe at that time:

> How the person came to have that evidence, whether by conscientious inquiry or by avoiding potentially troublesome information, is irrelevant to this epistemic fact. Similarly, how the person ought to proceed in the future is also irrelevant. Evaluations of this behavior can be made, of course, but these evaluations are of a different nature than those made on evidentialist grounds [prudential or moral evaluations] (Feldman, 2000, p.190).

I think Feldman is right that how a person actually came to possess the evidence has no bearing on what the person ought to believe at that moment, unless of course the person has *evidence* that the evidence was gathered in some flawed way, e.g., in a biased manner. Having evidence *e* that you gathered your evidence in a biased way *changes* your total body of evidence *E1* to include *e* which gives you a new total body of evidence *E2*, so that your prior belief *p*, which was based on *E1*, no longer *fits* your new body of evidence *E2*. Given *E2*, you ought not believe *p*. However, suppose that you gathered *e* in a biased way but you have no evidence that *e* was gathered in a biased way. If this is the case, then it seems right to say that you ought to continue believing *p* based on *E1*. This is correct. In this latter case, historical facts about how you gathered your evidence have no bearing on whether *p* is justified for you. However, all of this is consistent with the evidentialist view we've considered. That is, in any given circumstance at any time, a person epistemically ought to believe what that person's evidence supports. However, this (like I discussed earlier in this chapter) is consistent with the view of epistemic justification that I have

been advocating. More precisely, I have been advocating a view of epistemic justification in which a belief is categorically justified for a person if that person's current total body of evidence supports that belief. That is, if my total evidence supports *p*, then I'm justified in believing *p*, full stop. Kelly, Dougherty, and Conee and Feldman would all agree with this claim. Where I depart from their evidentialist view is when I make the further claim that this view is compatible with the view that reliability can still play a salient role in fixing the degree to which a belief is justified. They would deny this claim. They think that justification is *solely* a function of evidential fit.

So, in contrast to Kelly, Dougherty, and Conee and Feldman, I endorse the view that even though historical facts having to do with how I gathered my evidence have no bearing on whether my belief that *p* is well-founded (justified *period*), such historical facts can and often do have a direct bearing on the *degree* to which my belief is justified. So, I don't think that DeRose's arguments impugn evidentialism as a *synchronic* account of *categorical* justification. DeRose's conditional claim that <if a person would have and/or should have gathered different/additional evidence, then that person would have and/or should have believed differently> has no bearing on what evidence the person *actually* has nor does it have any bearing on what the person ought to believe given the person's current total body of evidence. However, I do think whether the method of inquiry the person *actually* used to gather her current body of evidence was virtuous or vicious does have a bearing on the degree to which that belief is justified (other things being equal). In other words, whether the person could have or should have gathered different/additional evidence is irrelevant to a belief's justificatory status, but the *actual* mode of inquiry used to gather the evidence is often relevant to the degree to which a belief is justified (if the mode of inquiry was virtuous). Thus, DeRose's argument doesn't

impugn evidentialism as a synchronic account of categorical justification, but this evidentialist view doesn't rule out that virtuous inquiry can play a role in determining the degree to which a belief is justified. I think this is the right conclusion.

I wish to make one more crucial point about the sort of epistemic 'ought' discussed in this section. This 'ought' has to do with what doxastic attitude you are epistemically permitted to take toward a proposition given the total body of evidence you currently possess. This is a similar point to the one made in Chapter One about epistemic justification. That is, when we are talking about epistemic 'ought' here, we are not talking about epistemic *blameworthiness* or *blamelessness*. That is, it's not the case that you ought to believe some proposition simply because you are somehow blameless for holding the belief. Again, this sort of doxastic justification is *subjective deontological justification*. But this is not the sort of 'ought' that is relevant to what I'm concerned with in this section. I'm concerned with a sort of 'ought' that is more objective than subjective deontological justification. Consider once again the example of someone who grows up in a religious society and is taught to listen to the deliverances of an oracle. We can't blame her for believing what she believes, but if her belief is not supported by her evidence, then she epistemically ought *not* believe what she believes. The same goes for the example of the paranoid man who believes he is being spied on but lacks any evidence supporting his belief. In this case, he *epistemically* ought not believe he is being spied on. In both of the above cases, the person is epistemically blameless for holding his/her belief, but it is still not what each person *epistemically* ought to believe. In both of these examples, the person doesn't have evidence that sufficiently supports his/her belief. Thus, in both examples, each person epistemically ought not believe what he/she believes because the evidence he has does

not support that belief. The evidence each person possesses determines what is epistemically *permissible* for him/her to believe.

III. Conclusion

In this Chapter I argued that epistemic justification is not just a species of instrumental rationality. I also argued that methodological epistemic rationality is not merely a form of instrumental rationality.

I also argued that evidentialist views such as Kelly's and Conee and Feldman's provide a plausible account of when a belief is epistemically justified for a person. More precisely, I've argued that evidentialism has the conceptual tools to explain when a belief is epistemically justified *period*. My conclusion from this discussion was that a belief is epistemically justified *period* if the total body of evidence a person possesses supports that belief at that time. Thus, evidentialist accounts of epistemic justification provide a satisfactory synchronic account of categorical justification.

I also argued against diachronic approaches to epistemic justification that rely on a distinction between personal justification and doxastic justification. I argued that personal justification is irrelevant to epistemic justification. Furthermore, I argued against diachronic approaches of epistemic justification that claim that a belief is not justified if the person holding the belief based on some body of evidence, could have and/or should have gathered different/additional evidence but failed to do so. However, I argued that it is compatible with a synchronic account of categorical justification that how a person *actually* gathered her evidence often does have a bearing on the *degree* to which a belief is justified. Specifically, if a person gathers her evidence virtuously (reliably), then (other things being equal) the degree to which

that belief is justified is increased and if a person gathers her evidence viciously (*un*reliably), then (other things being equal) the degree to which that belief is justified is decreased.

In Chapter Four I will continue to motivate and defend the evidentialist view I've been discussing in Chapter Three. I will provide a more detailed and fleshed out account of evidentialism that appeals to 'best explanations' to account for evidential fit, i.e.—*explanationist evidentialism*. This account provides both necessary and sufficient conditions for a belief's being categorically justified. In Chapter Five I will argue that the reliability or unreliability of one's inquiry can play a salient role in fixing the degree to which a belief is justified.

CHAPTER FOUR

Explanationist Evidentialism

In Chapter Three I gave numerous examples and arguments that motivated and defended a rough evidentialist view of necessary and sufficient conditions for *categorical* epistemic justification. This chapter will continue to motivate, develop, and defend this view but in much greater detail. I will do so by explicitly appealing to the work of Conee and Feldman and more recently Kevin McCain (2014). I will discuss what I think is the best synchronic approach to categorical doxastic justification: *explanationist evidentialism.* Explanationist evidentialism offers a full account of well-founded belief.

I will begin by discussing the nature of evidential fit or the epistemic support relation. Thus far I've mentioned the notion of 'evidential fit' numerous times, but I haven't yet discussed what it amounts to. Having a strong conception of evidential fit is important for understanding how explanationist evidentialism accounts for doxastic justification. I will endorse an *explanationist* account of evidential fit (a view argued for by Conee and Feldman and McCain) that I think is amicable to my own two-component view of justification. Roughly, an explanationist account of evidentialist fit holds that a belief *p* is justified for a person S if *p* is the best explanation S has for why S has evidence *E* or *p* is an available entailment of the best explanation of S's evidence (McCain, p.6). Hence, explanationist fit offers a full account of *propositional justification* which is crucial for understanding well-founded belief.

Before I discuss the notion of 'evidential fit' I will spend a little more time motivating the type of evidentialism I have been considering thus far. The type of evidentialism I have been

advocating is an *internalist* and *mentalist* form of evidentialism.[40] Fundamentally, an internalist view (at least most) holds that the *justifying* of a belief is done by things internal to the mental life of the subject, i.e.—mental states. On this view, evidence consists entirely of mental states and evidence is what justifies a belief, hence the mentalist component to internalism. Mentalism is a supervenience thesis which means that justification supervenes on mental states. This is best understood by the idea that the *beliefs* of any two individuals that are alike mentally are alike justificationally (Conee and Feldman, 2004a, p.56). Importantly, this internalist and mentalist conception of evidentialism is compatible with the fact that there are *non-mental/external* factors that are necessary for justification. For example, internalism and mentalism accept that *facts* about epistemic support (fit) are not internal to the subject—they are necessary truths about which propositions certain mental states make it rational to believe. That is, the *fact* that evidence E supports believing *p* is a necessary truth, and therefore external to S. However, the view is still internalist because E must itself be internal and it is still mentalist because E consists solely of mental states.

In order to motivate the discussion in this chapter, I will briefly present three cases that support an evidentialist synchronic approach to categorical justification. Case 1 is the now famous 'New Evil Demon' (NED) problem[41]:

> Imagine a world where an evil demon creates non-veridical perceptions of physical objects in everybody's minds. All of these perceptions are qualitatively identical to ours, but are false in the world in question. Hence, their perceptual belief-forming processes (as judged by the facts in that world) are unreliable; and their beliefs so caused are unjustified. But since their perceptual experiences—hence evidence—are qualitatively identical to ours,

[40] There are other types of externalist, non-mentalist, etc., evidentialist theories that are different from the version of evidentialism I'm concerned with here. There are important differences between them and I will touch on some of these distinctions later on, but a more thorough discussion of these distinctions goes beyond the scope of this dissertation.

[41] This is the problem presented by Cohen and Lehrer that I discussed in the previous chapter.

> shouldn't those beliefs in the demon world be justified? Evidently, then, reliabilism must deliver the wrong verdict in this case.[42]

Here a person in the demon-world is justified *period* since her belief fits the relevant perceptual evidence, *regardless* of the fact that in the demon-world *all* beliefs are formed by unreliable belief-forming processes. What is noteworthy is that this case doesn't rule out the possibility that reliability can play a role in determining the degree to which a belief is justified. Evidence justifies necessarily in both the demon-world and the actual world. In the actual world our perceptual belief-forming processes *are* reliable. So why not think that such processes can be relevant to the justificatory status of a belief in the actual world? This case simply tries to show that reliability isn't *necessary* for justification. It says nothing about reliability's *sufficiency* for increasing a belief's level of justification.

Cases 2 and 3 come from Conee and Feldman (2004a). Here is Case 2:

> Initially, Smith has excellent reasons to believe that Jones, who works in his office, owns a Ford. Smith deduces that someone in the office owns a Ford. The latter belief is true, but the former is false. Smith's reasons derive from Jones pretending to own a Ford. Someone else in the office, unknown to Smith, does own a Ford. The fact that Jones is merely simulating Ford ownership keeps Smith from knowing that someone in his office is a Ford owner, but it does not prevent Smith from being justified or diminish his justification. At a later time Smith gains ample reason to believe that Jones is pretending. At that point Smith is not justified in believing either that Jones owns a Ford or that someone in his office owns a Ford (Conee and Feldman, 2004a, p.60).

Case 2 is similar to cases we saw in Chapter Three in that Smith has excellent reasons to believe what he believes and his belief is therefore justified. Those reasons are in the form of evidence for his beliefs. Furthermore, Smith's belief that Jones owns a Ford is justified even though that belief is false (just like in NED). Finally, when Smith's evidence changes, so does the justification for his initial belief. In this case he is no longer justified in believing that Jones

[42] Taken from Goldman, Alvin and Beddor, Bob, "Reliabilist Epistemology", *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2015/entries/reliabilism/>.

owns a Ford (or that anyone in his office owns a Ford for that matter). Whether reliability plays a role in determining the justificatory status of Smith's belief is irrelevant. The point is that whether Smith's belief is categorically justified is a function of his total body of evidence. But again, this case doesn't rule out the possibility that reliability or unreliability can play a role in determining the degree to which a belief is justified.

Case 3 is similar to Case 1:

> Hilary is a brain in a vat who has been abducted recently from a fully embodied life in an ordinary environment. He is being stimulated so that it seems to him as though his normal life has continued. Hilary believes that he ate oatmeal for breakfast yesterday. His memorial basis for his breakfast belief is artificial. It has been induced by his "envatters" (Conee and Feldman, 2004a, p.60-61).

Conee and Feldman provide two versions of the relevant details. Hilary A's ". . . recollection is very faint and lacking in detail. The meal seems incongruous to him in that it strikes him as a distasteful breakfast and he has no idea why he would have eaten it" (Conee and Feldman, 2004a, p.61). Hilary B's ". . . recollection seems to him to be an ordinary vivid memory of a typical breakfast for him" (Ibid). Conee and Feldman say that even though both Hilary A's and Hilary B's breakfast beliefs are false and their bases are abnormal, Hilary A's belief is not well justified and Hilary B's belief is well justified. Hilary A's belief differs internally from Hilary B's belief because Hilary B's mental states include better evidence for his belief (Ibid). It's debatable whether Hilary A has sufficient evidence to justify his belief (the evidence might not sufficiently support his belief). However, intuitively it seems correct to say that Hilary B's belief is categorically justified given his memorial evidence and this is the case even though his belief is false and is not formed by any reliable belief-forming process (his memorial evidence is fed to him by an unreliable process that feeds him false information). But again, this example doesn't rule out the possibility that in a non-envatted world a person's belief can receive a justificatory

boost from the fact that the evidence that supports his belief was gathered using a reliable-belief forming process (such as virtuous inquiry). This example also doesn't rule out there being a justificatory difference between two person's beliefs who have precisely the same evidence, but one is envatted and one is in a verific world. They can both have precisely the same evidence and therefore their beliefs are well-founded, but the person in the envatted world will have evidence that results from unreliable processes while the person in the verific world will have evidence that results from reliable processes. If so, the person in the verific world will have well-founded beliefs that are justified to a higher degree than the envatted person's.

The diagnoses for each of these cases is compatible with my two-component view of justification in which 1) your belief is categorically justified if it fits your total body of evidence and 2) the degree to which your belief is justified is often largely a function of whether your evidence was gathered using virtuous (reliable) inquiry or vicious (unreliable) inquiry. These cases don't establish or even suggest that reliability or unreliability make a justificatory difference in some instances, but these cases don't rule out the possibility either. I will defend this possibility later on. For now it's important to note that in each case above, a change in *evidence* makes an epistemic difference. Furthermore, regardless of issues of demon-deception, brain stimulations, or false or misleading information, evidential relations between a belief and the evidence that supports it are necessary relations. As long as this necessary relation holds, a person's belief is categorically justified.

I. Evidential Fit

I now turn to a more fully fleshed out discussion of the components of an explanationist account of evidentialism. I begin by discussing the nature of evidential fit. Here is a theory neutral definition of evidential fit or what is known as the 'epistemic support relation':

Evidential Fit: The relation that holds between some evidence *e* and a proposition *p*, such that when person S has evidence *e*, S is justified in believing that proposition.

Furthermore, the justified doxastic attitude toward a proposition at a particular time is the one that *fits* the evidence that person has at that time. However, this generic definition doesn't help us understand what this epistemic support relation amounts to. I will defend an *explanationist* account of evidential fit which essentially holds that *for belief p to fit S's evidence at a given time is for p to be the best explanation of that evidence or for p to be an available entailment of the best explanation of S's evidence*. But before fleshing out the details of this explanationist account, I will explain some other accounts of evidential fit and the problems inherent to them. This will help to motivate the preferred explanationist account.

The first view of evidential fit is called *evidential proportionalism* (EP) which says that each and every body of evidence bears some logical or objective probabilistic relation to each proposition. There are two versions of EP. First, S's evidence supports belief *p* only when S's evidence *entails* that the proposition is true.

Logical Entailment (LE): *p* fits S's evidence at *t* iff S's evidence at *t* entails *p*.

Conee and Feldman reject this view because it ties epistemic relations too closely to logical relations (2008, p.94). They argue that LE requires a type of logical omniscience that justified believers needn't possess. That is, there are propositions a person may know, which logically entail other propositions the person cannot even grasp, and surely couldn't know, that follow from what she currently believes. Hence, just because a belief is logically entailed by one's evidence doesn't mean the person is justified in holding that belief. Kevin McCain (2014) argues that LE rules out extremely plausible instances of inductive support for propositions (p.57). It seems that S can have epistemic support for propositions *not* entailed by S's evidence.

That is, a proposition can fit one's evidence when it is the conclusion of an inductive inference from that evidence.

The second way to understand EP is that the evidence makes the proposition *probable* to some degree. This account is superior to LE because it allows for inductive support. This gives us the following account of evidential support (b = background information and *e* = S's evidence):

Probability (PB): *p* fits S's evidence at *t* iff Pr (p|e&b) > Pr (p|b) at *t*.

Conee and Feldman (2008, p.95) and McCain (p.57) argue that the same problem applies to PB that applies to LE. That is, there will cases where this probabilistic relation is beyond the grasp of S. Furthermore, McCain argues that PB allows that propositions that S does not grasp fit S's evidence as long as S's evidence increases the probability of these propositions. But this is counterintuitive.

There is a second possible account of evidential fit called *subjectivism*.

Subjectivism: evidence *e* supports *p* for person S iff S *believes* that *e* entails or makes sufficiently probable *p*.

The idea here is that the support relation depends upon S *believing* that a satisfactory relation holds between *e* and *p*. Conee and Feldman provide two reasons why this account fails. First, it requires a 'meta-level requirement' that justified belief always requires believing that some satisfactory objective relation holds between one's evidence and the content of one's beliefs. But this is the wrong picture of epistemic justification because children and unsophisticated believers don't possess these explicit beliefs about the relationship between their evidence and their conclusions. Additionally, if simply believing that the relation obtains were sufficient, then very 'wild and unjustified beliefs' to the effect that some evidence supports a conclusion would render these conclusions justified. But that's implausible. Finally, if there will be cases in which

the justification of a conclusion depends upon a belief that the evidence is properly related to the conclusion, then this belief about the connection must itself be justified. However, this will result in an intolerable infinite regress and therefore an actual belief in such a connection cannot be required in all cases. Therefore, *Subjectivism* fails as an adequate account of evidential fit (Conee and Feldman, 2008, p.95).[43]

A third possible view of evidential fit involves *non-doxastic seemings*. There are times when someone just reflects on his evidence and it "just seems" to the person that it supports some conclusion.

> **Non-Doxastic Seemings (NDS):** evidence *e* supports C for person S provided it seems to S that *e* has a proper logical or probabilistic relation to C.[44]

Thus, its 'seeming' to S that *e* supports *p* is a non-doxastic state, not a belief that *e* supports *p*. This avoids Conee and Feldman's objection to *Subjectivism* above (and yet NDS is still a brand of subjectivism). A seeming (unlike a belief) is not something that needs justification, so if a seeming can justify a belief then it stops the infinite regress problem stated above. Conee and Feldman reject this view. First, they argue that 'seemings' are not required for justification. If S has a good reason to believe something, then the conclusion is justified regardless of whether one has the additional non-doxastic seeming. Second, Conee and Feldman argue that the existence of these non-doxastic seeming states is unlikely. It might be easy to pick out spontaneous beliefs about connections between propositions and to identify feelings of confidence about these beliefs. However, it is very difficult to recognize any state other than a belief or a level of

[43] One particular brand of subjectivism is advanced by Richard Foley (1987). Foley's subjectivism says that a doxastic attitude fits some evidence that a subject S has iff S would, on reflection, endorse the epistemic principle linking that doxastic attitude to that evidence.

[44] Michael Huemer (2007) advances a similar view called *phenomenal conservatism* which says: If it seems to S that *p*, then, in the absence of defeaters, S thereby has at least some degree of justification for believing that *p* (p.30).

confidence that can be properly characterized as it seeming to one that *e* supports *p* (Conee and Feldman, 2008, p.95-96).

Michael Bergmann provides a nice explanation of how to characterize the sort of objection to NDS that Conee and Feldman support (Bergmann, 2013, p.169-70). 'Objectivity' rejects the view that a doxastic response to an experience (such as a seeming) is fitting as long as it *feels* right to the person.

> **Objectivity**: it is possible for S's belief response to her experience to seem to her to be entirely fitting and epistemically appropriate when in fact it is objectively unfitting (epistemically) (Bergmann, 2013, p.169).

And likewise

> It is possible for S's belief response to her experience (e.g., a sensory experience) to be objectively unfitting epistemically even when—in virtue of the fact that that experience causes the content of that belief to seem true to S—that belief response feels right to S (Bergmann, 2013, p.170).

Bergmann gives the following example of a belief response to an experience that is epistemically *un*fitting for us: <believing, in response to the sensory experience that one is having as one sits in his quiet office looking at a bookcase lined with books, that he's standing on a mountaintop enjoying a view of the clear night sky while listening to a choir singing Handel's Messiah>. *Objectivity* holds that this unfittingness does not dissolve (a) simply because of its coming to seem to the person that that belief response to that experience is epistemically fitting or (b) simply because that belief response to that experience feels right because that experience causes that belief content to seem true (Ibid). I think this is correct and from this point on I adopt the objectivity view.

A fourth possible view of the evidential support relation involves principles of evidence that do not involve principles of deductive and inductive logic. This view was advanced by Roderick Chisolm (1979, p.238). The idea is that we discover these principles by first starting

with the assumption that what a person knows at any time is justified by reference to what is in one's evidence base at that given time and that we know things from various sources. There are a variety of sources of our knowledge: perception, memory, introspection, reason, etc. "Moral consciousness", "intuitive understanding", and "religious consciousness" may also be sources of knowledge (Chisolm, 1982, p.114-15). With these assumptions, we engage in a Socratic project of formulating evidential principles that account for our knowledge derived from these various sources. Together, these principles are supposed to describe all of the cases in which particular propositions fit one's evidence. The problem, as Conee and Feldman argue, is that all of these principles are such that they "…do not derive from any more fundamental or more general ones" (Conee and Feldman, 2008, p.97). The problem isn't that these principles don't correctly describe when our evidence supports believing a particular proposition, rather the issue is that "…there is a troubling arbitrariness and specificity about his [Chisholm's] choice of principles" (Ibid). That is, if there is something about perception, memory, reason, etc., that justifies, then there is something about them that makes this the case. Chisolm's view fails to offer more fundamental principles to explain why these experiences are justifying.[45]

In part II below, I will argue against one more important account of the epistemic support relation that incorporates reliability—a.k.a., 'reliabilist fit'. It is important to discuss this view because it tries to account for the role that reliability plays in epistemic justification. That is, 'reliabilist fit' conflicts with my own two-component view of justification, so it's important to see why reliabilist fit doesn't properly account for the role that reliability plays in justifying a belief. This helps to clear the way for my own account which I will develop later on.

[45] For a more through discussion of Chisolm's view as well as phenomenal conservatism, see Kevin McCain (2014) p.58-62.

II. Juan Comesana and 'Evidentialist Reliabilism'

In this section, I will argue against one other account of the epistemic support relation. Juan Comesana defends a theory of epistemic justification that combines evidentialism and reliabilism. In doing so, he utilizes an account of the epistemic support relation that he calls 'reliabilist fit'. Evidentialist reliabilism replaces talk of evidence *fitting* certain propositions with talk of evidence being reliably connected to the truth of certain propositions. I will argue that Comesana's *evidentialist reliabilism* (ER), although a promising account of epistemic justification, fails because it relies on a reliabilist notion of epistemic support as a necessary component. Remember that on my view, although reliability isn't necessary for epistemic justification, it can add to the degree to which a belief is justified.

II.1 - Against Evidentialism

Comesana begins by arguing that evidentialism is an incomplete theory of doxastic justification. Remember that Conee and Feldman define epistemic justification in the following way:

> **EJ:** Doxastic attitude *D* toward proposition *p* is epistemically justified for *S* at *t* if and only if having *D* toward *p* fits the evidence *S* has at *t*.

Conee and Feldman characterize evidentialism as an *internalist* theory which is the view that a person's beliefs are justified only by things that are *internal* to that person's mental life. However, notice that EJ as defined above is a theory of *propositional justification* and *not* a theory of *doxastic justification*. Thus, EJ cannot provide a *full* account of epistemic justification. EJ only tells us when a doxastic attitude is justified for S (whether S adopts that doxastic attitude or not). It does not tell us what it is for a person's existing doxastic attitude to be justified. For a full account of justification we need to know what it is for S's *existing* doxastic attitude to be justified. Furthermore, doxastic justification must include the fact that S's doxastic attitude is

based on the evidence in S's possession. To account for doxastic justification, Conee and Feldman offer an account of 'well-foundedness' which says:

> **WF:** S's doxastic attitude D at *t* toward proposition *p* is well-founded if and only if:
> (i) having D toward *p* is justified for S at *t*; and
> (ii) S has D toward *p* on the basis of some body of evidence *e*, such that
> a) S has *e* as evidence at *t*;
> b) having D toward p fits *e*; and
> c) there is no more inclusive body of evidence *e*' had by S at *t* such that having D toward *p* does not fit *e*'[46]

Thus, in order for a belief *p* at t to be well-founded, S must have evidence that on balance supports believing that *p*, this evidence in support of *p* must not be defeated by other evidence S has, and S has to believe *p* on the *basis* of S's supporting evidence (Comesana, p.573-74).

Even with WF, Comesana argues that evidentialism is an incomplete theory of epistemic justification for one important reason: a satisfactory account of evidential fit cannot be formulated if we adopt a mentalist account of evidentialism.

> **Evidential mentalism:** If any two individuals are exactly alike mentally, then they are exactly alike evidentially.[47]

If two individuals are exactly alike evidentially they not only have the same evidence but their evidence also justifies the same attitude towards the same propositions. How someone is mentally fixes what evidence that person *possesses*, but also what that evidence is *for*. Comesana argues that evidential mentalism cannot answer the following question:

> **Q:** What necessary and jointly sufficient conditions must evidence e and proposition p satisfy for it to be the case that adopting doxastic attitude D towards p fits *e*?

He says any answer to Q can be thought of as a completion of the following schema:

> **Doxastic attitudes:** Necessarily, adopting D towards *p* fits *e* if and only if…

[46] Kevin McCain points out that clause (i) in WF is actually redundant with the conjunction of a), b), and c) if EJ is assumed (McCain, p.3).
[47] Note that this should refer to *beliefs* being alike justificationally, not individuals.

This schema is a conjunction of these other three schemas:

> **Belief:** Necessarily, believing that *p* fits *e* if and only if…
>
> **Disbelief:** Necessarily, disbelieving that *p* fits *e* if and only if…
>
> **Suspension of judgment:** Necessarily, suspending judgment about *p* fits *e* if and only if…

Q, then, is a precise version of the question '*what is it for a doxastic attitude towards a proposition to fit some evidence?'* (Comesana, 2010, p.574).

According to Comesana, we can categorize the possible answers to Q in two ways. First, it is compatible with evidentialism (including mentalist versions) that we cannot have a full or even a partial answer to Q because it is compatible with evidentialism that epistemic facts are brute facts (Comesana, 2010, p.276).

> **Brutal fit:** Q doesn't have a full, correct and informative answer, because epistemic facts of the form "adopting D towards p fits E" are brute facts (Comesana, 2010, p.580).

Second, there is a class of answers that satisfy the following restriction:

> **Mental fit:** Q can be given a correct, full and informative answer in exclusively mental terms (Ibid).

Comesana points out that evidentialism doesn't entail mental fit. He says that this is clearly the case because evidentialism is compatible with brutal fit, which is itself *in*compatible with mental fit. Thus it follows that evidentialism is compatible with the *negation* of mental fit. Comesana believes that any answer to Q that tries to satisfy mental fit will encounter a fundamental problem:

> …the fundamental problem is that for a doxastic attitude towards a proposition to fit certain evidence, that evidence must be *a good reason* to take that attitude towards that proposition—and, whereas it is plausible to suppose that evidence itself is constituted by mental states, it is not plausible to suppose that what makes evidence *good* is itself constituted by mental states (Comesana, 2010, p.581).

Later I will examine the argument he provides that supports this claim, but for now it will suffice to motivate examining a different answer to Q.

Now, before I explain what Comesana thinks is a satisfactory answer to Q, I will briefly discuss the three problems he thinks reliabilism faces. Then I will explain why he thinks ER can provide a satisfactory answer to Q. I will then discuss why Comesana thinks that this proposed "hybrid" view also avoids the problems that reliabilists face.

II.2 - Reliabilism

Comesana also acknowledges three problems for reliabilism. Comesana offers the following general account of reliabilism:

Reliabilism: S's belief that p is justified if and only if:

1. S's belief that p was produced by a belief-independent process that is reliable; or
2. S's belief that p was produced by a belief-dependent process, the beliefs on which the process depends are justified, and the process is conditionally reliable.

This definition of reliabilism respects the distinction between belief-dependent processes (processes whose inputs include beliefs) and belief-independent processes (processes whose inputs *do not* include beliefs). It is not necessary that belief-dependent processes be *unconditionally* reliable in order to justify a belief. That is, they only need to tend to produce mostly true beliefs if the beliefs that are part of their inputs are true. A belief produced by a belief-dependent process is justified just in case the process is thus *conditionally* reliable and the input beliefs are themselves justified (Comesana, 2010, p.577).

Comesana identifies three main problems with reliabilism. First, that reliabilism isn't *necessary* for justification. Second, that reliability is not *sufficient* for justification. Third, that reliability is *unintelligible*.

The first problem, that reliabilism isn't *necessary* for justification, is expressed by the New Evil Demon Problem (NED) that I discussed earlier.

> **NED:** The objection starts by having us consider victims of a Cartesian evil demon. Everyone agrees that these victims don't have knowledge (if only because most of their beliefs are false, but not only because of that), but, Cohen argues, we should admit that they can be just as justified as we are. To be sure, some of these victims engage in wishful thinking, hasty generalization, and other epistemic vices—their beliefs are not justified. But some other victims are careful reasoners and reasonable "perceivers": they take their experience at face value when there is no conflicting information, they commit as few fallacies as we do, etc. We better say that the beliefs of these victims are justified—for what is the relevant difference between them and us? But reliabilism seems to have the consequence that the beliefs of the victims of an Evil demon can never be justified, because what the demon does is to make the belief-forming processes of his victims nearly completely unreliable (Comesana 2010, p.577-78).

NED is supposed to show that reliabilism seems to have the consequence that the beliefs of a victim of an evil demon cannot be justified because they are not reliably produced. But intuitively the victim's beliefs can be justified if the victim in question is a careful epistemic agent.

Comesana offers what he calls 'indexical reliabilism' as a means to deal with NED's challenge to reliabilism's necessity for epistemic justification (2010, p.579).

Indexical Reliabilism (IR): S's belief that p is justified if and only if:

1. S's belief that p was produced by a belief-independent process that is *actually* reliable;

   or

2. S's belief that p was produced by a belief-dependent process, the beliefs on which the process depends are justified, and the process is conditionally *actually* reliable.

Here, 'actually' is to be understood as having a two-dimensional semantics. That is, actual reliability expresses two different propositions associated with any attribution of justification: the "diagonal proposition" which holds that the belief was produced by a method that is reliable in S's world (the world where it is believed), and the "horizontal proposition" which holds that the

belief was produced by a process that is reliable in the world where the proposition is being considered (the world in which justification is being *ascribed* to S's pertinent belief via an utterance of the sentence 'S's belief that p is justified'). To have knowledge, a belief must be diagonally justified. But even subject S, who is the victim of an evil demon, can be justified in the horizontal sense. Victims of an evil demon can have justified beliefs if those beliefs are produced by processes that are reliable in our world.

The second problem, that reliability is not *sufficient* for justification is expressed in the Norman the Clairvoyant case. Norman's belief is not justified since he has no evidence or reasons of any kind for or against the possibility of such a cognitive ability, or evidence of the President's whereabouts (Comesana, p.578).

The third problem, that reliability is unintelligible, is best expressed by the *generality problem*. The generality problem is the problem that any *token* process that produces a belief will be an instance of indefinitely many *types* of processes. Each type will have its own level of reliability, normally distinct from the levels of reliability of other types. The problem is determining which repeatable type should be selected for purposes of assigning a determinate reliability ratio to the process token (Ibid).

In order to understand Comesana's solutions to the second and third problems, we have to first understand his account of ER. To understand ER Comesana argues for a proper account of 'fit'.

II.3 - 'Fit'

Comesana proposes the following answer to what 'fit' consists in:

Anti-mental fit: Q can be given a correct, full and informative answer, but not in mental terms exclusively.

Comesana proposes ER as a way to answer Q that satisfies anti-mental fit. In developing this answer Comesana appeals to the following 'full and informative' answer to Q (Ibid):

Reliabilist fit (RF):

Belief: Necessarily, believing that *p* fits *e* for subject S if and only if:

1. *e* doesn't include any beliefs of S and the connection between S's having *e* and *p* is actually reliable; or
2. *e* includes beliefs of S, all of these beliefs are justified, and the connection between S's having *e* and *p* is conditionally actually reliable.

Disbelief: Necessarily, disbelieving that *p* fits *e* for subject S if and only if:

1. *e* doesn't include any beliefs of S and the connection between S's having *e* and *not-p* is actually reliable; or
2. *e* includes beliefs of S, all of these beliefs are justified, and the connection between S's having *e* and *not-p* is conditionally actually reliable.

Suspension of judgment: Necessarily, suspending judgment with respect to *p* fits *e* for subject S if and only if neither believing nor disbelieving that *p* fits *e* for subject S.

RF replaces talk of evidence *fitting* certain propositions with talk of evidence being *reliably connected to the truth of certain propositions*. Comesana says that given the fact that we want to avoid postulating primitive (brute) epistemic facts, and given the shortcomings of mental fit, he believes that reliabilist fit is a promising complement to evidentialism.

II.4 - Proto-Evidentialist Reliabilism

Adding RF to IR helps deal with the challenge to the sufficiency of reliability for justification. Remember that IR consists of a two-dimensional semantics for 'actual'. Comesana argues that the sentence "clairvoyance is actually reliable" (uttered in a context where Norman is being discussed) expresses two propositions: a true diagonal proposition (because clairvoyance is reliable in Norman's world), and a false horizontal proposition (because clairvoyance is not reliable in our world). But again, Norman has no *evidence* for or against his clairvoyant ability *or* regarding the whereabouts of the President. The belief just pops into his head. This gives us

the intuition that Norman is not justified in believing that the President is in New York City, since his belief isn't based on any evidence (Comesana, 2010, p.582).

Now, in order to account for the dependency of justification on evidence, IR can be revised in the following way:

> **Proto-evidentialist reliabilism (Proto-ER):** S's belief that p is justified if and only if that belief was produced by a process P which includes some evidence e and:
>
> 1. e doesn't include any beliefs of S and P is actually reliable; or
> 2. e includes beliefs of S, all of these beliefs are justified, and P is conditionally actually reliable.

Proto-ER properly diagnoses the Norman case because now his belief is unjustified because it was produced by a process that doesn't provide any evidence (evidence of the possibility of Norman's clairvoyant ability or evidence regarding the President's whereabouts). Thus, Proto-ER has advantages over both reliabilism and evidentialism. Unlike reliabilism, Proto-ER is successful at dealing with counterexamples like Norman the Clairvoyant, and unlike evidentialism, it includes a non-pessimistic answer to Q. However, in order to deal with the third problem, the unintelligibility of reliabilism, a further integration of evidentialism and reliabilism is required (Comesana, 2010, p.583).

II.5 - Evidentialist Reliabilism

Proto-ER provides a partial solution to the generality problem because it rules out process types that don't include evidence. However, there are still too many process types that include evidence under which any given token process of belief-production falls, and it is not guaranteed that they will all have the same reliability ratio. For example, the token process that produces <S's belief that there is a computer in front of him> falls under all the following processes: <perceptual process>, <visual process>, <visual process while sober>, <visual process while sober and not wearing socks under bad lighting conditions>, etc. Comesana says that intuitively,

the answer to the generality problem includes pointing out that some of the process types are descriptions of the psychological processes that actually produced the belief, while other types include irrelevant information, or leave out relevant information. For example, whether S is wearing socks has nothing to do with the formation of S's belief and any type that doesn't include information about features of S's experience that S takes notice of in forming his belief is likely to leave out relevant information (Ibid).

Importantly, Comesana points out that any complete epistemological theory will include a theory of doxastic justification. Evidentialism is not incomplete because it has a theory of well-foundedness, which appeals to the notion of *basing* a belief on some evidence. In order for a person's belief to be justified, he must *base* his belief on his evidence that justifies it, i.e.—his belief is well-founded. Comesana further points out that *any* theory that accounts for doxastic justification must include a basing relation (Comesana, 2010, p.584).

Comesana believes that the evidentialist theory of well-foundedness provides an important insight into how to solve the generality problem:

> Now, the evidentialist theory of well-foundedness has it that every justified belief will be based on some evidence. If so, then the token process that produced that belief will always be an instance of a type of process of the form producing a belief that p based on evidence e—for instance, producing a belief that there is something gray in front of me based on the fact that it looks as if there is something gray in front of me, etc. That type, I propose, is the one whose reliability is relevant for the justification of a belief (Comesana 2010, p.584).

Comesana thus provides what he thinks is an adequate solution to the generality problem by providing an account of how to specify which token process is always involved in producing the relevant belief. In each token process there is some evidence that produces S's belief. If we can identify the evidence that the belief is based on then we can identify the token process that produced the belief. The token process always being identifiable as the specific evidence e (*it*

looks as if there is something gray in front of me) that produces the specific belief p (*there is something gray in front of me*).

More clearly expressed, the solution to the generality problem (based on the evidentialist definition of well-foundedness) is ER. ER is a form of evidentialism where the epistemic notion of *fit* is replaced by the notion of *reliability*.

ER: A belief that *p* by S is justified if and only if...

1. S has evidence *e*;
2. the belief that *p* by S is based on *e*; and either
    a. *e* doesn't include any beliefs and the type *producing a belief that p based on evidence e* is actually reliable; or
    b. *e* includes other beliefs of S, all of those beliefs are justified and the type *producing a belief that p based on evidence e* is conditionally actually reliable (Comesana 2010, p.584).

ER is no longer Proto-ER because ER fully incorporates the evidentialist insight that the evidence possessed by S plays a crucial role in determining whether a particular belief is justified for S. That is, Proto-ER only requires that there is some evidence in the belief-forming process, while ER requires that the justification of the belief depends on whether the subject *used* that evidence in forming the pertinent belief. For ER the justification of the belief depends upon whether the subject *used* the evidence to form the belief (Ibid).

II.6 - Objection to Evidentialist Reliabilism

Although Comesana's account of ER seems to be a promising account of doxastic justification, ultimately it doesn't work. The following objections are aimed at both RF and ER.[48] The first objection is aimed directly at ER.

Consider the following case as a counterexample to ER. Vic is the unfortunate and unwitting victim of a Cartesian Evil Demon. Both Vic and the Evil Demon are the only things

[48] These objections are attributed to Michael Pace (Chapman University) in a communication I had with E.J. Coffman.

contained in the world they both inhabit, world 'W'. The Evil Demon manipulates Vic's sensory experience such that all of the beliefs Vic forms on the basis of these experiences are false beliefs. For example, the Evil Demon gives Vic a misleading visual experience as of hands (even though Vic has no hands). On the basis of this misleading visual experience, Vic forms a non-inferential belief that he has hands. Upon reflecting on his "hands" belief, he attributes justification to it by uttering (thinking it in his own mind) the following sentence: 'My non-inferential belief that I have hands is justified.'

Now, since ER incorporates IR, we are to understand 'actually' as having a two-dimensional semantics. There are two different propositions associated with any attribution of justification: the diagonal proposition, which says that the belief was produced by a method that is reliable in the world where it is believed, and the horizontal proposition, which says that the belief was produced by a process that is reliable in whichever world the proposition is *considered*. Therefore, Vic's utterance 'My non-inferential belief that I have hands is justified' is true only if 1) forming a belief on the basis of visual experience is reliable in the world where it is believed (Vic's world W) or 2) forming a belief on the basis of visual experience is reliable in the world in which the proposition is considered, i.e.—in the world where the justification-attribution is uttered (also W since Vic is the person attributing justification to his belief that he has hands). Therefore, ER holds that Vic's utterance 'My non-inferential belief that I have hands is justified', is true only if visual experience is reliable in world W (the world Vic inhabits). However, visual experience is *not* reliable in W. Hence, ER holds that Vic's utterance 'My non-inferential belief that I have hands is justified' is false, i.e.—neither the diagonal nor the horizontal proposition is true. But this result is implausible because it seems that Vic's utterance

‘My non-inferential belief that I have hands is justified’ is true. Therefore, this seeming truth of Vic’s utterance weighs significantly against ER.

The idea behind this objection is that even though visual experience isn't reliable in W, the evidence Vic has (seeming to have hands) justifies his non-inferential belief ‘I have hands' because his visual experience (absent any defeating evidence) gives him a very good reason to believe that he has hands. If true, this is compatible with the view that you can have a justified false belief because you have strong evidence that supports that belief.

This same objection can also be used against the left-to-right conditional of RF. In particular, the objection indicates that the left-to-right conditional of RF is false. Simply put, we need to modify Vic’s case so that Vic ascribes to his hand belief the property of *fitting* the evidence it’s based on – viz., visual experience as of hands. So suppose that Vic utters (in the same manner as before) ‘My hand belief *fits* my visual experience as of hands’. Intuitively, there seems to be something correct about Vic’s utterance. That is, it seems to be expressing a true statement. But, RF implies that Vic’s utterance expresses a truth only if “hand-ish experience” reliably indicates “hand presence” in W (the world Vic inhabits). However, by hypothesis “hand-ish experience” is not a reliable indicator of “hand presence”. Therefore, RF yields the counterintuitive result that Vic’s utterance is false, which strongly weighs against RF. This strongly supports that reliability is not a necessary condition for a belief to be justified. It also supports the evidentialist view that a belief is justified if it properly fits one’s evidence that it is based on. If so, then evidentialist fit is a sufficient condition for a belief to be justified *period* (at the very least).

III. Externalist Evidentialism

One important point needs to be made before I move on to discussing accounts of *explanationist evidentialism*. Explanationist Evidentialism (EE) is a *mentalist* and an *internalist* theory of epistemic justification. It is a mentalist theory of justification because it holds the following two principles:

1. The justificatory status of a person's doxastic attitudes strongly supervenes on the person's occurrent and dispositional mental states, events, and conditions.
2. If any two possible individuals are exactly alike mentally, then they are alike justificationally, e.g., the same beliefs are justified for them to the same extent (Conee and Feldman, 2004a, p.56).

Number 1 says that justification supervenes on mental states and 2 is an implication of 1. EE is an internalist and mentalist account of justification. However, not all accounts of evidentialism are internalist. Timothy Williamson offers an externalist evidentialist view. He identifies one's evidence with one's knowledge. That is, one's evidence consists of the totality of propositions that one knows. Famously, he holds that E=K (evidence = knowledge). Williamson's account of evidentialism is externalist because knowledge consists of true propositions and truth is external to the thinker. So if E=K, then evidence consists of only true propositions. Thus, a belief's justification depends upon truth, which is external to the thinker. Williamson's view is known as *propositionalism* which says that all evidence is made up of propositions. Very briefly I would like to discuss some strong reasons to doubt that propositionalism is true.

There are a number of good reasons to doubt propositionalism. For the sake of brevity, I will provide what I think are the two biggest reasons to doubt propositionalism.[49] McCain argues that propositionalism is too restrictive as an account of evidence since it is highly

[49] For a more thorough discussion of the issues with propositionalism and Williamson's E=K thesis, see McCain (2014) Chapter 2 and Feldman and Conee (2008).

plausible that your *experiences* count as evidence. "Your experience of being in pain is evidence for you that you are in pain. Your experience of being hungry is evidence for you that you are hungry. Your experience of a book looking blue is evidence for you that the book is blue. And so on" (McCain, p.19). Conee and Feldman write, "Experience is our point of interaction with the world—conscious awareness is how we gain whatever evidence we have" (Conee and Feldman, 2008, p.87). McCain points out that none of this means that only experience counts as evidence. It is very plausible that things like beliefs count as evidence as well. However, if a belief is to count as evidence for anything other than the fact that the person holds that belief, then it must be a justified belief, i.e.—it must be a belief for which the person has good evidence. This good evidence could be other justified beliefs or experiences, but eventually it is very likely that the evidence will 'bottom out' in experiences (perceptual, introspective, memorial, intuitive, etc.). Thus, beliefs can be evidence, but it's likely that they are not 'ultimate evidence'. McCain thinks it is very plausible (and likely) that all ultimate evidence is experiential and that "…all other evidence is evidence in virtue of bearing appropriate relations to ultimate evidence" (McCain, p.19-20).

One other reason to doubt propositionalism and therefore E=K is that evidence need not be factive. If evidence consists solely of mental states, then the type of mental states in question are non-factive mental states (McCain, p.10). Non-factive mental states are mental states someone can be in even if they misrepresent the world. Factive mental states, on the other hand, are mental states that actually represent the world. For example, seeing that there is a tree is a factive mental state since S can only be in the mental state of seeing that there is a tree if there actually is a tree. However, seeming to see that there is a tree is a non-factive mental state because S can seem to see that there is a tree even when there isn't one (e.g., the seeming is a

hallucination). We can see that evidence is non-factive because it doesn't always point to the truth. Good evidence can support a false proposition. In other words, even good evidence is fallible.

IV. Explanationist Accounts

Given the misgivings of the above accounts of the epistemic support relation, I want to now discuss what I think is a far more promising account: *explanationist fit*. However, in order to have a full account of categorical doxastic justification, we need to know not only what it is for a doxastic attitude towards some proposition to *fit* some evidence that the subject possesses, but what it is for a subject to *base* his doxastic attitude toward some proposition on that evidence (well-foundedness). So, explanationist fit is only one important component of a fully fleshed out evidentialist theory of categorical doxastic justification, i.e.—well-foundedness.

To begin my discussion of explanationist fit, I propose the following working definition, which will need a great deal of unpacking:

> **Evidential Fit:** For belief *p* to fit S's evidence at a given time is for *p* to be the best explanation of that evidence.

Having a strong conception of evidential fit is important for understanding how evidentialism accounts for epistemic support. Conee and Feldman say the following about epistemic support:

> We believe that the fundamental epistemic principles are principles of best explanation. Perceptual experiences can contribute toward the justification of propositions about the world when the propositions are part of the best explanation of those experiences that is available to the person. Similarly, the truth of the contents of a memory experience may be part of the best explanation of the experience itself. Thus, the general idea is that a person has a set of experiences, including perceptual experiences, memorial experiences, and so on. What is justified for the person includes propositions that are part of the best explanation of those experiences available to the person. Likewise, one's inferences justify by identifying to one further propositions that either require inclusion in one's best explanation for it to retain its quality or enhance the explanation to some extent by their inclusion (Conee and Feldman, 2008, p.97-98).

Thus, according to Feldman and Conee, the fundamental support principle of evidentialism is a principle of best explanation. We can formulate this basic principle like so:

> **Best Explanationism:** Belief *p* is justified for S if and only if *p* is the best explanation of the evidence S has at a given time.

This principle is a principle of propositional justification only, so later I will discuss an account of the basing relation in order to turn explanationism into a theory of categorical doxastic justification (in this case, well-foundedness).

Now, there are a number of things that need to be addressed in order to understand Conee and Feldman's view. As they point out, we need a better understanding of what it means to have a best "available" explanation. Is this "availability" qualifier referring to the best explanation available *simpliciter* or the best explanation available to the *subject*? If the former, then there will be worries about the right-to-left direction of 'Best Explanationism'. For example, the best scientific explanation of S's current experience will include detailed scientific theories or logical consequences that S is unable to grasp or understand. So, there will be certain explanations that will constitute the best explanation for why S has that evidence, but aren't available to S as an explanation. Even if Conee and Feldman add that the type of availability in question includes the best explanation available to the *subject*, then this qualifier itself will need to be explained (Conee and Feldman, 2008, p.98).

Conee and Feldman make the following claim about the "availability" qualifier:

> The best available explanation of one's evidence is a body of propositions about the world and one's place in it that make best sense of the existence of one's evidence. This notion of making sense of one's evidence can be equally well described as fitting the presence of the evidence into a coherent view of one's situation. So it may be helpful to think of our view as a non-traditional version of coherentism. The coherence that justifies holds among propositions that assert the existence of the non-doxastic states that constitute one's ultimate evidence and propositions that offer an optimal available explanation of the existence of that evidence (Ibid).

However, there are still important questions that remain. If it is this 'coherence' that justifies, then is it an objective coherence that holds among one's body of propositions that does the justifying or is it the fact that it *seems* to the subject that such coherence obtains? It suffices to say that more needs to be said about this "availability" qualifier. Conee and Feldman admit that much more work needs to be done to develop their view, but they believe they have cast doubt on competing views of epistemic support and offered a more plausible account. In what follows, I will discuss how Conee and Feldman's position has been developed in a recent work by Kevin McCain (2014).

IV.1 - Kevin McCain's Account of Evidential Fit

Conee and Feldman suggest that what is justified for a person are propositions that are part of the best explanation of the person's experience available and the conclusions of inferences that would enhance that explanation. They admit that "…there are important details of this account that are yet to be developed" (Conee and Feldman, 2008, p.98). Kevin McCain attempts to develop this explanationist account.[50]

The first step McCain takes to develop this theory is to define the following important terms:

Explanation:

p fits S's evidence, *e*, at *t* iff *p* explanatorily coheres with *e* at *t*.

Explanatory Coherence:

p explanatorily coheres with *e* at *t* iff *p* is part of the best explanation available to S at *t* for why S has *e*.

By combining these two we get

[50] McCain's account of evidentialism is quite lengthy, so I'm offering what I think are the essential elements of his theory that are directly relevant to my project.

Best Explanation (BE):

> *p* fits S's evidence, *e*, at *t* iff *p* is part of the best explanation available to S at *t* for why S has *e*.

According to BE, the basic epistemic principle is a principle of best explanation (McCain, p.63). But McCain points out the inadequacy of BE.

McCain points to two counterexamples by Keith Lehrer (1974) and Alvin Goldman (2011) that he thinks convincingly show that BE is an unsatisfactory account of evidential fit. Lehrer's example involves using the Pythagorean Theorem to figure out that a particular mouse is five feet from a particular owl:

Pythagorean Theorem:

> Imagine that I am standing with my toe next to a mouse that is three feet from a four-foot-high flagpole with an owl sitting on top. From this information concerning boundary conditions and the Pythagorean Theorem, which we here construe as an empirical law, we can deduce that the mouse is five feet from the owl (McCain, p.64; Lehrer, p.166).

Lehrer claims that this counterexample works against explanationist accounts of justification because he is "…completely justified in his belief that the mouse is five feet from the owl", but he "…has no explanation of why the mouse is five feet from the owl" (Lehrer, p.178). His justification for this belief doesn't rely on any explanatory relations whatsoever. Instead, "…it is enough that the man knows the Pythagorean Theorem, the distance to the pole, and the height of the pole, and deduces the conclusion" (Ibid). The assumption is that logical relations are *not* explanatory relations.

Alvin Goldman makes a similar point:

Animals:

> I think there are two squirrels on my deck, and I think there are two birds. So I infer that there are (at least) four animals. Presumably, this arithmetic inference is justified. Is it a case of explanatory inference? Surely not. How does there being four animals *explain* there being two squirrels and two birds? It doesn't. Still, here is a justified belief that some

> epistemic principle must cover. But that principle, in turn, cannot be grounded in terms of best explanation (McCain, p.64; Goldman, 2011, p.277-78).

In this case there is justification for thinking that there are four animals on the deck, but explanationist accounts of epistemic support (such as BE) cannot account for this justification. The proposition <there are (at least) four animals> is a logical consequence of Goldman's evidence that there are two squirrels and two birds on his deck, but there being four animals on his deck does not explain there being two squirrels and two birds on his deck. Thus, this is a case in which a proposition is justified for someone because it fits the person's evidence, but the proposition is not itself part of the best explanation of the person's evidence. Therefore, BE fails as an adequate account of evidential fit (McCain, p.64-5).

In response to the inadequacy of BE, McCain offers a new account of evidential fit that respects the fact that not all fundamental epistemic principles are principles of best explanation. An adequate account of evidential fit cannot simply appeal to explanation. Any account of evidential fit must appeal to logical consequences of best explanations. McCain thus reformulates epistemic support as

> *Explanationist Fit* (EF):
>
> *p* fits S's evidence, *e*, at *t* iff either *p* is part of the best explanation available to S at *t* for why S has *e* or *p* is available to S as a logical consequence of the best explanation available to S at *t* for why S has *e* (McCain, p.65).

Before moving on, McCain makes some important points about EF. First, when we discuss whether S is justified in believing that *p* because it fits S's evidence, "S's evidence" is S's *total evidence* (McCain, p.65). One's total evidence is a subset of one's *total possible evidence*. A person's total possible evidence is all and only the information a subject has stored in his mind at a given time (conscious, unconscious, retrievable, irretrievable, etc.). One's total evidence is the

subset of S's total possible evidence that is relevant for justification. One's total evidence includes any defeating evidence one might have.

Second, in order for *p* to be part of the *best* explanation available to S at *t*, there must not be any contrary explanation with respect to *p* available to S at *t* that is as good of an explanation as, or better than, the best explanation of which *p* is a part. But this doesn't mean that there must be one unique best explanation. That is, S can have two explanations available to her that are both *equally* good explanations of why S has the evidence she does. What is key is that there are no explanations available to S that do not contain *p* and are just as good as the best explanation she has available. Finally, what determines whether one explanation is better than another is a matter of which one is more explanatorily virtuous, i.e.—explanatory power, simplicity, scope, etc. (McCain, p.65-66).

Third, there can be cases in which the best explanation available is a very poor explanation. Belief *p* may fit S's total evidence but it may not fit her evidence all that well. In which case, the amount of justification *p* has for S is minimal. It might be just above the justification she has for believing ~*p*. This simply means that one can be more or less justified (to different degrees) in believing some proposition. The better the explanation that *p* is a part of, the better *p* fits S's evidence and therefore the higher degree of justification *p* has for S (McCain, p.66). However, I don't think McCain is right when he says that a very poor explanation can provide some minimal amount of justification for a belief. Before moving on I'd like to explain why I think McCain is wrong about this.

Like McCain, Ted Poston (2014) adopts an explanationist view of epistemic justification in which explanationism is a mentalist evidentialist account of epistemic justification. Poston holds the position that "…one's normative standing in the space of reasons is constituted by

one's explanatory position" (p.69) and that the justificatory status of one's belief is a function of the explanatory virtues or merits of one's system of beliefs (p.80).

Poston argues that the virtues relevant to explanationism and justification are *conservativism, simplicity, and explanatory power*.[51] Briefly, 'conservatism' is the view that belief confers some positive merit on the proposition that is believed.[52] 'Simplicity' roughly holds that given two theories which both explain some phenomena, the simpler theory is the better of the two. 'Explanatory power' involves a theory's ability to explain a wide range of phenomena. These virtues are central in understanding Poston's explanatory view of epistemic justification which says:

> **EX-J**: S has justification for believing p if and only if p is a member of a sufficiently virtuous explanatory system, *E*, and *E* is more virtuous than any p-relevant competing system *E'* (Poston, p.90).[53]

Poston makes an important point about EX-J that is relevant to the present discussion and I think it bolsters McCain's view. Poston says that for p to be justified for S it must be part of a *sufficiently* virtuous explanatory system. This condition deals with the objection that sometimes the best explanation isn't a very good explanation. In other words, in some cases an explanation may be the best among a set of really lousy explanations in which case it doesn't seem like the pertinent belief is justified. Poston says that such explanations aren't sufficiently virtuous and therefore such explanations aren't sufficient for justification. That is, in cases where the best explanation is a lousy explanation, EX-J doesn't entail that p is justified (Poston, p.88). I think this is the correct analysis. This is an important condition that I think any explanationist view

[51] I will not discuss whether there this is an acceptable list of virtues. I'm simply providing an example of how one explanationist view makes sense of explanatory virtues. I don't claim that this view is definitive.
[52] Conservatism is a very controversial view in epistemology and a discussion of the relevant controversy goes beyond the scope of this dissertation. I simply offer Poston's explanationist view as a way for the reader to see what a plausible view of best explanations might look like. For a more thorough discussion of the issues surrounding conservativism see Poston (2014) Chapter 2.
[53] Again, this should say '*p* is justified for S' and not 'S has justification for believing *p*'.

should have. It is also in line with Feldman's point that mere evidential support isn't enough for propositional justification. Recall the example Feldman gives to make this point clear: if you learn from a reliable weather report that there's a 51% chance of rain tomorrow, then while your total evidence now barely supports the proposition that "It'll rain tomorrow", you're not yet justified in believing it'll rain tomorrow. Evidential support must be *sufficiently* strong for propositional justification.[54] That's not to say that the strength of evidential support must be extremely high in all cases in order for a belief to be justified; it just means that a belief needs stronger evidential support than what McCain allows. I think this same condition applies to best explanations. That is, for a best explanation to justify a belief, that explanation must be sufficiently virtuous. Thus, not all best explanations are justifying explanations. McCain's view on this matter is too permissive.

Finally, McCain provides an account of what it means for an explanation of why S has *e* to be *available* to S at *t* and what it means for a proposition to be *available* to S as a logical consequence of the best explanation available to S at *t* for why S has *e*. First, McCain says that at *t* S has *p* available as part of the best explanation for why S has *e* at *t* if and only if

> At *t* S has the concepts required to understand *p* and S is disposed to have a seeming that *p* is part of the best answer to the question "why does S have *e*?" on the basis of reflection alone (p.67).

By "seeming" McCain adopts an account defended by Chris Tucker (McCain, p.67; Tucker, 2013) called the "Experience View", which says that seemings are mental states with propositional content and a particular phenomenology—they "have the feel of truth, the feel of a state whose content reveals how things really are" (Tollhurst, 1998, p.298-299). On the 'Experience View', seemings are not beliefs or inclinations to believe but are either *sui generis*

[54] See footnote 31 in Chapter Three.

propositional attitudes or experiences. The idea is that if seemings are a *sui generis* propositional attitude, then they are the sort of thing that is like a belief, desire, etc., that takes a proposition as its object, but they don't fall under any of the other categories of propositional attitudes.

Importantly, the "availability" qualifier does not *require* that S actually have the seeming in question. Instead, S must have the *disposition to have* such a seeming. McCain cites Peter Markie's (McCain, p.67; Markie 2013) "Knowledge How Proposal" to explain how this disposition should be understood. The idea is that S has the pertinent disposition by virtue of having certain information as background evidence. Markie says that this background information helps determine the character of the seeming and helps sustain it (McCain, p.67; Markie, 2013, p.263). What this means is that this disposition to have the necessary seeming is a type of evidence that connects or links *p* to *e*. Hence, this disposition to have a seeming is part of S's total evidence. In other words, the disposition to have this sort of seeming is *intermediate* evidence that connects *p* to *e*. The *ultimate evidence* connecting *p* to *e* is made up of the background evidence S has that sustains this disposition to have the pertinent seeming (McCain, p.81).

To be clear, McCain is not adopting a subjectivist view along the lines of phenomenal conservatism. As we saw earlier, McCain rejects such a view. McCain is not claiming that if it seems to S that *p* is part of the best answer to the question "Why does S have *e*?", then it is part of the best answer. McCain says that "…it is an open question whether *p* is in fact part of the best answer to this question" (Ibid). He says "It is plausible that S could have a mental state in which it seems to her that *p* is part of the best answer to "why does S have *e*"? but in fact it is not" (Ibid). The idea with availability here is that S doesn't have to consciously at *t* recognize/think of/believe/have a seeming/etc., that *p* is part of the best explanation of her

evidence in order for *p* to be justified for her. Such a requirement would over-intellectualize justification. Instead, S has to have a disposition upon reflection to have a seeming that *p* is part of the best explanation of her evidence in order to have *p* available in the right way. It has to be possible that S have the seeming that *p* is part of the best explanation for why she has *e*. This weaker requirement is one that can be met by even unreflective agents (such as children). So, when S has *p* available as a best explanation in this sense, then she has (at least some) justification for believing that *p* even if *p* isn't part of the actual best explanation of *e*. Roughly, the idea is that one doesn't need to be infallible when it comes to recognizing best explanations in order to have justification. An example might help here. Suppose that S has three potential explanations of *e* that she can understand, E1, E2, and E3. Further, suppose that *p* is part of E1. It is possible for S to be justified in believing *p* because it is available to her as part of the best explanation of *e*—even if in fact according to the correct (objective) standards of evaluating explanations, E2 is a better explanation. Thus, for S to have *p* available as part of the best explanation for why S has *e* at *t*, then (in addition to having the relevant concepts) it must simply be possible for S to have the experience (seeming) that *p* is part of the best explanation for why S has *e* at *t*.[55]

Now, in terms of logical entailment, McCain says: "…*p* is available to S as a logical consequence of the best explanation available to S at *t* for why S has *e* if and only if

> S has the concepts required to understand *p* and S is disposed on the basis of reflection alone to have a seeming that *p* is entailed by the best explanation available to S at *t* for why S has *e*" (p.68).

[55] I thank Kevin McCain (in communication) for clarifying this availability qualifier.

S has the seeming that *q* (for example) follows from the best available explanation of her evidence in the sense that someone who does not have the concept of *modus ponens* can have the seeming that if *p* and *if p, then q* are true, then *q* must be true (Ibid).

McCain provides a couple of examples to show that EF is an acceptable account of evidential fit. In the first example he says that in normal conditions (and S has no defeaters) S's visual experience as of a red block provides her with epistemic support for the proposition <there is a red block>. He argues that EF delivers precisely this result. That is, the proposition <there is a red block> is part of the best explanation of why S has an experience as of a red block that is available to S. S has the concepts required to understand <there is a red block> and is disposed to have a seeming that <there is a red block> is part of the best answer to the explanatory question, "Why am I having this visual experience?" or "Why am I having *this*?" on the basis of reflection alone. McCain says that in a normal case there will not be an equally good or better explanation that does not include <there is a red block> available to S. So, <there is a red block> will fit S's (assuming S is an ordinary person) evidence. Therefore, S will have justification for believing that <there is a red block> when S has the visual experience as of a red block. The same goes for other cases of perception (McCain, p.70).

In order to understand how EF works in accounting for logical entailment (in addition to seeing how EF works more generally), McCain applies it to Lehrer's 'Pythagorean Theorem' and Goldman's 'Animal' examples. Let's look at each of these.

In Lehrer's example, he has evidence that justifies him in believing that the mouse is three feet from the flagpole, the owl is on top of the four-foot tall flagpole, and the Pythagorean Theorem is true. Additionally, he is presumably justified in believing other relevant information (what three squared equals, what four squared equals, etc.). Given EF, the following

propositions are all available to Lehrer: <the mouse is three feet from the four foot tall flagpole>, <the owl is on top of the four foot tall flagpole>, <the Pythagorean Theorem is true>, <the Pythagorean Theorem says that the hypotenuse of a triangle that has a three foot long side and a four foot long side is five feet long>, etc. These propositions are part of the best explanation available to Lehrer. Thus, <the mouse is five feet from the owl> is a logical consequence of this explanation. McCain says that it is very plausible that <the mouse is five feet from the owl> is available to Lehrer as a logical consequence of the best explanation available to him at *t* for why he possesses the evidence he has. After all, McCain says, Lehrer claims that he deduces this conclusion. "So, *EF* yields the result that <the mouse is five feet from the owl> fits Lehrer's evidence, so he is justified in believing that the mouse is five feet from the owl" (McCain, p.74).

McCain doesn't think 'Animals' is a problem for EF either. Part of the best explanation available to Goldman for his current visual experience and his background evidence in this case is that <there are two squirrels on the deck>, <there are two birds on the deck>, and <squirrels and birds are both animals>. Goldman has the concepts required to understand <there are four animals on that deck> and he is disposed based on reflection alone to have a seeming that <there are four animals on the deck> is entailed by the best explanation available to him at *t* for why he has the evidence that he possesses. Thus, Goldman has justification for believing that <there are four animals on the deck> because this proposition fits his evidence (McCain, p.74-5).

EF provides a piece of McCain's full account of epistemic justification. It provides necessary and sufficient conditions for S to have justification for adopting a particular doxastic attitude (propositional justification):

EF-EJ:

I) Believing *p* is epistemically justified for *S* at *t* if and only if at *t* S has considered *p* and:

1) *p* is part of the best explanation available to S at *t* for why S has her occurrent non-factive mental states and the non-factive mental states that she is disposed to bring to mind when reflecting on the question of *p*'s truth

OR

2) *p* is available to S as a logical consequence of the best explanation available to S at *t* for why S has her occurrent non-factive mental states and the non-factive mental states that she is disposed to bring to mind when reflecting on the question of *p*'s truth

II) Withholding judgment concerning *p* is epistemically justified for S at *t* if and only if at *t* S has considered *p* and neither believing *p* nor believing ~*p* is epistemically justified for S (McCain, p.79).[56]

However, in order to have a full account of doxastic justification (well-foundedness), McCain offers an account of what it takes for S to have a particular doxastic attitude *based on* her evidence. I turn toward that discussion now.

IV.2 - McCain's Account of the Basing Relation

McCain adopts a version of the *causal* account of the basing relation. Causal accounts usually entail that in order for S's belief to be based on S's evidence, that evidence must contribute *causally* to S's having that belief. In contrast, McCain rejects *doxastic* accounts of the basing relation because he thinks that such accounts engage in over-intellectualization and/or fall prey to a vicious regress.[57] The idea behind doxastic accounts is that the basing relation requires a meta-belief about the support that the evidence provides for the belief. That is, such accounts require S to have the meta-belief that her evidence supports believing that *p* in order for her belief that *p* to be based on that evidence. McCain says that such accounts are guilty of over-

[56] Behind EF-EJ is McCain's principle (MVP*) that says:

> S has *p* available as evidence relevant to *q* at *t* iff at *t* S is currently aware of *p* or S is disposed to bring *p* to mind when reflecting on the question of *q*'s truth (McCain, p.51).

[57] For a discussion of doxastic accounts see Keith Allen Korcz. 1997. "Recent Work on the Basing Relation." *American Philosophical Quarterly,* 34: 171-91; and Joseph Tolliver. 1982. "Basing Beliefs on Reasons*." Grazer Philosophische Studien,* 15:149-161.

intellectualization. For example, if S were to see a tree in perfect viewing conditions, her belief that <there is a tree> is not based on her evidence unless she also believes that her visual evidence is good evidence for thinking that her belief <there is a tree> is true. McCain thinks it's obvious that we have many beliefs that are based on the justifying evidence that don't require such meta-beliefs. This is the case for children and unreflective adults who lack the ability to form such meta-beliefs. Another problem with doxastic accounts is that they fall prey to a vicious infinite regress because the meta-belief itself will have to be justified by another meta-belief and so on *ad infinitum*. On the other hand, if the first meta-belief doesn't need to be justified then it's unclear how such a belief could be useful for basing (McCain, p.85).

In contrast to doxastic accounts of the basing relation, McCain opts for an *interventionist* account of causation that he thinks provides a satisfactory way of dealing with the basing relation.

McCain's interventionist account of the basing relation is a type of *manipulability* theory.[58] Generally, manipulability theories hold that when A is a cause of B our manipulating A in certain ways changes B. More precisely:

> A causes B if and only if B would change if an appropriate manipulation on A were to be carried out (McCain, 2014, p.86; Woodward, 2008a).

Interventions are manipulations on one or more variables in a system. More precisely, "…an intervention *I* on *X* with respect to *Y* will be such that *I* causes a change in *X*, *I* does not cause a change in *Y* via some route that does not go through *X*, and *I* is exogenous in the sense of not itself having a cause that affects *Y* via a route that does not go through *X*" (McCain, p.87;

[58] McCain appeals to James Woodward (2003, 2008a, 2008b) and Brad Weslake (2011) to explain his interventionist account.

Woodward, 2008b, p.213). Now, it is also important to understand what it is for X to be causally relevant to Y.

> **(M)** *X* causes *Y* if and only if were some intervention…that changes the value of *X* to occur, *Y* or the probability distribution of *Y* would change in some regular, stable way [at] least in some range of background circumstances *B* (Ibid).

M holds that a cause is a difference maker. M says that a cause must make a difference to its effect because in order for *X* to be a cause of *Y*, changes in the value of *X* must be correlated with changes in *Y* in a stable way (McCain, p.87).

One important feature of the interventionist view is an extension of M that accounts for a specific kind of *type-causal relation*, the relation of a *direct cause*.

> A necessary and sufficient condition for *X* to be a direct cause of *Y* with respect to some variable set V is that there be a possible intervention on *X* that will change *Y* (or the probability distribution of *Y*) when all other variables in V besides *X* and *Y* are held fixed at some value by interventions (McCain, p.87; Woodward, 2003, p.55).

To help understand what a direct cause is, McCain offers the following example: "…if we have a variable set that includes *X*, *Y*, and *Z*; *X* is a direct cause of *Y* if and only if there is a possible intervention on the value of *X* that will change the value of *Y* when we hold fixed the value of *Z"* (McCain, p.87). McCain also explains how we should determine which variable set to use when evaluating whether *X* is a direct cause of *Y*. We are concerned with when S's belief is based on her evidence, so the relevant variable set will consist of variables that can propositionally justify or defeat S's justification for believing a proposition. Hence, the pertinent variable set will be a set of non-factive mental states (McCain, p.88).

A direct cause is also different in an important respect from an *actual cause*. In order to understand what an actual cause is we need to first understand two other important notions. First, there is the notion of a *directed path*. A directed path from one variable *X* to another variable *Y* is a chain of direct causal connections from *X* to *Y* (McCain, p.88; Woodward, 2003,

p.59). That is, there is a directed path from *X* to *Y* if and only if each variable starting with *X* and ending with *Y* is a direct cause of the variable that immediately succeeds it. So, if there is a causal chain from *X* to *R* to *S* to *Y*, this chain is a directed path from *X* to *Y* if and only if *X* is a direct cause of *R*, *R* is a direct cause of *S*, and *S* is a direct cause of *Y*. The second important notion for understanding what an actual cause is, is the notion of a *redundancy range*.

> For a path P from X to Y in a causal model, define V1…Vn as all variables that are not on P. Values v1…vn are on the redundancy range for Vi with respect to P if no intervention on v1…vn while holding X fixed would result in a change to the actual value of Y (McCain, p.88; Weslake, 2011, p.8).

Together, the notions of a directed path and redundancy range help us understand an *actual* cause. When *X* is an *actual* cause of *Y*, *X*'s having its actual value is a cause of *Y*'s having its actual value. Thus, 'actual cause' refers to a relation of *token* causation. Basically, when *X* is an actual cause of *Y* it is possible to change the actual value of *Y* by changing the actual value of *X* while still keeping all other direct causes of *Y* that are not part of a chain of direct causes (directed path) leading from *X* to *Y* fixed at a value within their redundancy range (McCain, p.89; Woodward, 2003, p.84).[59]

Given the above features of the interventionist account of causation, McCain believes that he has the means to formulate the basing relation which "…bridges the gap between propositional justification and well-founded belief (doxastic justification)" (p.91). That is, for S's belief to be well-founded, S's belief must be propositionally justified by her evidence and she must base her belief on that justifying evidence. Here is McCain's interventionist account of the general basing relation:

***IB*:** S's belief that *p* at *t* is based on *X*, if and only if at *t*:

[59] For a nice (lengthy) example to help understand the various features of the interventionist account see McCain, p.89-90.

1) *X* is a direct cause of S's believing that *p*

AND

2) *X* is an actual cause of S's believing that *p* (McCain, p.93).

Now, we need to know whether S's evidence bears a causal relation of sufficient strength to her belief in order for that belief to be based on the evidence. Thus, McCain offers a refined version of the general relation:

***IB-R*:** S's belief that *p* at *t* is based on her evidence, *E*, if and only if at *t*:

1) Each $e_i \in E$ is a direct cause of S's believing that *p*

AND

2) Each $e_i \in E$ is an actual cause of S's believing that *p*

AND

3) It is not the case that intervening to set the values of all direct causes of S's believing that *p*, other than the members of *E*, to 0 will result in S's not believing that *p* when every $e_i \in E$ is held fixed at its actual value (Ibid).

In IB-R, *E* is a subset of S's non-factive mental states, which propositionally justifies *p*. S's belief that *p* has the connection to her propositional justification needed for well-founded belief only if there is some *E* such that S's belief that *p* is based on *E* in the manner described in IB-R and S's believing that *p* satisfies the conditions laid out in EF-EJ for *p* to be justified for S. Also, sets of variables are used when evaluating the causal relations between S's evidence and her beliefs. These variable sets include variables for S's belief, S's evidence (*E*), and any factor that might be a causal influence on S's believing as she does (e.g., wishful thinking, being struck by a ray from Alpha Centauri, etc.). Further, assigning values for the variables mentioned involves imposing a binary system in which each variable has a value of either 1 or 0. So, for any variable, the influence of that variable is either active or it is not. However, we are not limited to a binary schema. We might adopt a schema in which these variables get assigned

values based on the strength of the mental state, mechanism, or influence in question. Finally, IB-R describes a relation between a belief and its causal base at a given time *t*. Thus, IB-R offers an account of basing for both belief formation and sustention. This is crucial since the causal basis of a belief might fluctuate over time so that at *t*, *x* causes S to form the belief that *p*, but at t_2 S's belief that *p* is causally sustained by *y*. Thus, depending on the nature of *x* and *y* it is possible that S's belief that *p* is based on her evidence at one time, but not another. This is in line with the intuition that someone may form a belief that is not based on sufficiently strong evidence, but later on obtain evidence that sufficiently supports the belief and continue to hold the belief because of that evidence, so the later belief is based on sufficient evidence (McCain, p.91-93).

McCain provides a number of examples to show how IB-R works in application (see pp.93-99) and how IB-R handles standard objections to causal accounts of the basing relation (see pp.99-112). I would like to look at a couple of these examples, both of which arise from a particular objection to IB-R that is relevant to my overall project. This particular objection to IB-R involves a challenge by John Turri (2010) in which he challenges the 'orthodox view' which holds that well-founded belief should be understood in terms of propositional justification and not the other way around. Turri provides counter-examples which he thinks reveal that the orthodox view "…misses something *deep and important* about the relationship between propositional and doxastic justification" (McCain, p.109; Turri, p.317-18). More precisely, "The way in which the subject performs, the manner in which she makes use of her reasons, fundamentally determines whether her belief is doxastically justified. Poor utilization of even the best reasons for believing *p* will prevent you from justifiedly believing or knowing that *p*" (McCain, p.110; Turri, p.318). As an analogy, Turri offers the following example:

> In evaluating beliefs we are evaluating a kind of performance, the performance of a cognitive agent in representing the world as being a certain way, and when performing with materials (which, in cognitive affairs, will include reasons or evidence), the success, or lack thereof, of one's performance will depend crucially on the way in which one makes use of those materials. This is true for carpentry as well as cognition. Consider a carpenter, equipped with the finest tools and lumber. You want a deck built, and he is in a position to build a wonderful one for you. Despite the quality of his tools and lumber, unless he puts them together *in the right way*, you are not going to be happy with the end result. You want a *well built deck*, not just one built with tools and materials *fit for making a well built deck.* Merely having the right equipment for the job, and using it to perform the job, does not guarantee a job well done (Turri, p.315).

IB-R is part of the orthodoxy, so if Turri is right then there is something very wrong with it. However, McCain believes that IB-R is not threatened by Turri's objection.

Here is Turri's first counter-example. There are two jurors, Miss Proper and Miss Improper, who sit in judgment of Mr. Mansour. Both know a set of propositions (P1-P4) which propositionally justifies believing that <Mansour is guilty>. Here is how the case is described:

> **PROPER & IMPROPER**
>
> As it happens, each comes to believe <Mansour is guilty> as the result of an episode of explicit, conscious reasoning that features (P1–P4) essentially. Miss Proper reasons like so: (Proper Reasoning) (P1–P4) make it overwhelmingly likely that Mansour is guilty. (P1–P4) are true. Therefore, Mansour is guilty. Miss Improper, by contrast, reasons like this: (Improper Reasoning) The tea leaves say that (P1–P4) make it overwhelmingly likely that Mansour is guilty. (P1–P4) are true. Therefore, Mansour is guilty (McCain, p.109; Turri, p.316).

Turri says that this is a counter-example to orthodox views since (P1-P4) propositionally justify <Mansour is guilty> for both jurors and Miss Improper bases her belief on (P1-P4) *as evidence* for <Mansour is guilty>. It is part of the content of her reasoning that (P1-P4) make <Mansour is guilty> extremely probable. However, Miss Proper's belief is well-founded and Miss Improper's belief isn't (McCain, p.110; Turri, p.317).

Turri's second case is as follows:

PONENS & LACY

Mr. Ponens and Mr. F.A. Lacy each knows the following things:

(P5) The Spurs will win if they play the Pistons.

(P6) The Spurs will play the Pistons.

This is a paradigm case of propositional justification. <The Spurs will win> is propositionally justified for each man because he knows (P5) and (P6). From these two premises, and only these premises, each man draws the conclusion:

(P7) Therefore, the Spurs will win.

Ponens applies modus ponens to reach the conclusion. Lacy, however, applies a different inference rule, which we may call *modus profusus*: for any p, q, and r: $(p \wedge q) \rightarrow r$. Lacy's belief that the Spurs will win is definitely not doxastically justified; following that rule could never lead to a justified belief (McCain, p.110; Turri, p.317).

Turri thinks this is a problem case for orthodox views since Ponens applies a proper inference rule to reach the conclusion while Lacy used an invalid inference rule. Ponens belief is non-deviantly caused by his belief in P5 and P6, and the same is true for Lacy. Further, both Ponens and Lacy believe that the conjunction of P5 and P6 is a good reason to believe P7. However, Ponen's belief P7 is well-founded, but Lacy's belief isn't (Ibid).

McCain agrees with Turri that the manner in which one uses her evidence makes a crucial difference to whether her beliefs are based on her evidence in the sense required for a belief to be well-founded. However, McCain thinks that Turri is mistaken in thinking that *all* orthodox views of the basing relation fail to account for this truth. McCain claims that IB-R accounts for this truth.

Here is McCain's explanation of how IB-R yields the intuitive result that Miss Proper's belief that <Mansour is guilty> is based on her evidence while Miss Improper's isn't. Beginning with Miss Proper, the relevant variable set for her contains: p (Miss Proper's belief that <(P1-P4) make it overwhelmingly likely that Mansour is guilty>), q (Miss Proper's belief that <(P1-P4)

are true>, *d* (Miss Proper's disposition to have the seeming that <Mansour is guilty> is a logical consequence of the best explanation of her evidence), and *B* (Miss Proper's belief that <Mansour is guilty>). The key is determining whether *B* is based on *E*. The first step is to determine if each of the members of *E* (*p*, *q*, and *d*) are a direct cause of *B*. This seems to be the case since there are interventions that one could perform on each of them that would result in a change in *B*. For example, if we hold fixed *q* and *d* but lower the value of *p* to 0 via an intervention, Miss Proper wouldn't have a reason to infer *B* (B's value would change) since she is making an inference from <(P1-P4) make it overwhelmingly likely that Mansour is guilty> and <(P1-P4) are true> to <Mansour is guilty>. The same applies *mutatis mutandis* for *q* and *d*. Furthermore, each member of *E* is an actual cause of *B*. There are interventions on the actual values of the members of *E* that would each result in the actual value of *B* changing. For example, holding the values of the other variables fixed at a value within their redundancy ranges, but lowering the actual value of *p* to 0. With *p* set at 0, *B* will change to 0. Intuitively, if Miss Proper doesn't believe that <(P1-P4) make it overwhelmingly likely that Mansour is guilty> she will no longer infer that < Mansour is guilty>. The same applies *mutatis mutandis* for *q* and *d*. Therefore, IB-R yields the correct result that Miss Proper's belief *B* is based on her evidence (McCain, p.111).

McCain also argues that IB-R successfully rules out Miss Improper's belief that <Mansour is guilty> from counting as based on her propositionally justified evidence. According to Turri, the evidence that propositionally justifies Miss Improper in believing <Mansour is guilty> consists of her justified belief that <(P1-P4) are true>. Hence, *E* consists of *q* (<(P1-P4) are true>) and *d* (Miss Improper's disposition to have the seeming that <Mansour is guilty> is a logical consequence of the best explanation of her evidence). The other relevant variables are *t* (Miss Improper's belief that <the tea leaves say that (P1-P4) make it

overwhelmingly likely that Mansour is guilty>) and *B* (Miss Improper's belief that <Mansour is guilty>). McCain points out that in Miss Improper's case, the third condition of IB-R isn't met. Interventions that set the values of all direct causes of Miss Improper's believing that <Mansour is guilty>, other than the members of *E*, to 0 results in her not believing that <Mansour is guilty> when each member of *E* is held fixed, but sets *t* to 0 will result in *B* being 0. Now, because part of what leads Miss Improper to believe that <Mansour is guilty> is her belief that <the tea leaves say that (P1-P4) make it overwhelmingly likely that Mansour is guilty>, it seems plausible that removing this belief would lead to her not believing that <Mansour is guilty>. This would be the case even if she continued to believe that <(P1-P4) are true> and kept her disposition to have the seeming that <Mansour is guilty> is a consequence of the best explanation of her evidence. IB-R yields the result that Miss Improper's belief is not based on her justifying evidence. Therefore, IB-R yields the intuitively correct result that Miss Proper's belief is based on her justifying evidence while Miss Improper's is not. Hence, PROPER & IMPROPER fails as a counter-example to IB-R (McCain, p.111-12).

McCain also argues that PONENS & LACY poses no problem for IB-R either. Very briefly, Ponen's belief in the conclusion of his inference is based on his justifying evidence. However, Lacy does not base his belief on his justifying evidence (McCain, p.112).

Therefore, McCain argues, IB-R yields the intuitive result in both of Turri's counter-examples because IB-R emphasizes that the manner or way in which S makes use of her reasons, fundamentally determines whether her belief is well-founded. McCain concludes "If Turri is correct that all the other orthodox views of the basing relation fail to capture this important truth, then instead of an argument against the orthodoxy, Turri has helped provide a strong argument in support of IB-R as the correct account of the basing relation" (Ibid).

McCain believes that EF-EJ is a satisfactory account of propositional justification and that IB-R is a successful account of the basing relation. Thus, putting EF-EJ and IB-R together gives us a complete evidentialist theory of epistemic justification that McCain calls *Explanationist Evidentialism*. Explanationist evidentialism is a full and fleshed out account of Conee and Feldman's schema for an evidentialist account of well-founded belief.

EX-WF

At *t*, S's belief that *p* is well-founded if and only if:

At *t*,

(I) 1) Each $e_i \in E$ is a direct cause of S's believing that *p*
AND
2) Each $e_i \in E$ is an actual cause of S's believing that *p*
AND
3) It is not the case that intervening to set the values of all direct causes of S's believing that *p*, other than the members of *E*, to 0 will result in S's not believing that *p* when every $e_i \in E$ is held fixed at its actual value.

(II) *E* is a subset of S's occurrent non-factive mental states and the non-factive mental states that she is disposed to bring to mind when reflecting on the question of *p*'s truth.

(III) 1) *p* is part of the best explanation available to S for why S has *E*
OR
2) *p* is available to S as a logical consequence of the best explanation available to S for why S has *E*.

(IV) There is no set of S's occurrent non-factive mental states and the non-factive mental states that she is disposed to bring to mind when reflecting on the question of *p*'s truth, *E**, such that:
A) *E* is a subset of *E**
AND
B) *p* is not part of the best explanation available to S for why S has *E**
AND
C) *p* is not available to S as a logical consequence of the best explanation available to S for why S has *E** (McCain, p.118).

McCain explains the various conditions of EX-WF. Condition (I) holds that S's belief must be based on her justifying evidence in the manner required by IB-R in order to be well-founded. Condition (II) holds that the evidence that S's belief is based on at *t* must be evidence that she

possesses at *t*.[60] Condition (III) says that S has to satisfy the requirements for propositional justification (EF-EJ). Condition (IV) is a crucial 'no defeater' condition, which says that in order for S's belief that is based on *E* to be well-founded, it cannot be fully defeated by other evidence that S possesses. Together, these conditions yield a full account of well-founded belief and therefore a complete evidentialist theory of epistemic justification (McCain, p.119).

IV.3 - Objection to EX-WF

In this section I wish to begin a preliminary sketch of an argument that concludes that EX-WF is *not* a full account of epistemic justification because there is an important sense in which it fails to account for the prominent role that reliability often plays in epistemic justification. In order to see this it will be instructive to see how reliability *cannot* figure into EX-WF.

One objection that deserves consideration is that in the Turri cases that McCain discusses, the fact that Miss Improper and Mr. Lacy aren't justified is because they don't *properly* base their belief on the evidence. In order to properly base one's belief on the evidence one must use it the right way, e.g., by applying *modus ponens* correctly, etc. But by saying this, isn't McCain giving away the game to the reliabilist because 'using' one's evidence properly involves using conditionally reliable belief-forming processes such as good reasoning? That is, it would seem that McCain is ignoring an important *reliability component* in his account that plays a crucial role in the justificatory status of the relevant beliefs. The belief had to be *properly formed by a reliable belief-forming process*, i.e.—the process of utilizing one's evidence properly. If so, then there's an unacknowledged reliability-related element in McCain's account of well-founded belief.

[60] For an account of what it means to *possess* evidence, see McCain (2014) Chapter 3.

However, I think McCain has an effective response to this objection. Although McCain agrees with Turri that the manner in which S makes use of her reasons does matter for basing, there's no need to concede to the reliabilist with IB-R. The problem with Miss Improper and Lacey is that a *non-justifying reason* (the bad inference rule) is playing too strong of a causal role in their beliefs. Put another way, the problem is that the *evidence* which propositionally justifies Miss Improper and Lacey is not playing a sufficiently strong causal role in their forming the beliefs in question.[61] Reliability of a process doesn't play a role here.[62]

Although EX-WF needn't concede anything to the reliabilist with respect to IB-R, I still think there is an important sense in which EX-WF cannot account for the role that reliable evidence gathering plays in a belief's justificatory status.

McCain responds to objections that concern cases of bad evidence gathering. In such cases, S has good evidence for believing that *p,* but she has been negligent in her evidence gathering. If S had not been negligent then she would have gathered evidence that does not on balance support believing *p*. S's negligence is supposed to make it so that she lacks justification. EX-WF yields the result that S has justification for believing that *p* because all that matters is the evidence that S has at the time. This is exemplified in the GEORGE case. Remember that Baehr stipulates that George's belief about second-hand smoke is well-supported by the evidence, but that he only has this evidence as the result of his intellectual 'tunnel vision'. Baehr thinks that George is clearly unjustified because of this. Thus, Baehr thinks that views such as EX-WF need a constraint that one must be virtuous in her evidence gathering practices in order to have

[61] A similar point can be made about Engel's example of 'Sally the misinformed logic student' from Chapter Three. In this case Sally doesn't base here belief on the evidence. She is *propositionally* justified, but her belief isn't well-founded because her evidence is not playing a sufficiently strong causal role in forming her belief. Instead, her mistaken belief about *modus ponens* is playing too strong a role in causing her belief.

[62] I would like to thank Kevin McCain for providing helpful clarification on this point (in communication).

justification. However, McCain diagnoses this case differently. McCain thinks that despite George's obvious intellectual viciousness, George does not lack epistemic justification for his belief because his belief is supported by the evidence he has. Just as Conee and Feldman argue, McCain thinks that even if we acknowledge that George should gather more evidence, it does not follow that he should not believe what he does *right now*, given the evidence he currently possesses (McCain, p.150-51). "After all, if someone else were to have exactly the same evidence as George, we would be inclined to think that she should believe that secondhand smoke does not pose a significant health risk just as George does" (McCain, p.151).

Even worse than GEORGE, there are cases in which a person intentionally and actively tries to avoid the truth. Consider the following case as an example:

> **FOOTBALL JON**
>
> Jon hates to find out the results of a football game before he has had a chance to watch the game in its entirety. He missed yesterday's game, so Jon tries very hard to avoid learning the score. He refuses to consult any of the typical sources for this information. He will not read newspapers, watch the news, or even listen to the radio. In fact, he goes so far as to avoid talking to people he knows are likely to have watched the game. While running out of a room with his ears plugged to avoid hearing the latest sports report Jon runs into a newspaper rack and sees on the front page of a newspaper he knows to be extremely reliable that the Bears won the game yesterday (McCain, p.151).

McCain says that Jon has very good evidence for believing that the Bears won the game and that he obviously has justification for believing the results of the game. At the same time, it is clear that Jon is a poor epistemic agent with respect to this proposition. Jon is clearly worse than George in this respect. Regardless, Jon has justification for believing that the Bears won. Jon's evidence gathering methods do not change this fact. McCain writes, "While we might aptly criticize agents for not employing good evidence gathering methods, the fact that their evidence gathering methods are bad does not change what their evidence supports or what they have

justification to believe" (Ibid). McCain concludes that cases of bad evidence gathering pose no problem for EX-WF.

One possible difficulty with McCain's argument in this section of the book is that he holds seemingly inconsistent views. As we just saw, he writes (1) "While we might aptly criticize agents for not employing good evidence gathering methods, the fact that their evidence gathering methods are bad does not change what their evidence supports or what they have justification to believe". However, at the end of the previous paragraph in that section of his book, he writes (2) "Further, S's justification does not depend upon how she gathers that evidence *so long as she has no reason to think that her evidence gathering is flawed in some relevant way* [My italics]" (Ibid). If, as McCain claims, (1) is true, then having a reason to believe your evidence-gathering is flawed needn't have any impact at all on what belief your evidence justifies (~2). But this is a contradiction (at least it seems to be given that McCain doesn't elaborate on (2)). Hence, both (1) and (2) can't both be true.

If what was said in the last paragraph is correct, then it would seem that when evidentialists like McCain concede that *awareness* of bad evidence-gathering (unreliability) makes a justificatory difference, they unwittingly commit themselves to the view that *reliability itself* can make a justificatory difference. So, if it really is an epistemological truth that bad (unreliable) evidence gathering doesn't make a justificatory difference, then a self-consciously bad (unreliable) evidence gatherer should be able to apply that epistemological truth to his own situation and remain just as justified in the beliefs he bases on the badly (unreliably) gathered evidence as he was *before* he realized he was gathering evidence badly (unreliably). However, I think that McCain has the resources to explain this and this is one of the main issues I will discuss in detail in the next chapter. But, as I will argue in the next chapter, there is another

important reason to think that reliability still figures prominently in a belief's justificatory status, i.e.—other things being equal, if your evidence was gathered using a reliable belief-forming process (e.g., virtuous inquiry), then the probability that your evidence is not misleading increases and hence the degree to which your belief (which is based on that evidence) is justified is increased. The issue is that McCain's view doesn't account for the fact that a belief can be well-founded and still be justified to a higher or lesser degree depending on whether the person's evidence was gathered using reliable or unreliable evidence-gathering methods (virtuous or vicious inquiry).

Now, I emphasize that I'm speaking here of evidentialists *like* McCain. Other prominent evidentialists like Conee and Feldman do not think that awareness of bad evidence-gathering has any bearing on one's evidential support and makes no justificatory difference. For example, Feldman writes: "It would be a mistake to hold that the mere fact that I know that there is additional evidence…neutralizes the evidence I already have" (Feldman, 2004b, p.235).[63] Even though Conee and Feldman don't make the same claim that McCain makes about the justificatory difference that the awareness of bad evidence-gathering makes (and therefore don't contradict themselves), I will argue that they are still mistaken in claiming that such awareness makes no justificatory difference. I take this up in the next chapter as well.

V. Conclusion.,

In Chapter Four I motivated, presented, and provided a defense of what I think is the most promising and complete version of evidentialism—explanationist evidentialism. I

[63] Feldman is referring to the "Precarious Peak" example in which a hiker has evidence of the difficulty of a particular hike. In addition to this evidence the hiker has a book with facts about the relative difficulty of hikes (including the one the hiker is considering). However, the hiker fails to consult the hiking book. Feldman says that the evidence in the book is not psychologically available in the relevant sense. It is only available in the sense that the hiker could have obtained it.

discussed alternative accounts of the epistemic support relation and provided reasons for why each of these alternative accounts are unsuccessful. I also discussed how explanationist evidentialism handles cases in which it seems that flawed evidence-gathering or irresponsible inquiry are responsible for a belief. I then hinted at the fact that the explanationist response does not completely rule out the idea that reliability or unreliability can still play an important role in determining the degree to which a belief is justified.

In Chapter Five I will defend the explanationist view that flawed evidence gathering, vicious inquiry, etc., is problematic for a belief's justification if the person has evidence that she engaged in such unreliable belief-forming methods. That is, having evidence that one engaged in unreliable evidence-gathering methods can *decrease* the degree to which a belief is justified for a person and in some cases give a person reason to disbelieve some proposition. Then I will finally provide a more thorough presentation and defense of my two-component view of justification. I will argue that the reliability or unreliability of one's evidence gathering methods can increase or decrease the degree to which one's well-founded belief is justified (other things being equal) even if one has no evidence of this reliability or unreliability.

CHAPTER FIVE

Towards a Two-Component Theory of Epistemic Justification

In this chapter, I will finally put together my two-component theory of epistemic justification. I will argue that EX-WF gives us a complete synchronic theory of *categorical* epistemic justification, but that including reliability as an additional justificatory factor provides a more satisfactory account of justification. That is, S's belief that *p* is *categorically* justified (justified 'period' or 'full stop') when *p* meets the conditions of EX-WF, but that in addition, the reliability of one's evidence-gathering methods can also play a salient role in increasing the degree to which *p* is justified (other things being equal). Thus, I will argue that a belief can be justified along two-dimensions: the *fittingness-dimension* and the *reliability-dimension*, but that being justified along the fittingness-dimension is all that is necessary for a belief to be justified full-stop (this also allows for the strength of one's evidence to partly—sometimes wholly—determine the degree to which a belief is justified). Being justified along the reliability-dimension can only increase the degree to which the belief is justified.

First I will provide an account of how the evidentialist can make sense of the role that evidence-gathering plays in fixing a belief's justificatory status without appealing to notions of reliability. Then I will argue that there is an additional important externalist justificatory factor that is positively relevant to a belief's justificatory status, namely the reliability of one's method of inquiry (virtuous evidence-gathering methods). Specifically, I will argue that if the method of inquiry used to gather evidence is highly reliable, then the probability that that evidence is not misleading increases. If so, then that reliability has a *positive* epistemic relevancy to the justificatory status of the pertinent belief. The flipside is also true. That is, if the method of inquiry used to gather evidence is highly unreliable, then the probability that that evidence is

misleading increases. If so, then that reliability has a *negative* epistemic relevancy to the justificatory status of the pertinent belief.

I. Evidence *of* Flawed Evidence Gathering (EFEG) as an 'Undercutting Defeater'

I will now discuss how the explanationist evidentialist can make sense of how flawed evidence gathering methods can undermine the justification of a belief. In order for this to happen, a person must have *evidence* that her evidence gathering method was flawed (unreliable).

Epistemic justification is defeasible. That is, justified beliefs can be refuted or undermined by other beliefs, evidence, etc. Michael Bergmann (2005) distinguishes between what he calls *propositional defeaters* (which are propositions) and *mental defeaters* (which are either propositional attitudes or experiences or combinations thereof) (p.422). I will not discuss propositional defeaters. Instead, my focus will be on mental state defeaters. Bergmann defines mental state defeaters as follows:

> **D_1.** *d* is a defeater at *t* for S's belief *b* iff (i) *d* is an experience or propositional attitude or combination thereof; (ii) S comes to have *d* at *t*; (iii) as a result of S's coming to have *d* at *t*, *b* ceases to be justified.

Epistemologists typically distinguish between two ways that a belief can be defeated. First, there are defeaters that are reasons for believing that one's belief that *p* is *false*. Such defeaters give us reason to believe ~*p* (the negation of *p*). Second, there are defeaters that are reasons to lower the probability that one's belief that *p* is true. Such defeaters are reasons for no longer believing *p* or at least believing *p* to a lesser degree of confidence. The first type of defeaters are called *rebutting defeaters* and the second type are called *undermining* or *undercutting defeaters*.

John Pollock (1986) provides one widely used taxonomy of defeaters. A *rebutting defeater* for some belief *p* is a reason for holding the negation of *p* or holding some proposition *q* that is incompatible with *p* (Pollock, p.38). For example, from a distance Mary sees what seems to be a sheep in a field, but the owner of the farm tells her that there is no sheep in the field. This provides Mary with a reason to believe that there is no sheep in the field. Alternatively, Mary might walk up to the sheep in the field to discover it is a fake sheep made out of papier- mâché. In which case she has acquired a reason to believe *q* (the sheep is made out of papier-mâché), which is incompatible with *p* (there is a *real* sheep in the field). Bergmann provides a nice formal definition of a rebutting defeater that is suitable for our purposes:

> **D2**. *d* is a rebutting defeater for *b* iff *d* is a defeater for *b* which is (or is an epistemically appropriate basis for) the belief that *b* is *false* (Bergmann, p.424).

In contrast to a rebutting defeater, an *undercutting* defeater is a reason for no longer believing *p*. It is a reason to think that one's *ground* for believing *p* is not sufficiently indicative of the truth of *p* (Pollock, p.39). For example, suppose a person enters a widget factory where there is an assembly line on which there are widgets that appear red. It thus appears to the person that he sees red widgets. However, the factory manager explains to the person that the widgets are actually being illuminated by powerful red lights. In this case the person loses his reason for believing that the widgets are red instead of acquiring a reason for believing that they are not red.[64] Again, Bergmann provides a suitable definition of an undercutting defeater:

> **D3**..*d* is an undercutting defeater for *b* iff *d* is a defeater for *b* which is (or is an epistemically appropriate basis for) the belief that one's actual ground or reason for *b* is not indicative of *b*'s truth (Ibid).[65]

[64] Importantly, there are two further distinctions to be made. There are also *partial* and *full* defeaters. Full defeaters completely neutralize one's justification for believing *p*, while partial defeaters simply reduce one's justification for believing *p* to some degree (Bergmann, p.422).

[65] In a revised definition of D_3, Bergmann adds the following clause: "(ii) is (or is an epistemically appropriate basis for) an attitude of significant uncertainty about the proposition that one's actual ground or reason for b is indicative of b's truth, significant enough to withhold that proposition" (p.427). This added clause is not important to my present discussion. For Bergmann's argument supporting this clause see Bergmann (2006) pp.424-27.

So what kind of defeater is EFEG (evidence of flawed evidence-gathering)? EFEG is a reason to question the quality of one's evidence in support of *p* such that one is no longer justified in believing *p* or at least less justified to some degree. EFEG does not give one a reason to believe the negation of *p* or to believe some proposition *q*, which is incompatible with *p*. For example, suppose a private investigator collected evidence for a client that he thinks strongly supports the belief that his client's wife is having an affair. Based on the evidence the investigator gathered, the client believes that his wife is having an affair. However, after discussing the evidence with the investigator, the client learns that the investigator was sloppy in gathering the evidence. Although the client still has evidence that his wife is having an affair, he now has a reason to be less confident in believing the evidence is veridical. Hence, he has some justification for believing that his wife is having an affair, but the degree to which he is justified has been diminished considerably (possibly enough to suspend judgment on the matter). The client does *not* have reason to believe that his wife *isn't* having an affair (the negation of what he currently believes), nor does he have a reason to believe something incompatible with his belief, e.g., that his wife is simply having business-related encounters with another man. Thus, EFEG is an undercutting defeater rather than a rebutting defeater.

I.1 - 'Higher-Order' Evidence and Defeaters

Next, we need to understand what *type* of evidence EFEG is. One way to get at this is to distinguish between *first-order* and *second-order* evidence. First-order evidence is evidence for a "non-evidential" proposition, i.e.—a proposition that isn't itself about evidence. Second-order evidence, or what has been called *higher-order* evidence (HOE) (Christensen, 2010; Feldman, 2005; Kelly, 2005; 2010; Lasonen-Aarnio, 2014), is evidence for an "evidential" proposition, i.e.—a proposition that is itself about evidence. EFEG is higher-order evidence or HOE. HOE

often includes the following type of evidence: that one is the victim of some sort of cognitive malfunction, that one has made a crucial mistake in one's calculations or in properly cataloguing one's evidence, that one has failed to give some evidence its due, that some epistemic principle has been incorrectly used, that one has engaged in poor evidence gathering methods, etc. HOE are often defeaters that lead one to doubt one's epistemic rationality.[66] [67] For example,[68] suppose that a meteorologist analyzes available evidence that bears directly on tomorrow's weather and as a result concludes that it will rain tomorrow. The meteorological data is first-order evidence that bears directly on the meteorologist's conclusion that it will rain tomorrow. The fact that the meteorologist arrived at his conclusion on the basis of the available evidence is HOE because it is evidence about the content and import of the meteorological data (as evidence). However, the meteorologist might have some HOE that the cognitive process that produced the belief was malfunctioning, e.g., perhaps the meteorologist realizes that he was very drunk when he analyzed the data or perhaps he has strong evidence that he performed some miscalculations (maybe he remembers that he applied an incorrect formula). Such evidence is evidence that his conclusion about the weather might be seriously flawed to the extent that he is no longer justified or at least far less justified in holding that conclusion.

HOE is 'backward-looking' in the sense that it is evidence that concerns whether someone's belief was *ever* epistemically rational to begin with. It's not that the meteorologist fails to be justified once he had this HOE. Instead, such evidence indicates that the meteorologist was *never* justified in holding the relevant belief because at the time he formed the

[66] This is the majority view but there are some that disagree, e.g., Field (2000).
[67] Importantly, not all HOE are defeaters. For example, often times we get new evidence which indicates that some previously acquired evidence was *properly* gathered. However, I'm currently focusing my discussion on 'flawed' evidence-gathering.
[68] Kelly discusses a similar example. See Kelly (2014).

belief he was extremely drunk or performed some significant miscalculation. When we get HOE, it serves as evidence that the evidence we gathered is very likely misleading.

The above discussion doesn't completely exhaust the types of HOE. There are two more I will briefly mention. First, there is HOE which involves someone else's belief that some proposition is true. For example, suppose that Rick and Jeff have the same first-order evidence that bears on some proposition *p*. Suppose that Rick forms his belief *p* before Jeff forms the same belief, and Rick tells Jeff what his conclusion is. One question is whether Jeff should take Rick's conclusion to be additional evidence in support of *p*? Similarly, in cases of peer disagreement in which two experts in the same field with exactly the same evidence arrive at different conclusions, should the fact that another expert in your field disagrees with your evaluation of the evidence serve as evidence against your own conclusion? These are all important issues but they go beyond the scope of the present discussion. The takeaway from this section should be that HOE differs from first-order evidence in that the former either prevents us from giving some part of our first-order evidence its due or in some cases it might serve to increase our credence in some conclusion we've reached given the available evidence.[69] Further, HOE differs from first-order undercutting defeaters in that HOE does not disturb the connection between the evidence and the conclusion. Instead, the person is in a position where he can't trust his own appreciation of the first-order evidence. "Thus…rational accommodation of HOE can require a certain kind of bracketing of some of one's reasons, in a way that does not seem to occur in accommodating ordinary evidence, even when that evidence is an ordinary undercutting defeater" (Christensen, p.198).

[69] I will discuss this issue later on.

I.2 - Michael Bergmann—'Defeaters and Higher-Level Requirements'

In this section I will argue that EFEG is a higher-level belief that serves as an undercutting defeater for first-level (object-level) beliefs. That is, if one is aware that one's evidence gathering methodology was unreliable, then one has a defeater for the belief that the pertinent evidence bears directly on. However, in order for person S to have a justified belief, it is *not* a requirement that S satisfy the higher-level requirement of holding a belief about the reliability of one's evidence gathering methodology.

Michael Bergmann (2005) argues that a *higher-level requirement* is not required for justification. A higher-level requirement on justification is a requirement according to which a belief is justified only if the subject has a higher-level belief, i.e.—a belief about the epistemic credentials of a belief (I say more about this below). Bergmann argues that higher-level requirements have the following important asymmetry: analytically, a belief is not justified if we have a defeater for it, but contingently, it is often the case that to avoid having defeaters, our beliefs must satisfy a higher-level requirement (Bergmann, 2005, p.419). In other words, it is necessarily the case that if a belief has a defeater, then it fails to be justified (or I would add: at least justified to a lesser degree), but it is not necessary to meet a higher-level requirement for epistemic justification. A higher-level requirement would require that for a belief *p* that anyone holding *p* must have a certain *perspective* on it, i.e.—must conceive of *p* as having something going for it. For example, someone might think that in order for someone to justifiedly believe *p* he also believe that *p* was formed in a reliable way. A higher-level requirement for justification requires a higher-level belief about one's epistemic rationality. Bergmann rejects this higher-level requirement, but affirms the 'no defeater' requirement. I turn to this asymmetry argument now.

Bergmann argues that higher-level requirements are appealing because internalists often think that we ought to reflect critically on our beliefs, and therefore that

(*) *S takes up some doxastic attitude towards* p^*_s (the proposition that S's belief that *p* is formed in a reliable way[70])

is a necessary condition for the justification of S's belief that *p*. Similarly, internalists often think that if S disbelieves or suspends judgment on p^*_s, then S has a defeater for the belief that *p*. So, they conclude that S must believe p^*_s for belief *p* to be justified for S. Bergmann's point is that internalists who are drawn to higher-level requirements assume that if S's belief does not satisfy a higher-level requirement, S must either 1) question the reliability of his belief source or 2) be irresponsible by not thinking whatsoever about the reliability of that source (Bergmann, 2005, p.430).

Bergmann denies that * is a necessary condition for the justification of S's belief that *p* because there is no sense in which * contributes to the justification of *p*. How might * contribute to the justification of *p*? The line of thinking seems to be that by considering and taking some doxastic attitude towards p^*_s, S is thereby doing something that contributes to the justificatory status for his belief that *p*. On the other hand, if S disbelieves or suspends judgment on p^*_s, then doing so will not contribute to the justification of *p*. So, it is only the belief towards p^*_s that could possibly contribute to the justification of S's belief that *p*. Bergmann then asks whether just any belief, justified or not, does so (Bergmann, 2005, p.430-31). "Would an irrational, irresponsible or insane belief that p^*_s contribute to the justification of S's belief that *p*, any more than having no doxastic attitude at all towards p^*_S does?" (Bergmann, 2005, p.431). If we suppose that S's belief that *p* satisfies some externalist conditions and that S has no defeater for it

[70] Bergmann uses the example of a belief's reliability but any other relevant example of a higher-level belief would suffice to make his point.

and no doxastic attitude towards $p*_s$, then would the justification of S's belief that *p* be increased if S had an irrational, irresponsible, insane, or unjustified belief that $p*_s$? Bergmann thinks that the answer is certainly no. Thus, the only way in which * could possibly be required for the justification of S's belief that *p* is if we stipulate that it must be made true by S's having a justified belief that $p*_s$ (Ibid).

However, Bergmann argues that to require that * must be satisfied by S's having a justified belief that $p*_s$ brings up the following dilemma: 1) either this requirement applies to *all* beliefs, which leads to a vicious regress, or 2) it applies *only* to beliefs at the first/object level (and possibly some higher levels, but not all higher levels), then such a restriction is made *ad hoc* for no other reason than to avoid the regress. Bergmann concludes:

> I conclude, therefore, that although it is contingently the case that some beliefs are justified only if the person holding them has some higher-level belief, because it is contingently the case that some beliefs can avoid having a defeater only by this means, there is no general or principled higher-level requirement on justification (Ibid).

Thus, higher-level requirements on justification are *sometimes* required, but certainly not in all (if not most) cases. Given this conclusion we can put together Bergmann's asymmetry claim. First, recall that D_1 says:

> **D_1.** *d* is a defeater at *t* for S's belief *b* iff (i) *d* is an experience or propositional attitude or combination thereof; (ii) S comes to have *d* at *t*; (iii) as a result of S's coming to have *d* at *t*, *b* ceases to be justified.

On this account of a defeater it is *analytic* (i.e.—necessary and therefore uncontroversial) that a belief is justified only if there is no defeater for it (Bergmann, 2005, p.422). Secondly, as we just saw, higher-level requirements are not a requirement for all beliefs to avoid having defeaters. It is simply a *contingent* fact that *some* beliefs must satisfy a higher-level requirement to avoid defeaters.

Let me state precisely how this discussion relates to EFEG. EFEG is a higher-level belief that serves as an undercutting defeater for first-level (object-level) beliefs. That is, if one is aware that one's evidence gathering methodology was unreliable, then one has a defeater for the belief that the pertinent evidence bears directly on. On the other hand, in order for person S to have a justified belief, it is not a requirement that S satisfy the higher-level requirement of holding a belief about the reliability of one's evidence gathering methodology.[71]

Next, I turn to the issue of understanding in what sense being aware of the unreliability of one's evidence-gathering methods serves as an undercutting defeater.

I.3 - 'Defeat' In Explanatory Terms

On McCain's evidentialist model, we must think of 'defeat' in explanatory terms. Given some proposition *p* and some total body of evidence *E* that strongly supports *p*, the best explanation for why I have *E* is that *p* is true. For example, the proposition "It is snowing outside" is well-supported by the following evidence: *<There are white flakes falling from the sky>; <It is very cold outside>; <It is cloudy>; <The local meteorologist says it is snowing in parts of the city>*, etc. This evidence strongly supports the proposition "It is snowing outside". In this case, the best explanation of this evidence is that it happens to be snowing outside. However, suppose that I later find out that my roommate, who happens to work in the movie industry, played a trick on me and used a fake-snow making machine to create the appearance that it was snowing. My friend also played a fake weather forecast on my television. In this case, my *total body of evidence* has changed to include the fact that *<My roommate played a trick on me and used a fake-snow making machine>* and the fact that *<My friend tricked me with a fake forecast video>*. I now have evidence *E** and my new evidence of my roommate's

[71] However, I do think that *contingently, awareness* that one's evidence gathering methods were highly reliable does contribute to the justification of one's relevant belief (even if not necessary).

trickery serves as a defeater for my belief that the proposition "It is snowing outside" is true. Once I have this information about my roommate, *p* is no longer the best explanation of my total body of evidence. The best explanation of E^* is p_1: "My roommate tricked me into thinking it is snowing outside".

The above example is an instance of an *undercutting* defeater. Remember that a *rebutting* defeater gives one reason to believe ~*p*. For example, suppose I found that my roommate had tricked me and then I got on the internet to look up the local weather and found that all local news stations are reporting that it is *not* snowing. In this case, your roommate's admission of trickery gives you a strong reason to question whether *p* is true (an undercutting defeater), but not that ~*p* is true ("It is *not* snowing"). However, the information you got from the internet *does* give you good reason to believe that a rival ~*p* explanation is true (a rebutting defeater), i.e.—you get evidence of ~*p*. Now, higher-order evidence such as EFEG serves as an undercutting defeater, not a rebutting defeater. To see why, suppose I gather some body of evidence in order to find out who murdered the butler. I now have some body of evidence *E* that supports *p* (that the maid murdered the butler). In this case, let's suppose that the best explanation of *E* is that *p* is true. However, suppose further that I later learn that my evidence gathering methodology was biased in some significant way. Without this information, part of the best explanation for why I have *E* is that *p* is true. However, once I have this information concerning the biased nature of my evidence gathering methods, the truth of *p* is no longer clearly the best explanation. That is, I now have a rival explanation—I cherry-picked evidence so that I got the sort of evidence that I wanted. In this sort of case, I don't get evidence for thinking that a rival ~*p* explanation is true (say, that the maid was not the murderer or that the

gardener was the murderer instead), but I do get evidence that undercuts *E's* support for *p* by making the truth of *p* a weaker explanation of *E*.

The fact that EFEG is an undercutting defeater and that explanationist accounts of justification view defeaters in explanatory terms, means an explanationist evidentialist like McCain has the resources to respond to the criticism I made in the last chapter. Namely, the criticism that when evidentialists like McCain concede that *awareness* of bad evidence-gathering (unreliability) makes a justificatory difference, they unwittingly commit themselves to the view that *reliability itself* can make a justificatory difference. First, *total evidence* is always what matters. Once you have a reason to believe that your evidence gathering is flawed, you have *different* total evidence than what you had before. Suppose that you have evidence *E* which supports *p* <*the maid murdered the butler*>. However, unbeknownst to you, you gathered *E* in a flawed manner—perhaps you only consulted biased sources. In this case you hold that *E* supports *p*. And, assuming that you don't have any defeaters in your total evidence, you are justified in believing *p*. Now, add to the case that you come to have good reason to think that you gathered *E* in this flawed way. *E* still supports *p* (evidential support relations are necessary). However, your *total evidence* doesn't support believing *p*. The reason for this is that your total evidence now includes *E* as well as new information about your flawed inquiry (this information constitutes a reason to doubt that *p*'s truth is the best explanation of your having *E*, and thus prevents *E* from justifying *p* for you).

Second, the evidentialist could say that the *reliability* of one's evidence gathering methods itself (or its lack) doesn't make a justificatory difference. However, *awareness of reliability* or better yet 'evidence gathering methods that are apt to produce biased samples' does make a difference. How is this so? It's because *evidence* is what matters for justification. The

mere fact that something is unreliable doesn't matter—when it comes to defeaters, the mere fact that there is information in the world that would give me a defeater doesn't matter if I'm not aware of that information. Otherwise, any and all unpossessed evidence or facts about the world would serve as defeaters regardless of the fact that we are unaware of them. However, being aware that some process is unreliable does give me a defeater because it gives me evidence that undercuts the support that my evidence gives for various propositions. For instance, if S is unwittingly in the demon-world then she has just as much justification for the sorts of perceptual beliefs that we typically form. But, if S were to have sufficient evidence that she is in fact in the demon-world, then her justification would be defeated by this evidence. In both cases her perceptual faculties are equally unreliable, but only in the latter case does she fail to have (all-things-considered) justification.[72]

It is important to understand what this reply amounts to. EFEG doesn't do anything to change the *support relation* between *E* and *p*. As I said above, epistemic support relations are *necessary*. That is, *p* is justified for S at *t iff* S's evidence at *t* on balance support *p*. Evidence justifies necessarily, i.e.—the justificatory status of a proposition for S strongly supervenes on the body of evidence S has. Thus, as Conee and Feldman hold:

> Necessarily if S1's belief *p* is justified, and *E* is evidence that S1 has, then necessarily, (1) on balance *E* supports *p*, and (2) if *E* is the evidence that S2 has, then *p* is also justified for S2.[73]

[72] I would like to thank Kevin McCain (in communication) for helpful comments about the material in the above part of this section.

[73] I changed the original wording of this principle to entail that belief *p* is justified for S1 and S2, rather than Conee and Feldman's original phrasing that entailed that S1 and S2 are justified in believing *p*. As I've argued, it is important to remember that it is *beliefs* and not *persons* that are epistemically justified. Here is the original passage:

> Necessarily if S1 is justified in believing *p*, and *E* is the evidence that S1 has, then necessarily, (1) on balance *E* supports *p*, and (2) if *E* is the evidence that S2 has, then S2 is justified in believing *p* (2008, p.83).

EFEG does *not* undermine the epistemic support relation between evidence *E* and *p*. In other words, EFEG doesn't change the fact that *E* necessarily support *p*. What is peculiar about higher-order evidence like EFEG is that even though it does nothing to undermine the support relation between *E* and *p*, it does call into question whether the truth of *p* is the best explanation for why one has *E*. For example, you might have evidence *E* that on balance supports *p* <*the maid murdered the butler*> (because this support relation is necessary), but your total evidence *E**, which now includes EFEG, undercuts your belief that the truth of *p* <*the maid murdered the butler*> is the best explanation of *E*, even though *E* still supports *p*. In other words, you are given good reason to lower your confidence that the truth of *p* is what best explains *E*. *But it is the evidence that does the justifying*. Without the *awareness* of flawed evidence-gathering, flawed-evidence gathering has no bearing on the justificatory status of your belief that *p* is true. Hence, on the explanationist evidentialist account of justification, reliability or unreliability of one's evidence gathering methods plays no role in the *categorical* justification of your belief that *p* is true.

I.4 - The New Evil Demon Problem

Evidentialists reject the notion that reliability is necessary for justification. I think this is correct. One way to see this is by looking at the 'New Evil Demon' problem (NED). I discussed NED in earlier chapters, but it would be beneficial to lay out the problem once again because later on I will use it to highlight important applications of my two-component view of justification.

Here NED is presented somewhat differently but with the same basic elements. Suppose there are two individuals, person A and person B. Imagine that A and B believe precisely the same thing, undergo precisely the same experiences, have the same memories, have the same

intuitions, and reason in the same manner. Further, imagine that A and B have both had the same non-factive mental states their whole lives. Now, suppose that A lives in the *actual* world (the real world we live in), but B lives in a 'demon-world' (unbeknownst to him) where he is unwittingly being deceived and manipulated by an evil demon. Thus, A's beliefs are reliably produced by processes that reliably lead to the truth, but B's beliefs are formed by unreliable processes (demon-deceptions). For a reliabilist, you cannot have a justified belief unless that belief was formed by some reliable belief-forming process. If the reliabilist is correct, then none of B's beliefs are justified.[74] However, it is a commonly held intuition that the beliefs of A and B are both equally well justified. The evidentialist will say that because A and B have all of the same experiences, memories, reason in the same manner, etc., they are alike evidentially and because they are alike evidentially their beliefs are alike justificationally. That is, if both A and B base their beliefs on the same sufficiently strong evidence, then both of their pertinent beliefs are equally justified, *regardless of the fact that A's beliefs are reliably formed and B's beliefs are unreliably formed.*

Importantly, however, I don't quite agree with this 'commonly held intuition' that 'both A's and B's beliefs are *equally* well justified'. There is a crucial distinction to make here. What I do think is true is that both A's and B's beliefs are justified *period.* The fact that they both have the same supporting evidence guarantees that their beliefs (that are based on that evidence) meet the *bare minimum* (at the very least) required for a belief to be justified. That is, A's and B's beliefs that *p* are *categorically* justified, justified *full-stop,* justified *period,* etc. This is an important distinction because if NED is interpreted in this way, it leaves room for the possibility

[74] Of course, this isn't necessarily true since there may be some beliefs that are justified regardless of whether you are in the demon-world. For example, your introspective belief that you are currently thinking seems immune to demon-world deceptions. However, even if this is so, it doesn't change the implications of the case at hand since almost all of the processes that form beliefs are susceptible to the demon's manipulations.

that even though A's and B's beliefs that *p* are both categorically justified, one of their beliefs might possess a higher degree of justification depending on whether the evidence gathered to support that belief was gathered virtuously (reliably). This picture of categorical justification still preserves the intuition of NED that both A's and B's beliefs that *p* are justified. However, my view also allows for there to be a difference in the *degree* to which (possibly) A's belief that *p* is justified to a higher degree than B's.

Notice how my view of NED is consistent with the evidentialist analysis of EFEG I discussed earlier. Consider person A and B again, but this time they are both in the actual world. Suppose that A virtuously (reliably) gathers evidence X which strongly supports his belief *p*. On the other hand, suppose B viciously (unreliably) gathers the same evidence X which strongly supports his belief *p*. Now, both A and B have precisely the same evidence, namely X, so their beliefs that *p* are both categorically justified regardless of how they gathered their evidence. Further, suppose that A becomes aware of new evidence that his evidence gathering methods were highly reliable and B becomes aware of new evidence that his evidence-gathering methods were highly unreliable. As discussed in the last section, this new evidence about the reliability of one's method of inquiry changes the *total evidence* one has. So, A and B both have new evidence that changes the justificatory status of their beliefs. Person A now has higher-order evidence that *p*'s truth is the best explanation of A's having evidence X. Thus, there is an *increase* in the justificatory status of A's belief that *p*. More precisely, the fact <A inquired reliably> is new evidence that A has for believing that *p*'s truth is what best explains his having evidence X. Moreover, the new total body of evidence X* (evidence X + <A inquired reliably>) supports *p* even more strongly than does evidence X alone, so that A's belief is now (on the basis of evidence X*) even more justified than it was before he acquired the information about his

reliable inquiry. Person B's new higher-order evidence <B inquired *un*reliably> is new evidence B has for thinking that *p*'s truth is *not* what best explains his having evidence X. Moreover, the new total body of evidence X' (evidence X + <B inquired unreliably>) does not support *p* as strongly as does evidence X alone, so that B's belief is now (on the basis of X') less justified than it was before he acquired the information about his unreliable inquiry. Consequently, given A's and B's new evidence about the nature of their respective evidence-gathering methods, A's belief that *p* is far more justified than B's belief that *p*.

This evidentialist analysis is consistent with other problems related to reliable evidence gathering. Consider again the cases of GEORGE and FOOTBALL JON. The challenge from bad evidence-gathering is a challenge to one having propositional justification. In other words, the challenge from such cases is not that, Jon say, has sufficient justification for *p*, but he simply doesn't base his belief on his evidence in the right way. Instead, these cases are attempts to say that, for example, Jon lacks justification for *p* because of how he gathered the evidence. In a sense, the claim is that the evidence that the person has doesn't support *p*—even though, if someone else had that same evidence they would have justification for *p*. This is in line with the evidentialist analysis of what is going on with A and B in NED. What justifies the person's belief here isn't reliability, but instead that the person is basing his belief on sufficiently strong evidence. The evidence in cases like GEORGE and FOOTBALL JON is just as strong as it would be had George and Jon been good epistemic agents in gathering their evidence. That is, the epistemic support relation between George's and Jon's total evidence and their respective beliefs are unaffected by the unreliability of bad evidence gathering, whether they are aware of this unreliability or not. But there is more to say on behalf of the reliabilist about this issue.[75]

[75] I want to thank Kevin McCain (in communication) for helpful comments regarding GEORGE and FOOTBALL JON.

II. The Fittingness-Dimension

In this section I will very briefly argue that EX-WF provides a satisfactory account of one dimension of epistemic justification, i.e.—the fittingness-dimension (FD). What I say here is a culmination of arguments I've already discussed elsewhere. I will adopt McCain's EF (Explanationist Fit) as the main focus of this discussion.

Explanationist Fit (EF):

> *p* fits S's evidence, *e*, at *t* iff either *p* is part of the best explanation available to S at *t* for why S has *e* or *p* is available to S as a logical consequence of the best explanation available to S at *t* for why S has *e*.

EF provides a proper account of what it means for a proposition to be justified for a subject and is of course a key component in EX-WF (well-founded belief). I think the crucial thing behind EF is that it explains when a subject epistemically *ought* to believe a proposition. For example, in the GEORGE case, the evidentialist says that given George's *total* evidence, he is believing precisely how he ought to believe. That is, his total body of evidence supports his belief about second-hand smoke. There is nothing else George ought to believe given his evidence regarding the health effects of second-hand smoke. Absent any defeaters (first- or second-order) George's belief is justified *period*. Two people with the exact same evidence (who are alike mentally), will have the same belief justified for them if that belief is the best explanation available to them of that evidence. In this case their respective beliefs are both *categorically* justified. Let's look at an example.

Consider the following:

Biologists (BIO)

> Biologist A and Biologist B have gathered evidence X. But, Biologist A engaged in *virtuous* inquiry while gathering his evidence and Biologist B engaged in *vicious* inquiry while gathering his evidence. Suppose that Biologist A and B both form the same belief *p* based on evidence X and *p* strongly fits evidence X (i.e.—belief *p* is the best

> explanation of evidence X, etc.). Neither Biologist A nor B is aware of the virtuous or vicious nature of their respective inquiries.

Absent any first- or second-order defeaters, Biologist A and B believe precisely what they should believe given the evidence they have at that given time. Even though Biologist B believes *p* based on evidence that was viciously gathered and given he has no evidence that his methods of inquiry were vicious, what else should he believe (not in the blameless sense)? Belief *p* strongly fits evidence X, so there is a very good reason for Biologist B to believe *p*. Now, compare this example to one in which we plug in these same biologists to NED:

> *Biologists-NED* (BIO-NED)
>
> Biologist A who is in the actual world and his demon-world counterpart Biologist B have exactly the same experiences, apparent memories, intuitions, etc., (they have the same evidence) and in both worlds go through exactly the same processes of reasoning and evidence gathering, and form exactly the same belief *p*. Furthermore, Biologist A and Biologist B form the same belief *p* based on the same evidence and *p* strongly fits that evidence (i.e.—belief *p* is the best explanation of their evidence, etc.). Yet belief *p* is true in the actual world and false in the demon-world.

EF holds that Biologist A and B have the same total evidence. The important question to ask in both BIO and BIO-NED is what each biologist should believe. I have already assessed BIO, so I will now focus on BIO-NED.

In BIO-NED, Biologist B has all the same evidence as Biologist A and therefore has very good reasons to believe the same things that Biologist A does. In fact, Biologist B has *exactly the same reasons to believe the same things* (if 'reasons' is construed as 'evidence'). Therefore, Biologist A's and B's beliefs are alike justificationally along FD. What is crucial in accounting for justification in BIO and BIO-NED is that the beliefs in question are non-factive. This means that the beliefs they have may misrepresent the world. This is because even evidence can support false beliefs (fallibilism). A false belief can still be the best explanation available to a person given the evidence the person possesses at a given time. In both BIO and BIO-NED, the

best explanation Biologist A and Biologist B have of their evidence is belief *p*. However, in BIO it is *possible* that belief *p* is false, and in BIO-NED Biologist B′s belief *p is* false because he is in the demon-world (e.g., Biologist A is looking at actual birds and believes he is looking at actual birds, while Biologist B is looking at a demon-created illusion of birds and believes he is looking at actual birds). Nevertheless, both Biologist A and Biologist B in BIO and BIO-NED, believe exactly what they should believe. If neither of the biologists have evidence of the reliability of their evidence-gathering methods, it would be quite irrational for either of them to believe anything else than what their evidence supports at that time.

Importantly, the sort of justification I′m talking about in BIO and BIO-NED is an objective sort of justification. Remember back to Chapter One when I discussed the distinction between *objective* and *subjective deontological justification*. For example, in BIO NED Biologist B is not justified because he is *blameless* for holding his belief. He is justified because his belief is the best explanation he has available to him for why he has that evidence (even though he is in the demon-world). Being justified along FD is not a case of subjective deontological justification. EF offers an objective sense of justification because it is the type of justification such that Biologist A's belief that *p* could have it and Biologist B's belief that *p* could lack it even if both Biologist A and B are equally epistemically blameless in believing *p*. So, for example, Biologist A might be blameless for believing *p* but the reason that *p* is justified for Biologist A is because *p* is the best explanation available to him for why he has some body of evidence, not because he is blameless for holding *p*. Furthermore, Biologist B might be blameless for believing *p* but be *un*justified in believing *p* because his evidence doesn't support *p* (e.g., a schizophrenic who believes his delusions or someone in a religious community taught to believe what the oracle says).

Hence, if *p* is part of the best explanation available to S at *t* for why S has *e* or *p* is available to S as a logical consequence of the best explanation available to S at *t* for why S has *e*, then S's belief is epistemically justified *period*. If a belief meets this standard, then the person is believing precisely what he/she epistemically ought to believe. Of course, EF is only one important component in EX-WF, but it is the crucial component that a belief is required to meet to be justified for a person along FD. Being justified along FD is all that is required for a belief to meet the bare minimum required to be epistemically justified. However, this sort of categorical justification is also compatible with a belief having varying degrees of justification depending on the strength of evidence the person possesses. Thus, the best explanation available to a person may be a very good explanation and hence the belief is strongly justified, but a best explanation that is very weak may not provide much support for a belief, in which case the belief is weakly justified, but justified nevertheless.[76] Furthermore, being justified along FD means a belief is justified full-stop, regardless of any historical facts that led to the formation of the belief (e.g., vicious inquiry).

III. The Reliability-Dimension

In this section, I will argue that beyond the justification conferred by one's evidence along FD, *reliability*, when relevant, can make a positive difference on the justificatory status of one's beliefs. Thus, I will argue that a person's belief *p* can be justified along FD *and* along the *reliability-dimension* (RD). Again, the picture I have in mind is that a person's belief is epistemically justified to varying degrees depending on how strongly that person's belief fits the evidence that person has at a given time, but that in addition, there is an externalist justification factor (namely reliability) that is often directly relevant to the positive justificatory status of that

[76] Of course, if the best explanation is lousy then the belief isn't justified.

belief. Specifically, I will argue that if a person's evidence was gathered in a reliable way, then (other things being equal) the belief that evidence bears on will have an even higher degree of epistemic justification than it would if it were only justified along FD. Importantly, however, for a belief to be justified it is not necessary that it is justified along RD. On the other hand, for a belief to be justified it must *necessarily* be justified along FD. A belief being justified along RD *only* contributes to increasing the *degree* to which that belief is justified. That is, being justified along RD is not necessary for justification, but it is *sufficient* to increase the degree to which a belief is justified (other things being equal).

Although the evidentialist has strong resources to explain the role that EFEG plays in epistemic justification without appealing to reliability, I will advance an argument that I think supports the view that the reliability of one's evidence-gathering methods does often have a positive bearing on the justificatory status of many of our beliefs. This argument centers on what Comesana (2005) refers to as 'support facts' and what Goldman (2009; 2011) refers to as externalist 'justification factors' ('J-factors'). I will begin by discussing Comesana's view of 'support facts' and then turn toward discussing Goldman's view.

III.1 – Support Facts and J-Factors

According to Comesana, anytime a person is justified in holding some belief, there will be some *facts* in virtue of which that person is so justified. He says that these facts are the truth-makers for the claim that the subject is justified in holding the pertinent belief. More generally, we can distinguish between two types of facts: 1) '*evidence facts*' of the form <*the subject is justified in having some other belief(s)*>, or <*the subject has a certain experience*>; 2) '*support facts*' which are of the form <*the subject's belief(s) is (are) a good reason for thinking that r*> or

<the subject's undergoing a certain experience e is a good reason for thinking that r>

(Comesana, 2005, p.60-1).

To help us understand 'support facts', Comesana has us consider two examples of justified belief (2005, p.60):

1. *Sally believes that the streets are wet (call this proposition q) because she justifiably believes that it is raining and that if it is raining then the streets are wet (call the conjunction of these propositions p).*

2. *Steve believes that the streets are wet because he is looking at them.*

What are the factors that contribute to the justification of the belief that *q* for Sally and Steve? For Sally, one factor that justifies her belief that *q* must include her being justified in believing that *p*. Sally's being justified in believing that *p* constitutes her *evidence* for believing that *q*. For Steve, one of the factors that justifies his belief that *q* is his having a certain visual experience *e* as of the street's being wet. This experience constitutes his evidence for believing that *q*.

Comesana acknowledges that part of what justifies Sally and Steve's belief that *q* is their evidence. Thus, Comesana grants that mentalism is true with respect to evidence, i.e.—that the evidence Sally and Steve have is composed entirely of factors that consist of their mental states. But Comesana argues that Sally's and Steve's evidence doesn't exhaust all of the factors that contribute to their justification for believing that *q*. In addition to Sally's evidence, there is the fact that <*p supports q*>, i.e.—the fact that *p* is a good reason for thinking that *q*. In addition to Steve's evidence, there is the fact that <*e supports q*>, i.e.—the fact that if someone has an experience with the same content as *e*, then this gives the person a good reason to think that *q* (Ibid). Support facts so conceived must exist because if *p* were not a good reason to believe that *q*, then the fact that Sally is justified in believing that *p* would not justify her belief that *q*.

Support facts are of the form: <*p* supports *q* just in case *p* is a good reason to believe that *q*>. There is no mental state that corresponds to the fact that <Sally's being justified in believing that *p* is a good reason for her to believe that *q*>. The support fact in question is a *relation* between those two mental states and this relation is not reducible to either of the two *relata* (Comesana, 2005, p.61).

Now, because of support facts, Comesana thinks that mentalist accounts of epistemic justification will run into a dilemma. Any mentalist theory holds that there are mental factors M_S of a subject S, such that the justificatory status of S's doxastic attitudes is determined by the constituents of M_S. A mentalist who claims that support facts are mental will hold that they are constituents of M_S. If so, Comesana says we can ask the following question: "…is it the case that the obtaining of the facts mentioned in M_S is a good reason for holding the doxastic attitude in question?" (Ibid). For example, is it the case that the obtaining of Sally's mental facts is a good reason for her believing that *q*? Here is the dilemma: If the answer is no, then it appears that Sally is not justified in believing that *q*, and thus the theory has a counterexample. If the answer is yes, then the theory has failed to pick out an epistemically relevant factor, "…because the fact that the obtaining of *all* the epistemically relevant mental factors is a good reason for believing that *q* cannot itself be a mental factor" (Ibid). Therefore, any mentalist epistemic account of support factors is either going to have counterexamples or fail to pick out relevant epistemic factors.

Given Comesana's account of support facts, I think it's easy to see how he would construct an argument against EF (Explanationist Fit) since it's a mentalist account of evidential fit. Remember that EF says

> *p* fits S's evidence, *e*, at *t* iff either *p* is part of the best explanation available to S at *t* for why S has *e* or *p* is available to S as a logical consequence of the best explanation available to S at *t* for why S has *e*.

The fact <*p* is part of the best explanation available to S at *t* for why S has *e*> is itself a support fact. 'Best explanation' in EF is a relation between *p* and *e*, i.e.—*p* is part of the best explanation for why S has *e*. Furthermore, the fact <*p* is available to S as a logical consequence of the best explanation available to S at *t* for why S has *e*> is also a support fact. That is, 'is a logical consequence of' is also a relation between two relata: '*p*' and 'the best explanation available to S at *t* for why S has *e*'. Neither of these two relations is a mental state. Thus, EF, being made up of these two relations, is itself a relation. That is, S's belief *p fits* evidence *e* and thus the fact <a belief *fits* one's evidence> is a support fact of the form <*p* supports *q* just in case *p* is a good reason to believe that *q*>. In the case of EF, *p's* fitting S's evidence *e* is a good reason to believe that *p*. Therefore, there are facts that contribute to S's justification for believing that *p* that are not mental states, i.e.—the fact <S's believing *p* fits the evidence *e* that S has at *t*>, and therefore evidence *e* is a good reason to believe that *p*. If this picture is correct, then EF does not fully exhaust everything that contributes to epistemic justification (at least if EF is conceived as a mentalist account of epistemic support).

If we apply Comesana's dilemma to EF then we get the following: Is it the case that S's belief that *p* fitting S's evidence *e* (*e* being made up solely of mental facts) is a good reason (best explanation) for S's believing that *p*? Here is the dilemma: If the answer is no, then it appears that S is not justified in believing that *p*, and thus the theory has a counterexample. If the answer is yes, then the theory has failed to pick out an epistemically relevant factor, "…because the fact that the obtaining of *all* the epistemically relevant mental factors is a good reason for believing that *p* cannot itself be a mental factor". Therefore, EF, as an epistemic account of support factors

is either going to have counterexamples or fail to pick out relevant epistemic factors. Either way it is an inadequate theory of epistemic justification.

Although Comesana makes some interesting points, I'm not entirely convinced his argument works. First, explanationist evidentialism is compatible with there being externalist support facts. Remember the point made at the beginning of Chapter Four. Fundamentally, an internalist view (at least most) holds that the *justifying* of a belief is done by things internal to the mental life of the subject, i.e.—mental states. On this view, evidence consists entirely of mental states and evidence is what justifies a belief, hence the mentalist component to internalism. Mentalism is a supervenience thesis which means that justification supervenes on mental states. This is best understood by the idea that the *beliefs* of any two individuals that are alike mentally are alike justificationally. Importantly, this internalist and mentalist conception of evidentialism is compatible with the fact that there are *non-mental/external* 'support facts' that are necessary for justification. For example, internalism and mentalism accept that support *facts* about epistemic support (fit) are not internal to the subject nor are these facts mental states—they are necessary truths about which propositions certain mental states make it rational to believe. That is, *facts* like <evidence *e* supports believing *p*>, the *fact* that <*p* is available to S as a logical consequence of the best explanation available to S at *t* for why S has *e*>, and the *fact* that <*p* is part of the best explanation available to S at *t* for why S has *e*> are therefore external to S. However, EF is still internalist because *e* must itself be internal and it is still mentalist because *e* consists solely of mental states.

The above non-mental/externalist factors are compatible with EX-WF, but Comesana is on the right track. Non-mental/externalist facts are important for understanding why EX-WF is not a full account of justification (even though it is a full account of categorical justification). To

show why EX-WF is not a complete account of justification we need to see how reliability (a non-mental/externalist fact) can play a positive role in justification, which is not something EX-WF countenances. If this is the case, then EX-WF does not work as a full account of doxastic justification by itself (even though it is a complete theory of categorical justification), because it doesn't exhaust all of the relevant justificatory factors that contribute to a belief's epistemic justification.

Alvin Goldman (2009) says that whether a theory of justification is classified as internalist or externalist depends on how that theory answers the question: "What kinds of states of affairs determine, or make a difference to, the justificational status of a belief (or other doxastic attitude)?" (p.309). Goldman calls factors that contribute to some extent in fixing a belief's justificatory status *justifiers* or *J-factors* (2009, p.310).

> (J) A justifier of any belief or other doxastic attitude is any property, condition, or state of affairs (and so on) that is positively or negatively relevant to the justificational status of that attitude (Goldman, 2009, p.311).[77]

Again, evidentialist theories such as Conee and Feldman's *well-foundedness* and McCain's EX-WF are internalist and mentalist versions of evidentialism. Such theories are compatible with there being externalist J-factors. Thus, if it can be shown that there are J-factors (namely reliability) that contribute to the justificatory status of a belief to some degree and are not mental states, then two individuals who are alike mentally do not *necessarily* have the same belief justified to the same *degree*. Why? Because there are other non-mental J-factors (reliability) that contribute to some extent in fixing a belief's justificatory status.

[77] The term 'relevant' is not to be understood as purely *causal*. Goldman believes that some, but not all, J-factors will be causally relevant. Instead, 'relevant' should be understood in *explanatory* terms. *Explanatory relevance* is defined in the following way: "X is a J-factor of a given belief's justificational status if and only if X *helps explain* why the beliefs justificational status is what it is."

Goldman offers what he calls 'the right-rule architecture of epistemic justification' (2009, p.313). This 'architecture' is influenced by the idea that justification is a normative term. When dealing with normative domains we often think in terms of 'rules' or 'principles' that govern that domain. Goldman considers the plausible view that the justificatory status of beliefs (and other doxastic attitudes) is (at least partly) constituted by certain conditions, states of affairs, circumstances, or causes, *plus* the governing of epistemic rules. There are objective rules that confer (objective) 'rightness' or 'wrongness' on epistemic conduct. These rules are *right* or *correct* rules of epistemic conduct. Epistemic 'acts' are objectively right or wrong (at least partly) in virtue of the fact that they fall under or violate some correct or legitimate rule.

Goldman formulates the following 'linkage' principle to articulate what he has called 'right-rule architecture':

> (L) *S* is justified in holding doxastic attitude *D* toward proposition *p* at time *t* if and only if there are some conditions *C* such that
> (i) *S* is in conditions *C* at (or before) *t*, and
> (ii) one or more right J-rules jointly permit a subject who is (or was) in *C* to form or retain attitude *D* toward *p* at *t*. (2009, p.314).

(L) says that "…the justificational status of a subject's doxastic attitude toward a proposition is a matter of whether the attitude conforms to what right J-rules authorize in light of the subject's epistemic situation or activity" (Ibid). Importantly, conforming to rules here does not mean *rule-following*, where doing so consists in conforming to rules as a result of *mentally representing* them and verifying that the epistemic act conforms to that rule.

Goldman distinguishes between *inferential* and *non-inferential* J-rules. Inferential rules sanction new beliefs (and other doxastic attitudes) based on logical, inductive, abductive, or probabilistic relations to pre-existing beliefs (and so on). Non-inferential rules sanction new

beliefs in virtue of the occurrence of non-doxastic states or events such as perceptual experiences and apparent memories (Goldman, 2009, p.315). He offers the following as schemas for J-rules:

> (INF) If subject *S*, at time *t*, holds doxastic attitudes $D_1, D_2, \ldots, D_m$ toward propositions *q*, *r,…,s,* respectively, and if proposition *p* bears relation *R* to the conjunction of *q*, *r*, *s*, then *S* is permitted to adopt attitude D_n toward proposition *p* at *t* (Goldman, 2009, p.316).
>
> (NONINF) If subject *S*, at time *t*, is in nondoxastic mental conditions *X*, *Y* and *Z*, then *S* is permitted to adopt doxastic attitude *D* toward proposition *p* at *t* (Ibid).

These schemas are only approximations (and not necessarily exhaustive) of what specific J-rules would need to look like. However, Goldman thinks they suffice as a good starting point.

Beginning with inferential J-rules. INF has a conjunctive antecedent. The first conjunct concerns S's occurrent doxastic states, and the second conjunct concerns some relation *R* of entailment, inductive support, etc., which holds between the propositional contents of those doxastic states and the target proposition *p*. Beginning with the second conjunct, Goldman says, "The fact that a certain logical, inductive, or probabilistic relation holds between specified propositions is a fact independent of any individual's mind" (2009, p.317). For example, the relation of entailment between the conjunction 'All men are mortal' and 'Socrates is a man', on the one hand and 'Socrates is a mortal', on the other, is *not a mental fact or condition*. That is, this relation holds independently of any minds whatsoever. Such conditions, states of affairs, factors, etc., satisfy an antecedent of an inference rule, and therefore will include conditions, states of affairs, factors, etc., that are not mental states belonging to S (Ibid).

There is also another category of external facts that will play a justificatory role. Goldman calls these *instantiation facts* (2009, p.320). To understand what these instantiation facts amount to, Goldman has us consider the following example:

> Melanie believes that Albert has just testified to *P*, that he has no motives for being disingenuous about *P*, and so forth. Does Melanie's cognitive situation permit her to infer

> *P*? Suppose there is a right inferential rule that authorizes one to believe a proposition if it is the best explanation of other things one already believes. And suppose that the truth of *P* is, as a matter of fact, the best explanation of Albert's testifying to *P* and other things Melanie believes about Albert. Then Melanie is indeed justified in believing *P* (Ibid).

What factors are relevant to Melanie's being justified? The answer is similar to the point I made about support facts such as best explanations and logical entailments. First, they include the antecedent beliefs from which Melanie infers *P*. Second is the fact that Melanie's antecedent beliefs and prospective conclusion belief (*P*) jointly *instantiate* the right inferential rule posited. This instantiation fact is highly relevant to the justificatory status of Melanie's belief *P*. "So the instantiation fact helps justify her (that is, helps make her justified) in believing *P*. This instantiation fact, however, is an external fact, not a mental fact. So its status as a justifier cannot be accommodated by mentalism" (Ibid). Importantly, Goldman reminds us what (L) says. A belief (or other doxastic attitude) is justified if and only if it *actually conforms* to the right rule. (L) doesn't require epistemic agents to be aware that their inferential behavior conforms to the right J-rule (Ibid). An awareness requirement would be too strong of a requirement on epistemic justification since most inferential relations are not available to most epistemic agents.[78]

Goldman also discusses non-inferential J-rules. However, I'm not going to discuss these in any detail.[79] Instead, I will focus on Goldman's argument for the *historicity thesis*. This thesis is a *diachronic* thesis about the nature of inferential J-factors. It holds that the justificatory status of a belief held at time *t* partly depends on what transpired in the subject's cognitive history prior to *t*. Therefore, some of the J-factors relevant to a given belief held at *t* will be conditions, states of affairs, etc., obtaining prior to *t* (Goldman, 2009, p.323). Goldman's argument for the historicity thesis hinges on the role of *preservative memory* in helping to

[78] For a more detailed defense of this point see Goldman (2009) Section III.
[79] For a detailed discussion of non-inferential J-rules see Goldman (2009) section VI.

determine a belief's justificatory status. However, I think that the evidentialist has a good way to respond to Goldman's account of preservative memory and therefore I don't think his argument for the historicity thesis works. On the other hand, I think there is a way to support Goldman's historicity thesis that doesn't involve an appeal to preservative memory, but rather one that appeals to the role of virtuous inquiry as a reliable evidence-gathering method.

Goldman's account of preservative memory is supposed to be a problem for explanationist accounts of epistemic support such as those offered by Conee and Feldman and McCain. Goldman says

> If S was justified in believing p earlier, and S retains her belief in p now via preservative memory, then S is prima facie justified in believing p now. Like introspection, however, preservative memory lacks any type of mental experience or episode that invites explanation. In particular, there is no (conscious) act of 'recollection' that invites explanation. An epistemic principle that covers preservative memory, then, cannot be rationalized by the best-explanation approach (Goldman, 2011, p.277).

However, this case fails to impugn explanationist accounts of the support relation. In a direct reply to Goldman, Conee and Feldman (2011b) argue that in cases of preservative memory, by stipulation, S has not forgotten the preservatively remembered proposition *q*. Hence, S "…has the potential to bring *q* to mind with the phenomenology of activating a memory, specifically, the memory that *q*" (Conee and Feldman, 2011b, p.304). S's ability to bring *q* to mind in this way is not "…merely because he has the stored belief that *q*. Many of our stored beliefs are recalled as things we merely believe, not as things we know" (Ibid). This 'potential' to bring a proposition to mind in this way in cases of preservative memory are "dispositions to recollect", which "…is a disposition to bring to mind the proposition as known" (Ibid).

Conee and Feldman hold that "…there is something among his mental states at the time that provides justification for this stored belief. The justifying is not done by any active conscious occurrence at the time, evidentialists need not appeal only to conscious occurrences as

justification" (Ibid). Thus, there is plausibility to the view that S has dispositional evidence. There is further plausibility to the view that the evidence that S has in this example supports his belief in the remembered proposition. Since in a case of preservative memory S has a disposition to recall the proposition remembered "with the phenomenology of activating a memory", it is plausible that S also has a disposition to have a seeming that the truth of the proposition is part of the best explanation S has of why he recalls the proposition with that particular phenomenology (Ibid). In his own response to Goldman, McCain argues that

> …according to *EF* [Evidential Fit], these dispositions (or at least the stored evidence that serves as the grounds for such dispositions) are sufficient for S to have justification for believing the proposition that is retained in preservative memory because they are part of S's evidence, which the proposition fits, so *EF* yields the intuitively correct result in cases of preservative memory (2014, p.76).

Conee and Feldman further argue that having one of these dispositions to recollect is not the same thing as being in a *factive* state. These dispositions to recollect can justify false beliefs that we only seemed to have learned when the dispositions were formed. This is not a case of preservative memory. Goldman says that 'Preservative memory does not create or generate justifiedness from "scratch", but instead transmits a belief 's justifiedness (or unjustifiedness) from one time to a later time' (Goldman, 2009, p.259-60; Conee and Feldman, 2011b, p.304-5). On the other hand, "…a disposition to recollect is a potentially momentary state" (Conee and Feldman, 2011b, p.305). Normally, "…it exists by its having been formed in the past when a proposition was learned, or at least the proposition seemed to have been learned, and then retained by a process of long-term memory consolidation" (Ibid). However, "…a recollective disposition toward a proposition might come about from brain malfunction or tampering" (Ibid). Furthermore, Conee and Feldman argue that

> Whatever causes a disposition to recollect does create evidence for the content proposition from scratch. Having the disposition constitutes having some defeasible evidence for its

> content, *whatever its historical origins* [my italics]. When a recollective disposition resulting from malfunction or tampering leads to a justified true belief, it may nevertheless be a Gettier case rather than a case of knowledge (Ibid).

Thus, contrary to what Goldman thinks, preservative memory does not support the historicity thesis.

I want to offer up another potential way to support the historicity thesis that appeals to a non-mental/externalist J-factor, i.e.—reliability. Suppose the body of evidence *e* that S bases his belief *p* on was gathered in a virtuous (reliable) manner. The next day, S surveys evidence *e* and forms a belief *p* based on this evidence. In this case, the belief *p* that S forms based on evidence *e*, fits the evidence, i.e.—*p* is the best explanation S has for why S has *e*. EX-WF holds that S's belief *p* is justified period because it is the best explanation S has for why S has *e*. What EX-WF says is that the only relevant J-factor is that *p* is the best explanation one has of *e*. But my proposal is that since S had gathered his evidence *e* in a virtuous manner (using reliable evidence-gathering methods), then this reliability positively contributes to increasing the degree to which his belief *p* is justified. I turn to that discussion now.

IV. A Two-Component View of Epistemic Justification

What I now wish to argue is that there is (at least) one important externalist cross-temporal (diachronic) J-factor that bears on the justificatory status of a person's belief that *p*. More precisely, there is a sense in which the reliability of one's evidence-gathering is justificationally *relevant* to the justificatory status of one's attitude toward *p*. Importantly, I'm not arguing that the reliability of one's evidence-gathering methods fully determines the justificatory status of *p*. Rather, as Goldman says, "…being a J-factor only requires being positively or negatively *relevant* to justificational status, not being decisively determinative of it" (2009, p.324). The upshot of my argument is that at least one *diachronic* factor is relevant to

justification after all. More precisely, in some cases a doxastic attitude's justificatory status at *t* is partly determined by the reliability or unreliability of one's evidence-gathering methods. Hence, historicity is relevant to epistemic justification.

In making my argument I will adopt a view of epistemic justification loosely similar to Goldman's (2011) two-component theory of epistemic justification. As I've already discussed, my two-component theory has a *fittingness-dimension* (FD) and a *reliability-dimension* (RD). On FD, a belief is justified if it fits the evidence one has at time *t*.

> *p* fits S's evidence, *e*, at *t* iff either *p* is part of the best explanation available to S at *t* for why S has *e* or *p* is available to S as a logical consequence of the best explanation available to S at *t* for why S has *e*.

But RD is also justificationally relevant. In this case, how the evidence was gathered can be positively or negatively relevant to the justificatory status of a belief. If S's evidence supporting S's belief *p* was gathered *virtuously* (i.e.—reliably), then the evidence-gathering method used will be (other things being equal) *positively* relevant to the justificatory status of *p*. On the flip-side, if S's evidence supporting S's belief *p* was gathered *viciously* (i.e.—unreliably), then the evidence-gathering method used will be (other things being equal) *negatively* relevant to the justificatory status of *p*.

The view I will offer is consistent with cases in which a doxastic attitude's justificational status is *wholly* determined by events occurring at *t* (synchronic). That is, one might have evidence that justifies some belief *p* at *t* where the person didn't engage in inquiry (e.g., the example of your friend blurting out the end of a movie or being overcome with the belief that it's dark when the lights go out). But, in many cases a doxastic attitude's justificatory status will be *partly* determined by the reliability or unreliability of one's evidence-gathering methods. I've spent a lot of time discussing how evidence can support a belief and thereby render that belief

justified period (FD), so I will turn my focus to arguing for RD as it relates to virtuous and vicious inquiry.[80]

IV.1 – Truth-Conduciveness and Justification

Two important things need to be put into place. First, I adopt a view advanced by Goldman (2011, p.278) about *why* inference-to-the-best-explanation is a good pattern for non-deductive inference. It is a good pattern of inference because it is conducive to true belief. This does not mean that no explanation genuinely *explains* anything unless it is true. This last point is consistent with explanationist accounts of justification such as Conee and Feldman's and McCain's. I agree with what Goldman says about explanatoriness. He says that explanatoriness is a good rationalizing property of a type of inference only because it is an excellent indicator of truth. Even if we accept explanatoriness as a mark of justification, the fact that it is such a mark derives from its correlation with truth-conduciveness. Truth-conduciveness is the fundamental principle of epistemic justification, which is completely in step with reliabilism in general. Goldman says "…we do not choose epistemic rules in which the support relation is weak probabilistic support, or inference-to-the-*worst*-explanation! Why not? Because these rules would yield a much lower rate of true-belief formation than the preferred principles" (2009,

[80] I think my view entails that 'non-inquiry' related reliability can play a salient role in fixing the degree to which a well-founded belief is justified. For example, in the case of being overcome by the belief 'It is dark' when the lights are turned out, the visual experience of darkness is evidence that it is dark. My belief that it is dark is the best explanation of my visual experience. Thus, along FD my belief is well-founded (assuming I base my belief on my visual experience). Additionally, however, the fact that my vision is reliable contributes to the degree to which my well-founded belief is justified *even though no inquiry is involved*. This is consistent with my two-component view. It is also consistent with my view to say that even if my vision is not a reliable belief-forming mechanism (because I'm in the demon-world), my belief still fits my evidence and therefore it is justified. I have not focused on reliability more broadly because my main interest has been to try to understand the role that *inquiry* plays in epistemic justification. However, I have not completely ignored this either. For example, I have mentioned the role that reliable vision plays in justification. Careful observation using my vision is a reliable belief-forming process that usually contributes to the degree to which my visual beliefs are justified (not always though because sometimes I'm not very attentive!). Of course, someone might think that engaging in careful observation is engaging in inquiry. I'm fine with this too. My point is simply that there are likely other ways my view is compatible with reliabilism. I count this as a virtue of my view.

p.336). He goes on to say that the standard for J-rule rightness is the tendency of such rules to produce a fairly high ratio of true beliefs. If this is true then the J-factor at the level of the rightness criterion is externalist. Importantly, Conee himself argues that evidence is indicative of the truth: "…all evidence for a proposition, however weak, is some indication that the proposition is true. Thus, the sort of justification that is constituted by evidence always bears on the truth of what is justified" (Conee, 2004, p.253).

IV.2 – Virtuous Inquiry as a J-Factor

So, what is the particular externalist J-factor pertinent to virtuous inquiry? Answer: reliability. Virtuous inquiry is a reliable evidence-gathering process. A subject who engages in virtuous evidence-gathering methods is far more likely to gather a body of evidence that is not *misleading* than a subject who engages in *vicious* inquiry. By 'misleading evidence' I mean:

> ***Misleading Evidence:*** *E* is misleading evidence for *p* just in case (i) *E* is evidence for *p* and (ii) *p* is false.

This should seem relatively uncontroversial. For example, a team of homicide detectives that meticulously surveys a murder scene, properly catalogues relevant evidence, utilizes state-of-the-art forensics techniques and laboratories, etc., is far more likely to gather non-misleading or *veridical* evidence than a homicide team that engages in sloppy, hasty, and less than thorough evidence-gathering. By 'veridical evidence' I mean:

> ***Veridical Evidence:*** *E* is veridical evidence for *p* just in case (i) *E* is evidence for *p* and (ii) *p* is true.

Thus, there is a crucially important *metaphysical fact* related to virtuous inquiry:

> **Metaphysical Fact$_1$:** *other things being equal*, virtuous evidence-gathering significantly increases the probability that the evidence gathered is veridical.

There is also a crucially important *epistemic fact* that falls out directly from this metaphysical fact:

> **Epistemic Fact$_1$:** *other things being equal*, the metaphysical fact that virtuous evidence-gathering significantly increases the probability that the evidence gathered is veridical, means that the *degree* to which a well-founded belief *p* is justified for S *increases* when the evidence supporting *p* was virtuously (reliably) gathered.

If S bases his belief *p* on evidence *E* and *E* sufficiently supports *p*, then *p* is categorically justified for S. But in addition, if S gathered *E* virtuously (reliably), then the *degree* to which *p* is justified *increases* because the likelihood of *E* being veridical increases.

Importantly, reliability is not a *necessary* condition for epistemic justification. FD says that subject S's belief *p* can satisfy the minimal requirement for epistemic justification if evidence *E* is the best explanation S has available for why S has *p*. This is the bare *minimum* required for S's belief *p* to be epistemically justified. Thus, if a belief meets this minimum requirement for justification, then that belief is justified *period.* However, the fact that *E* was reliably gathered is, other things being equal, *sufficient* to increase the *degree* to which *p* is justified for S. However, this is not the only way that a belief's justification can be increased. The strength of one's evidence also determines the degree to which a belief is justified. A belief can be justified to a higher degree if the evidence supporting that belief is strong. On the other hand, a belief that is supported by weaker evidence is less justified than a belief that is supported by stronger evidence. Thus, strength of evidence and reliable inquiry can both increase the justification level of a well-founded belief.

On the flip-side, *vicious* inquiry is an unreliable evidence-gathering process. A subject who engages in *vicious* evidence-gathering methods is far more likely to gather a body of evidence that is *misleading* than a subject who engages in *virtuous* inquiry. Again, a team of homicide detectives who engages in sloppy, hasty, and less than thorough evidence-gathering, etc., is far more likely to gather misleading or *non-veridical* evidence than a homicide team that

engages in virtuous inquiry. Thus, there is a crucially important *metaphysical fact* related to vicious inquiry:

> **Metaphysical Fact$_2$:** *other things being equal*, vicious evidence-gathering significantly increases the probability that the evidence gathered is misleading.

There is also a crucially important *epistemic fact* that falls out directly from this metaphysical fact:

> **Epistemic Fact$_2$:** *other things being equal*, the metaphysical fact that vicious evidence-gathering significantly increases the probability that the evidence gathered is misleading, means that the *degree* to which a well-founded belief *p* is justified for S *decreases* when the evidence supporting *p* was viciously (unreliably) gathered.

If S bases his belief *p* on evidence *E* and *E* sufficiently supports *p*, then *p* is categorically justified for S. But, if S gathered *E* viciously (unreliably) then the *degree* to which *p* is justified *decreases* because the likelihood of *E* being veridical decreases.

Importantly, vicious inquiry cannot render a well-founded belief unjustified. This is because reliability is not a necessary condition for a belief's justification. Satisfying FD is a necessary condition for justification. In other words, having evidence that on balance supports one's belief is a requirement for justification and is also sufficient for a belief to be justified, i.e.—if a belief meets this requirement it is justified *period*. There is no further requirement that a belief be reliably formed for it to be justified. So, if a belief is justified along FD it cannot be rendered unjustified by vicious inquiry.

IV.3 – An 'Awareness' Requirement for Justification?

At this point, it is important to look at a potential problem for my view. McCain argued that virtuous or vicious inquiry only matter to epistemic justification if the subject is *aware* that he/she engaged in one or the other types of inquiry. In other words, if a person has *evidence* that his evidence-gathering method was unreliable (careless, biased, etc.), then this evidence will

weigh against his belief's justificatory status. This is correct. But without such evidence, how one gathered evidence has no bearing on a belief's justificatory status. We are now in a position to see why this latter claim is wrong.

It is true that in order for *evidence* to support, undermine, or rebut some belief *p* for person S, S must be *aware* (occurrent or dispositional evidence) of that evidence. Thus, in order for the *evidence* that S engaged in virtuous/vicious inquiry to have an effect on the justificatory status of S's belief that *p*, S must be aware of that evidence. But this only applies when we consider FD. That is, when considering FD we are concerned solely with how *evidence* provides justification for a belief. In order for S's *evidence* that he engaged in virtuous/vicious inquiry to have any bearing on whether S's belief that *p* is justified, *p* must fit the evidence S possesses and S must also properly base his belief *p* on the relevant evidence, i.e.—S's belief must be well-founded. Additionally, remember that evidence need not be *occurrent*, since there is also *dispositional* evidence. Thus, S need not have occurrent evidence in order for *p* to fit S's evidence since *p* can fit S's dispositional evidence as well. Regardless of this fact, an *awareness requirement* cannot apply to all J-factors even if it does apply to evidence. Of course, EX-WF is consistent with this claim since it is compatible with there being non-mental/externalist J-factors that a person is unaware of. For example, a belief *p* is justified for S if *p* is the best explanation for why S has evidence *E*, regardless of whether S is *aware* that *p* is the best explanation for why he has evidence E. Having such an awareness requirement would overintellectualize justification because then beliefs held by unreflective adults and children wouldn't be justified. But EX-WF doesn't allow for reliability to play a role in the justificatory status of a belief unless the person holding the belief is aware of the reliability. It's this claim that I will argue against next.

Bergmann (2006) argues that there cannot be an *awareness requirement* on epistemic justification. That is, there cannot be an epistemic requirement that a person be either actually or potentially *aware* of the relevant J-factor (or as Bergmann puts it "justification-contributor") that justifies a belief in order for that belief to be justified for that person. Bergmann argues that if there is such an awareness requirement on justification, then we run into a dilemma. The awareness requirement must be either *strong* or *weak*. *Strong awareness* involves *conceiving* of the J-factor that is the object of awareness as being in some way relevant to the justification or truth of the belief. *Weak awareness* does *not* involve conceiving of the J-factor that is the object of awareness as being in some way relevant to the justification or truth of the belief. Bergmann's claim is that if the awareness requirement for justification is strong, then vicious regress problems arise and if the awareness requirement is weak, then the motivation to think that justification is internal is lost. Either way there cannot be an awareness requirement on epistemic justification. Let's look at his argument in more detail.

One of Bergmann's main goals is to argue against internalism. In doing so, he argues that all forms of internalism require, for a belief's justification, that the subject holding the belief be aware (or potentially aware) of that which contributes to its justification.[81] That is, if there is no actual or potential awareness of anything that contributes to the belief's justification, then the belief is not justified. Bergmann calls this the *awareness requirement*.

> *The Awareness Requirement*: S's belief B is justified only if (i) there is something, X, that contributes to the justification of B—e.g. evidence for B or a truth indictor for B or the satisfaction of some necessary condition of B's justification—and (ii) S is aware (or potentially aware of X (Bergmann, 2006, p.9).

[81] For Bergmann's defense of this claim see Bergmann (2006) Chapter 3. My purpose here is not to discuss Bergmann's argument against internalism. Instead, I'm focusing on Bergmann's important argument against any awareness requirement for epistemic justification. My two-component view of justification is compatible with there not being such a requirement and actually relies on there not being one.

Bergmann cites Bonjour's famous 'Norman the Clairvoyant' case as the most influential motivation for an awareness requirement (and one of the most influential arguments against externalism) (2006, pp.11-12). Recall that Bonjour thinks Norman lacks justification for his clairvoyant belief about the whereabouts of the President of the U.S. Norman's belief lacks justification because he is unaware of any relevant J-factor. From Norman's perspective, his belief is accidental, arbitrarily formed, or the result of luck (or something).

Bergmann calls Bonjour's objection to externalism the *Subject's Perspective Objection* (SPO).

> SPO: If the subject holding a belief isn't aware of what that belief has going for it, then she isn't aware of how its status is any different from a stray hunch or an arbitrary conviction. From that we may conclude that from her perspective it is an accident that her belief is true. And that implies that is isn't a justified belief (2006, p.12).

Bergmann points out that the SPO is used to argue against externalism but it is also used to motivate internalism. He thinks that this motivation for internalism is also the strongest motivation for the awareness requirement. Now let's see how all of this fits in with Bergmann's dilemma.

This first horn of Bergmann's dilemma is that a strong awareness requirement on justification gives rise to vicious regress problems leading to radical skepticism. Strong awareness involves conceiving of the J-factor that is the object of awareness as being somehow relevant to the justification of the belief. There are different types of strong awareness: doxastic and non-doxastic versions of strong awareness and actual and potential strong awareness. Doxastic strong awareness involves the *belief* that the object of awareness in somehow relevant to the justification of the belief in question. Non-doxastic strong awareness is simply strong awareness that isn't doxastic. An actual strong awareness requirement holds that the person be aware, while the potential strong awareness requirement holds that the person be able on

reflection alone to be aware. Each type of strong awareness falls prey to the vicious regress problem (Bergmann, 2006, p.14).

Bergmann first considers the *Actual Doxastic Strong Awareness Requirement* (ADSAR).

> ADSAR: S's belief B is justified only if (i) there is something, X, that contributes to the justification of B and (ii) S is actually aware of X in such a way that S justifiedly believes that X is in some way relevant to the appropriateness of holding B (Bergmann, 2006, pp..14-15).

ADSAR leads to regress problems. In order for S's belief B to be justified, S must have the further justified belief (with respect to something, X_1, that contributes to the justification of B) that:

> P_1: X_1 is somehow relevant to the justification of believing B.

According to ADSAR, in order for S's belief that P_1 to be justified, S must have the additional justified belief (with respect to something, X_2, that contributes to the justification of P_1) that:

> P_2: X_2 is somehow relevant to the justification of believing that [X_1 is somehow relevant to the justification of believing B].

And in order for S's belief that P_2 to be justified, S must have the further justified belief ((with respect to something, X_3, that contributes to the justification of P_2) that:

> P_3: X_3 is somehow relevant to the justification of believing that [X_2 is somehow relevant to the justification of believing that [X_1 is somehow relevant to the justification of believing B]].

This regress continues *ad infinitum*. The implication, given ADSAR, is that B is justified for S only if S actually holds an infinite number of justified beliefs of ever-increasing complexity. The issue is that it is difficult enough for anyone to grasp a proposition like P_5 or P_6, let alone believe it with justification. Thus, ADSAR leads to the conclusion that none of our beliefs are justified. Therefore, we have good reason to reject ADSAR.

A similar line of reasoning applies if we switch from actual strong awareness to merely requiring potential strong awareness. I will not go into Bergmann's argument against potential strong awareness, but briefly the issue is that it falls prey to the same regress in which S must hold an infinite number of beliefs of ever-increasing complexity (see Bergmann, 2006, p.16). This vicious regress also occurs if we switch to a non-doxastic awareness requirement. Instead of beliefs, the issue becomes that a person would have to have an infinite number of *concepts* of ever-increasing complexity. Again, I will not discuss Bergmann's argument for this either (see Bergmann, 2006, pp.17-18). For my purposes, the above discussion suffices to show that there is very good reason to doubt that a strong awareness requirement is necessary for a belief to be justified. But what about a *conceptual* weak awareness requirement? This brings us to the second horn of Bergmann's dilemma.

First, it will be instructive to see why a non-conceptual weak awareness requirement won't work. If the sort of awareness that is required is non-conceptual, then S can have the required awareness of the relevant J-factor without conceiving of it in any way (without classifying the J-factor according to any classificatory scheme). But then S can be non-conceptually aware of the relevant J-factor without conceiving of it as relevant in any way to the justification of his belief. SPO, says however, that if S does not conceive of the relevant J-factor as something relevant to the justification of his belief, then from S's perspective his belief is an accident, lucky, arbitrary, etc. Thus, in order for a weak awareness requirement to work it must be conceptual. However, Bergmann argues that such a switch doesn't work either since S could meet this requirement simply by being aware of the relevant J-factor and applying some concept or other to it (such as the concept of *being a physical token process*). This would mean S can have a conceptual weak awareness of the relevant J-factor without conceiving of it as relevant in

any way at all to the justification of his belief. However, the result is that a conceptual weak awareness requirement would be vulnerable to the SPO, but this violates the very intuition that motivated internalism to begin with. According to the SPO, it would still be an accident from S's subjective perspective that his belief is true. Even though S applies a concept to the relevant J-factor, it's not the correct type of concept. That is, the concept is irrelevant to whether the J-factor contributes to the justification of his belief. Thus, the only way to guarantee that S applies the correct concept to the relevant J-factor is to have his belief satisfy a *strong* awareness requirement. But if satisfying a strong awareness requirement is what is demanded, then we must deal with the vicious regress problem (Bergmann, 2006, pp.20-21).

There are good reasons to think that Bergmann's dilemma isn't successful, but a full exploration of this issue goes beyond the scope of this dissertation. However, I want to briefly discuss why Bergmann's dilemma isn't a problem for EX-WF. First, McCain's account of justification doesn't require strong awareness. He points out that A person S need only be disposed on the basis of reflection alone to recognize (have a seeming) that *p* is part of the best (maybe the only one she is aware of) answer to the question "why E?" or "why do I have E?" or "why do I have *this*?" (where 'this' refers to some relevant body of evidence). Having a strong awareness requirement would over-intellectuallize justitication. A person need not have robust concepts like "evidence", "entailment", "best explanation", "logical consequence", etc. McCain is correct in pointing out that it is plausible to think that children and unreflective adults (I would argue that *most* adults don't reflect in the way that philosophers might about justification) possess evidence that supports propositions for them, but surely they don't form higher-level beliefs about their epistemic support. Additionally, it is plausible that a person's evidence can support believing a proposition for him even if he hasn't consciously considered the support

relation between his evidence and that proposition. Thus, direct awareness of this support relation doesn't seem necessary for justification (McCain, p.78).

Even if some sort of weak awareness is required, I think EX-WF is safe from Bergmann's dilemma. At most, McCain's view requires that S be disposed to either have a seeming that *p* is part of the best explanation available to S at *t* for why S has the evidence he does or have a seeming that *p* is logically entailed by the best explanation available for why S has the evidence he does in order for *p* to be supported by his evidence. Lacking this 'seeming' requirement would seem to imply that S could have epistemic support for a proposition that he lacks the requisite concepts to even understand. It's highly implausible that S can have epistemic support for propositions he can't understand (Ibid).

Jason Rogers and Jonathan Matheson (2011) argue that 'weak awareness internalism' does not succumb to the SPO. They provide an account of 'seemings' in which they are neither conceptual nor beliefs. Rogers and Matheson offer an account of 'seemings' that appeals to 'seemings' as *justification-contributors*. McCain uses 'seemings' in a different sense. On McCain's account, 'seemings' are *not* justification-contributors. On McCain's view, 'seemings' allow us to have *weak* awareness of the relationship between some body of evidence and a proposition. However, McCain can still appeal to the general idea behind Rogers and Matheson's view to support his own. Rogers and Matheson argue that "A subject may "host" a seeming without conceiving of—or even having higher-order awareness of—the seeming *itself* in any way. Likewise, the seeming may result for the subject as a result of merely *weak* awareness of some object of first-order awareness" (Rogers and Matheson, pp.60-1). The idea is that S can be in a state wherein he hosts a seeming while remaining in a state of weak—non-conceptual—awareness of that seeming, or while having no higher-order *awareness* of the

seeming whatsoever, and all while staying in a state of weak awareness concerning the object of first-order awareness that causes that seeming (Rogers and Mathson, p.61).

Rogers and Matheson think that such states of seeming can make a relevant difference from the subject's perspective, i.e.—that the relevant belief is appropriate to hold.[82] Apply this to McCain's view, i.e.—hosting a seeming state, such as one that involves it seeming that *p* is the best explanation of some body of evidence, renders *p non-accidental* from S's perspective. Norman the Clairvoyant's belief just 'pops' into his head. From his perspective his belief is accidental, arbitrary, or something else. On the other hand, S's seeming state makes it the case that, from his perspective, *p* isn't accidental, arbitrary, etc. Having the seeming can make *p* reasonable from S's perspective. From Norman's perspective, his belief is inexplicable. Hence, no (higher-order) awareness of or conceptualization of the seeming itself is required for S's belief to be relevantly "non-accidental" from his perspective (Rogers and Matheson, pp.61-2). Consequently, McCain's view can avoid the SPO.

Thus, Bergmann's dilemma does nothing to undermine the view that a belief's being the best explanation a person has for why he has some body of evidence is, at the very least, minimally justifying. That is, if a belief *p* is the best explanation the person has for why he has evidence *e*, then *p* is justified full-stop for S. This is because evidential relations hold necessarily. What is at issue is whether rebutting evidence or undermining evidence are the only things that can bear negatively on the justificatory status of a well-founded belief. I have argued that this is not the case. I have argued that whether someone engages in virtuous or vicious inquiry can make a justificatory difference in the *degree* to which a belief is justified. On my two-component view there need not be an awareness requirement on the reliability of inquiry in

[82] Again, Rogers and Matheson are concerned with a somewhat different sense of 'seeming', but the basic idea of a 'seeming' and therefore the general approach to resisting Bergmann's dilemma is applicable to McCain's view.

order for it to play a role in epistemic justification. In order for reliability of inquiry to play a role along FD, a person must be aware (have evidence) that the evidence he gathered wasn't gathered reliably. However, along RD the person need not be aware of the reliability (or unreliability) in order for it to play a role in the belief's justificatory status. It is this latter view that Bergmann's arguments support.

Again, if a person has evidence that he engaged in virtuous inquiry then he has good reason to believe that the evidence he gathered is not misleading. And if the person has evidence that he engaged in vicious inquiry then he has good reason to believe that the evidence he gathered is more than likely misleading. Such evidence will affect the degree to which that belief is justified for that person. But McCain says that if a person is not aware (actually or potentially) of any evidence regarding the nature of the inquiry he engaged in, then the reliability (or lack thereof) cannot bear on the justificatory status of his belief. I agree with this but only as it applies to FD. *Evidence* can't bear on the justificatory status of a belief unless the person is aware of that evidence. However, as I've said before, whether someone engaged in virtuous inquiry *can* have an effect on the justificatory status of a belief along RD. Again, there is the *metaphysical fact* that, other things being equal, virtuous evidence-gathering significantly increases the probability that the evidence gathered is veridical. There is also the crucially important *epistemic fact* that falls out directly from this metaphysical fact: *other things being equal*, the metaphysical fact that virtuous evidence-gathering significantly increases the probability that the evidence gathered is veridical, means that the *degree* to which a well-founded belief *p* is justified for S *increases* when the evidence supporting *p* was virtuously (reliably) gathered (and vice versa for vicious inquiry). Bergmann's argument against an awareness requirement on justification helps support my claim. That is, it is not required that a

person be aware of this metaphysical fact (or the epistemic fact) in order for it to have an epistemic bearing on the justificatory status of his belief. Thus, the reliability of one's inquiry can have a positive bearing on the degree to which a person's belief is justified *even if the nature of the inquiry has no bearing on evidential relations*. Let me explain.

Evidential relations hold regardless of the reliability of the inquiry the person used to gather his evidence. That is, *p* either fits a body of evidence or it doesn't. If *p* fits a body of evidence, then it does so necessarily. Reliability does nothing to undermine this relation between a belief and the evidence which supports it. This means the *evidence* one has justifies the same *regardless of how it was gathered*. Therefore, absent any *defeating* evidence, a person's belief is justified *period* as long as that belief fits the evidence he has. But the degree to which the belief is justified can be affected by virtuous (reliable) or vicious (unreliable) inquiry. To see how this works it will be instructive to apply my view to some cases.

IV.4 – Applications

What will be helpful to understand my two-component theory is to see how my view applies to two persons in similar and dissimilar epistemic conditions. Person A and Person B are both epistemic agents who I'll place in different epistemic conditions. Both A and B can be in either the 'actual world' (the verific world we inhabit) or the 'demon-world' (the world in which an evil-demon manipulates all of the evidence one has by always providing misleading evidence, i.e.—evidence that points to the truth of some proposition, but that proposition is false). Further, A and B can either engage in virtuous (reliable) inquiry or vicious (unreliable) inquiry when gathering evidence. Importantly, A and B are equally capable cognitive agents.

My two-component theory of epistemic justification holds that, *other things being equal*, someone's well-founded belief *p* that is based on evidence that was gathered via *virtuous* inquiry

is justified to a higher degree than someone's well-founded belief *p* that is based on evidence that was gathered via *vicious* inquiry. The key is the 'other things being equal' requirement. In a demon-world, other things are *not* equal. The demon manipulates the evidence you gather no matter what method of inquiry you use. So, *all things considered*, how you inquired in the demon-world has no positive epistemic bearing on whether your evidence is veridical because it is *always* misleading (I will say more about this below). In fact, virtuous inquiry is *always* an unreliable evidence-gathering process in the demon-world. On the other hand, in the 'actual' world it *does* matter if you inquire virtuously because virtuous inquiry is normally a reliable evidence-gathering process, since normally other things are equal.

I'd like to make one more point about virtuous inquiry before I discuss the different applications of my two-component view. It might seem inconsistent to say that virtuous inquiry can be unreliable since I've stipulated that virtuous inquiry is a reliable belief-forming process. This is true, but only when we fail to consider the 'other things being equal' condition. In a verific world (the actual world), other things being equal, virtuous inquiry *is* reliable. Key to understanding this is recognizing that inquiring virtuously entails performing activities such as reasoning open-mindedly and fair-mindedly, reasoning consistently, engaging in careful visual observation, etc. Thus, in the actual world, other things being equal, performing these types of activities (among others) is a reliable way to gather and sort through evidence. But in a demon-world these activities are rendered completely unreliable. This is because all things considered performing these activities is always an unreliable way to gather evidence. Thus, there is no inconsistency in saying that virtuous inquiry can be unreliable. It is equivalent to saying that X is a reliable belief-forming process under normal conditions, but that X can be unreliable in

conditions that are not normal. Now I will look at how my two-component view applies to different cases.

My two-component theory commits me to the following list. Assume that in each case, A's and B's beliefs are equally well-founded (they have precisely the same total body of evidence) and are therefore equally justified along FD. Further, assume that they have no evidence of whether their inquiries were virtuous or vicious.

1. A and B inquired *virtuously* and both are in the actual world = A's belief *p* and B's belief *p* are equally justified.
2. A and B inquired *viciously* and both are in the actual world = A's belief *p* and B's belief *p* are equally justified.
3. A and B inquired *virtuously* and both are in the demon-world = A's belief *p* and B's belief *p* are equally justified.
4. A and B inquired *viciously* and both are in the demon-world = A's belief *p* and B's belief *p* are equally justified.
5. A inquired *viciously* and B inquired *virtuously* and both are in the demon-world = A's belief *p* and B's belief *p* are equally justified.
6. A inquired *viciously* and B inquired *virtuously* and both are in the actual world = B's belief *p* is more justified than A's belief *p*.
7. A inquired *viciously* and is in the actual world, while B inquired *virtuously* and is in the demon-world = A's belief *p* and B's belief *p* are equally justified.
8. A inquired *virtuously* and is in the actual world, while B inquired *viciously* and is in the demon-world = A's belief *p* is more justified than B's belief *p*.
9. A inquired *virtuously* and is in the actual world, while B inquired *virtuously* and is in the demon-world = A's belief *p* is more justified than B's belief *p*.
10. A inquired *viciously* and is in the actual world, while B inquired *viciously* and is in the demon-world = A's belief *p* and B's belief *p* are equally justified.

Notice that in 1-10 the evidentialist will say that A and B are equally justified because evidence is all that matters for justification and by stipulation they have precisely the same evidence. But on my view, some of these will be diagnosed differently. I think numbers 1-4 are uncontroversial since A and B are in identical epistemic conditions and inquired the same way. I also think numbers 8 and 10 are diagnosed correctly on my view. Number 6 seems diagnosed correctly on my view as well, which is also most important to my theory (I will discuss this below). Numbers 5, 7, and 9 are the most controversial and will require more support.

Number 5 says that A and B are both in the demon-world, have the same well-founded belief *p*, but A inquired viciously while B inquired virtuously. The reason that their respective beliefs *p* are equally justified is because we have to consider both FD and RD. On FD, A and B have equally well-founded beliefs. However, on RD, A's and B's inquiries are both unreliable because they are in a deeply impoverished epistemic condition, i.e.—they are both in a demon-world. Thus there is no epistemic difference between A's belief and B's belief along RD. Virtuous and vicious inquiry are equally unreliable in the demon-world. One initial objection here might be that 5 conflicts with 6. That is, if virtuous inquiry makes a positive epistemic difference in the actual world, then it should make an epistemic difference in the demon-world. But this objection fails to take into consideration two important points. First, in the demon-world other things are *not* equal. That is, this objection fails to grasp just how epistemically impoverished someone's condition is in the demon-world. This is where the 'other things being equal' constraint is crucial. The Evil Demon manipulates all of your evidence all of the time, making sure that your evidence is *always* misleading *regardless* of whether your inquiry was conducted in what is a normally reliable manner. So, in this case all things are not equal. *All things considered*, in the demon-world it is *not* a metaphysical fact that virtuous inquiry increases the likelihood that your evidence is veridical because one's evidence is *always* misleading in the demon-world. Consequently, in the demon-world, virtuous inquiry cannot increase the degree to which a belief is justified. Second, and related to the first response, when we talk about epistemic justification we are talking about a positive normative status that makes a belief a much better candidate for amounting to *knowledge* than it would be otherwise. Whether it is such a candidate has a lot to do with truth-conducivity. In the demon-world,

virtuous inquiry just can't get your belief any closer to knowledge because you're in a global deception scenario. Thus, my diagnosis of 5 holds.[83]

Of course, one might ask why it is that justification still holds along FD in a demon-world if the evidence one has can never amount to anything like knowledge. That is, if when we talk about epistemic justification we are talking about a positive normative status that makes a belief a much better candidate for amounting to *knowledge* than it would be otherwise, how would this apply to evidence in the demon-world that is by its very nature misleading? Conee (2004) argues that a proposition is epistemically justified for a person when it is evident to the person that the proposition is true. Further, any support a person has for a proposition is some sort of indication to the person that the proposition is true. Evidence is what indicates the truth of a proposition. The stronger a proposition's truth is indicated to a person, the better that proposition is epistemically justified for that person. Epistemic justification of a proposition is evidence of its truth. The *relation* of evidential support is its truth connection (Conee, 2004, pp.252-53).[84]

Furthermore, evidential relations are *necessary*, regardless of what world you are in. In the actual world a body of evidence can still be strongly supportive of some proposition even if that evidence is misleading. Thus, belief *p* can still be the best explanation available to a person for why he has evidence *e* even if that evidence is misleading. Likewise, being in a demon-world has no bearing on whether evidential relations hold. Evidence still justifies the same in the

[83] There is another possible world that is relevant to mention here. There could be a world that is antithetical to the demon-world, say a 'Good Angel' world, in which a good-angel makes all inquiries reliable. That is, there could be a world in which a good-angel makes vicious inquiry always reliable so that even vicious inquiry always yields veridical evidence. This would be the opposite scenario that we see in number 5. Thanks to EJ Coffman for suggesting this 'angel-world'.

[84] I should point out that truth is the unifying concept in my two-component theory of epistemic justification. That is, along FD the relation of evidential support is its truth connection. Along RD, the truth connection has to do with the fact that a reliable process yields true beliefs or veridical evidence. Thus, FD and RD do not rely on disparate fundamental concepts. Instead, FD and RD both rely on the same fundamental unifying concept, i.e.—truth.

demon-world as it does in the actual world. That is, no matter what world you are in, evidence still indicates the truth of a proposition. On the other hand, the reliability of some method of inquiry does not stay consistent between the actual world and the demon-world. Epistemic justification along FD is still an important sort of normative status required for a belief to *ultimately* amount to knowledge, even if that evidence isn't strong enough to support a belief to the degree required for knowledge. An unreliable method of inquiry can never confer the type of normative status a belief needs to amount to knowledge. Another way of putting this is to say that a well-founded belief is a rational belief, but a belief that isn't based on any evidence is not rational. Think about what evidence does for a believer. Evidence indicates that something is true (evidence is evidence that something is true, i.e.—evidence is always truth indicative). When you have evidence for a proposition (even misleading evidence) you have something that indicates the truth of that proposition. That is, when you have evidence of *p* you have reason to believe *p*. Having a reason to believe *p* is necessary for knowledge. It doesn't make sense to say that 'I know *p* but I have no reason to believe *p'*. Thus, having evidence (even when that evidence is always misleading—and you have no evidence of its misleading nature) does make a belief a better candidate for knowledge than it would be otherwise.

What about 7? Number 7 says that A is in the actual world and inquired viciously, while B is in the demon-world and inquired virtuously. Their beliefs are equally justified along FD because both of their beliefs are equally well-founded and thus justified period. Now, along RD, A engaged in vicious inquiry which decreases the degree to which A's well-founded belief is justified. Now, because B is in the demon-world, his virtuous inquiry does not increase the degree to which his belief is justified. In the demon-world, virtuous inquiry is completely unreliable and hence it decreases the degree to which B's well-founded belief is justified, i.e.—

all things considered, in the demon-world it is *not* a metaphysical fact that virtuous inquiry increases the likelihood that your evidence is not misleading because one's evidence is *always* misleading in the demon-world. In the demon-world virtuous inquiry is unreliable, so it *is* a metaphysical fact that virtuous inquiry *increases* the likelihood that your evidence *is* misleading. Furthermore, A's engaging in vicious inquiry in the actual world is also a form of evidence-manipulation, albeit not of the demon variety. That is, in an important sense, a person who engages in vicious inquiry in the actual world or the demon-world is having his evidence systematically manipulated. Rather than a global deception like that in the demon-world, a person in the actual world who engages in vicious inquiry is part of a localized 'deception'. Just like with 5, in order to appreciate 7, one must appreciate the impoverished epistemic condition one is in when in the demon-world and how similarly manipulated the vicious thinker's evidence is in the actual world. Thus my diagnosis of 7 holds.

What about 9? The correctness of number 9 should be easy to see by now. Number 9 involves person A inquiring virtuously in the actual world and person B inquiring virtuously in the demon-world. Obviously, person A's well-founded belief *p* is justified to a higher degree than person B's well-founded belief *p* because virtuous inquiry is a reliable evidence-gathering process in the actual world but is an unreliable evidence-gathering process in the demon-world. Thus, the degree to which B's belief *p* is justified *decreases* in the demon-world because virtuous inquiry is completely unreliable.

Now, of primary importance to my two-component theory is the truth of number 6 (A inquired *viciously* and B inquired *virtuously* and both are in the actual world, so B's belief *p* is more justified than A's belief *p*). This is because what we are primarily concerned with is the epistemic status of our beliefs in the actual world and my view makes good sense of how our

beliefs get justified in the actual world. In the actual world, other things are *normally* equal. I say 'normally' because someone might have the concern that people can, and often do, manipulate evidence in the actual world (a localized deception), so other things are not *always* equal in the actual world. The concern is that sometimes the actual world can be similar to the demon-world in this respect. I am happy to grant this fact (and I think it's true), but this doesn't impugn my view. If someone is manipulating the evidence in the actual world, then other things are not equal, so all things considered, engaging in virtuous inquiry may not increase the likelihood that the evidence is veridical. However, unlike in the demon-world, it is very possible to discover (and often is discovered) that the evidence has been manipulated (i.e.—tampered with or planted, etc.). This is not possible in the demon-world because the person has no way to confirm or disconfirm that he is in a demon-world or that his evidence is being manipulated. So, in the actual world, virtuous inquiry can still be effective in sorting through manipulated evidence and therefore can have a positive epistemic bearing on one's belief in the actual world.

What's important to notice about number 6 is that A and B have the same belief *p*. This is crucially different than a case in which A and B are in the actual world but have different beliefs about the same subject matter. For example, A might have belief *p1* about some subject X and B might have belief *p2* about subject X (suppose both beliefs are well-founded). Further, A inquired viciously when gathering his evidence *E1* in support of *p1* and B inquired virtuously when gathering his evidence *E2* in support of *p2*. It is highly probable that *E2* provides stronger evidential support for B's belief *p2* than *E1* provides for A's belief *p1*. For example, take Gerry's case in which his belief about second hand smoke (that it's not harmful) fits his evidence, but he inquired viciously. Now suppose Gerry* gathered his evidence virtuously and thinks that second hand smoke is harmful. It is highly probable that Gerry's belief will have

much weaker evidential support than Gerry*'s belief. Gerry's belief is well-founded but his supporting evidence is very likely not as strong as Gerry*'s. Thus, along FD, B's belief *p2* is better justified than A's belief *p1*. However, we must also consider how A's and B's beliefs are justified along RD. Along RD, B's belief *p2* receives a justificatory boost because he inquired virtuously and the justificatory status of A's belief *p1* is diminished because he inquired viciously. Therefore, my two-component view yields the verdict that B's belief *p2* is much better justified than A's belief *p1*. In contrast to this case, I think number 6 requires more defense. I turn to this now.

Let's look at the Biologist case again as a way to understand number 6. I want to look at two versions of the Biologist case that might be a potential problem for my diagnoses of number 6. In each case, the biologist is in the actual world.

> ***Virtuous* Biologist:** A field biologist is trying to explain a change in the migration patterns of a certain endangered bird species. Collecting and analyzing the relevant data is tedious work and requires a special eye for detail. The biologist is committed to discovering the truth, and so spends long hours in the field gathering data. He remains focused and determined in the face of various obstacles and distractions (e.g., conflicting evidence, bureaucratic roadblocks, inclement weather conditions, boredom, etc.). He picks out important details in environmental reports and makes keen discriminations regarding the composition and trajectory of numerous observed flocks. As a result of his determination and careful methods of inquiry, he uncovers strong evidence supporting a particular hypothesis for why the birds have altered their migration patterns.

Call the reason that the birds have altered course '*p*'. The best explanation available for why the Virtuous Biologist has the evidence he has is *p*. Furthermore, the biologist's belief is well-founded. Thus, along FD, *p* is a belief that is strongly justified for the biologist. Although the biologist's belief is strongly justified along FD, the degree to which *p* is justified is increased because he engaged in virtuous inquiry. Engaging in virtuous inquiry makes it more likely that the evidence he gathered is not misleading (i.e.—veridical) evidence. Thus, engaging in virtuous inquiry increases the degree to which his belief *p* is justified. This increase in epistemic

justification occurs whether or not the biologist is aware of the virtuous nature of his inquiry.

Now let's consider a different case.

> This next case is a different version of the Biologist case above.
>
> ***Vicious* Biologist:** A field biologist is trying to explain a change in the migration patterns of a certain endangered bird species. Collecting and analyzing the relevant data is tedious work and requires a special eye for detail. The biologist wants to discover the truth, but he is lazy and so spends very little time in the field gathering data. He remains unfocused and lacks determination in the face of various obstacles and distractions (e.g., conflicting evidence, bureaucratic roadblocks, inclement weather conditions, boredom, etc.). He fails to pay close attention to important details in environmental reports and makes poor discriminations regarding the composition and trajectory of several observed flocks. Despite his lack of determination and careless methods of inquiry, *he uncovers the same evidence as the Virtuous Biologist.*

In this case, the Vicious Biologist engaged in vicious inquiry. Yet suppose he uncovers the same evidence that the Virtuous Biologist uncovers and forms the same belief *p* on the basis of this evidence. Also, neither biologist is aware of the virtuous or vicious nature of his own inquiry. In this case, both biologists' beliefs are equally justified along FD. However, things are different along RD. Other things being equal (there is no evidence tampering/manipulation or other factor such as mental illness, etc.), belief *p* is justified to a higher degree for the Virtuous Biologist than for the Vicious Biologist because the former inquired virtuously and the latter inquired viciously.

Here's the concern with the above comparison cases and therefore a problem for my view: If the Virtuous Biologist and the Vicious Biologist uncover *precisely* the same evidence in normal circumstances, then the reliability or unreliability of their inquiries cannot make a justificatory difference. That is, the Virtuous Biologist inquired virtuously (reliably) and uncovered evidence *E1* in support of belief *p,* and the Vicious Biologist inquired viciously (unreliably) and also uncovered *E1* in support of belief *p*. On my view, the fact that the Virtuous Biologist inquired virtuously *increases* the degree to which *p* is justified because virtuous inquiry increases the probability that the evidence is veridical. But also on my view, the fact that the

Vicious Biologist inquired viciously *decreases* the degree to which *p* is justified because vicious inquiry decreases the probability that the evidence is veridical. The issue then is how *E1* can be both misleading and veridical at the same time under normal conditions. This seems to be an inconsistency. There is a simple reply to this concern.

First, the claim that the Virtuous and Vicious Biologist would have uncovered the same evidence lacks plausibility. It is implausible that these two biologists could have gathered the same evidence when trying to answer a difficult scientific question having to do with the complexities of the migratory pattern of some bird species. Along FD, the two biologists will have beliefs with very different degrees of justification because they will surely have different evidence. The Vicious Biologist's belief is not likely to be justified whatsoever given that his belief will likely have poor evidential support. But suppose for the sake of argument that the Vicious Biologist does have sufficiently strong evidence to justify his belief along FD. If so, then along RD, the unreliability of the Vicious Biologist's inquiry will increase the probability that his evidence is misleading and therefore the degree to which his belief is justified will decrease. On the other hand, the Virtuous Biologist's belief will be better justified than the Vicious Biologist's belief along both FD and RD. The Virtuous Biologist will have stronger supporting evidence and the reliability of his inquiry will increase the probability that his evidence is veridical and therefore increase the degree to which his belief is justified.

Even in cases where two people are gathering evidence in order to answer *simple* questions, the reliability or unreliability of their evidence gathering methods will often make a justificatory difference. For example, suppose person A and B are inquiring as to whether there is a cookie in the cookie jar. If person A carefully examines the content of the cookie jar and person B gives a very hasty or cursory examination of the contents, it could result in very

different evidence. Person A might perform a careful visual examination of the cookie, feel it in his hands, smell it, attempt to bite and taste the cookie, etc., and in the process discover the cookie is fake (made of plastic). Person B might open the cookie jar briefly and inattentively glance at the cookie and conclude that it is fake. Even two simple but different inquiries can yield different evidence. Person A clearly has stronger supporting evidence than person B. Furthermore, person A's inquiry was highly reliable and therefore the probability that his evidence is veridical is high. On the other hand, person B's inquiry was unreliable (or at best only somewhat reliable) and therefore the probability that his evidence is veridical is much lower than A's.

Second, in cases where two people get precisely the same evidence (where inquiry is involved) it is because they engaged in sufficiently similar inquiries (whether virtuous or vicious) or because the reliability or unreliability of their respective inquiries didn't play a sufficiently strong role in determining what body of evidence they gathered. Either way there is not going to be a justificatory difference between their beliefs along either FD or RD. Thus, my diagnosis of number 6 holds.

My two-component view contrasts with the evidentialist's intuition that if two people are alike mentally (i.e.—have the exact same evidence), then their beliefs are *always* alike justificationally (again, we are concerned with well-foundedness). That is, the evidentialist says that Person A's belief *p* and Person B's belief *p* are equally justified if they have the same body of evidence, regardless of the difference in how they gathered that evidence. On my view, this is not the case. On RD, the degree to which *p* is justified for Person A or Person B, can depend (other things being equal) on whether they engaged in virtuous or vicious inquiry, which in turn means there can be a justificatory difference between their two beliefs even though they have the

exact same evidence. My view also entails that vicious inquiry cannot render a well-founded belief unjustified. This is in contrast to any view of reliabilism which says that if a person engaged in vicious inquiry (i.e.—unreliable evidence-gathering methods), then belief *p* is not justified for that person, regardless of whether *p* is the best explanation available to that person for why he has that body of evidence. What is important to remember is that when considering the justificatory status of a belief, we have to evaluate it along both FD and RD. Of course, the belief must be justified along FD for it to be justified at all, so if it's not justified along FD, then there's no reason to evaluate it along RD.

My view preserves the intuition behind the New Evil Demon Problem that because Person A (actual world) and Person B (demon-world) possess the same strong evidence, their respective beliefs *p* are both justified, but my view adds that the belief is justified *period* (categorically justified along FD). I think any theory of epistemic justification worth its salt, must yield the result that belief *p* for Person B is justified. Regardless of the epistemic condition Person B is in (being deceived by an evil demon), there is no other belief he should rationally hold than what his evidence supports (in the explanationist sense). In this case B ought to believe *p*. What my view does not preserve is the intuition that the justificatory status of A's and B's beliefs are the *same* as long as they have (are either actually or potentially aware of) the same evidence. This is because we have to consider RD. If Person A (who is in the actual world) engaged in virtuous inquiry, while Person B (in the demon-world) also engaged in virtuous inquiry, then there will be a justificatory difference in their respective beliefs even though they have precisely the same body of evidence. Belief *p* for Person A will be justified to a higher degree than for Person B because unlike in the demon-world, engaging in virtuous

inquiry (reliable evidence-gathering) in the actual world is likely to yield veridical evidence. Virtuous inquiry in the demon-world is unreliable.

I'd like to make one more point. There is an important *asymmetry* between the virtuous and vicious inquirers who have the same evidence in the actual world. The asymmetry is that the virtuous thinker used a reliable process to gather evidence and the vicious thinker did not. Where there is a *symmetry* between the virtuous and vicious inquirers in the actual world is in the *availability* of reliable processes to each thinker. They both have the opportunity to utilize reliable evidence-gathering methods. It just so happened that the vicious thinker inquired viciously *but he could have inquired virtuously*. Importantly, this is the advantage that the vicious inquirer in the actual world has over *both* the vicious and the virtuous inquirer in the demon-world. That is, the vicious thinker in the actual world had the opportunity to use a reliable method of inquiry, but failed to take advantage of that opportunity. On the other hand, the thinkers in the demon-world did not have the same opportunity because any process of inquiry they use is unreliable. However, the vicious thinker's belief in the actual world is not thereby better justified because she *could have* inquired virtuously, it's just that he *would have* been better justified *if* he had availed himself of the reliable method. On the other hand, the virtuous thinker in the actual world had the opportunity and *did* use a reliable evidence-gathering method, so this does make a justificatory difference in his resulting belief.

V. Conclusion

The goal of this dissertation has been to defend a two-component view of doxastic justification that makes sense of the role that evidence and virtuous (reliable) inquiry play in fixing the justificatory status of a belief. In particular, the goal has been to defend the view that the justificatory status of a belief can be fixed along two dimensions: 1) the *fittingness-dimension*

which provides necessary and sufficient conditions for categorical doxastic justification and along 2) the *reliability-dimension* which often plays a salient role in increasing the degree to which a belief is justified.

In Chapter Two I distinguished between two different kinds of intellectual virtue properties—roughly, properties of inquirers (e.g., being open-minded) vs. properties of inquiries (e.g., being conducted open-mindedly). I raised objections to theorizing justified belief in terms of "inquirer-focused" properties. I argued that these objections do not foreclose the possibility of theorizing justified belief at least partly in terms of "inquiry-focused" properties. Further, I gave a rough, preliminary statement of an evidentialist view of justified belief in which the exemplification of "inquiry-focused" properties can increase the degree to which a doxastic attitude is justified. In Chapter Three I argued that epistemic justification is not just a species of instrumental rationality. I also argued that methodological epistemic rationality is not merely a form of instrumental rationality but also plays a salient role in fixing the justificatory status of a belief. I also argued that evidentialist views such as Kelly's and Conee and Feldman's provide a plausible account of when a belief is categorically epistemically justified for a person. I also offered a preliminary defense of a diachronic account of epistemic justification. In Chapter Four I defended *explanationist evidentialism* as providing necessary and sufficient conditions for categorical doxastic justification.

Finally, in this final chapter I provided a more thorough defense of my two-component view of epistemic justification. In part I, I discussed how the evidentialist can make sense of how flawed evidence-gathering methods can undermine the justification of a belief. Specifically, I argued that evidentialists can appeal to EFEG as an undercutting defeater. Further, I argued that EFEG is a higher-level belief that serves as an undercutting defeater for first-level (object-

level) beliefs. That is, if one is aware that one's evidence gathering methodology was unreliable, then one has a defeater for the belief that the pertinent evidence bears directly on. I then provided an account of how evidentialists can make sense of defeaters in explanatory terms. Finally, I discussed how NED captures the intuition that two people who are alike mentally will both have the same categorically justified belief. I then made the argument that this outcome allows room for reliability to play a further role in increasing the degree to which a belief is justified.

In part II of this final chapter I briefly argued that EX-WF provides a satisfactory account of one dimension of epistemic justification, i.e.—the fittingness-dimension (FD). I argued that as long as a person's belief is justified along FD, then that belief is justified *period*. The point being that being justified along FD is all that is required for a belief to be justified, i.e.—it is the bare minimum for a belief to be justified.

In part III of this final chapter I argued that beyond the justification conferred by one's evidence along FD, *reliability*, when relevant, can make a positive difference on the justificatory status of one's beliefs. Thus, a person's belief *p* can be justified along FD *and* along the *reliability-dimension* (RD). That is, a person's belief is epistemically justified to varying degrees depending on how strongly that person's belief fits the evidence that person has at a given time, but that in addition, there is an externalist justification factor (namely reliability) that is often directly relevant to the positive justificatory status of that belief. Specifically, I argued that if a person's evidence was gathered in a reliable way, then (other things being equal) the belief that evidence bears on will have an even higher degree of epistemic justification than it would if it were only justified along FD.

The upshot of my two-component theory of epistemic justification is that it provides strong reasons to believe that an evidentialist-cum-reliabilist view of epistemic justification is a tenable one. Even if the reader isn't entirely convinced of my two-component view, I hope that my insights have applications into other areas of epistemology.

There are still important issues that need to be worked out. One issue that needs to be addressed more carefully is the specifics behind how virtues construed as inquiry-focused properties increase the reliability of inquiry. For example, how precisely does conducting your inquiry open-mindedly, consistently, fair-mindedly, etc., contribute to the reliability of inquiry? Additionally, it would be helpful to understand when such virtues of inquiry contribute to a decrease in the reliability of an inquiry and even render an inquiry unreliable. For example, when is conducting an inquiry open-mindedly detrimental to the reliability of that inquiry? Also, are there intellectual virtues that aren't strongly correlated with reliable inquiry whatsoever (e.g., inquiring courageously)?

Another issue to explore is how my two-component view would deal with some of the traditional objections to evidentialism and reliabilism (e.g., the problem of forgotten evidence, the generality problem for reliabilism, etc.). My hope is that many of the problems that plague evidentialists and reliabilists could be fixed by a view in which both evidence and reliability play a salient role in epistemic justification. This would go a long way in strengthening the case for my two-component theory of epistemic justification. Also, I think more work needs to be done in accounting for the connection between best explanation and justified belief. That is, more work needs to be done on understanding the nature of best explanation and how it determines the justificatory status of a belief. Incidentally, this would also be another way to understand the role that virtue plays in epistemic justification. More precisely, such a project would help us

better understand the role that virtues as properties of *explanation* (simplicity, explanatory power, scope, etc.) play in epistemic justification.[85]

Ultimately, my hope is that I've made progress toward developing a two-component or "hybrid" view of epistemic justification and that my insights will motivate others to continue to develop such a view.

[85] Ted Poston (2014) offers a thoughtful discussion on the connection between best explanation and epistemic justification (including insights into the role that explanatory virtues play).

LIST OF REFERENCES

Axtell, Guy. 1997. “Recent Work in Virtue Epistemology.” *American Philosophical Quarterly*, 34: 410–430.

_____ 2011. “Recovering Responsibility.” *Logos & Episteme,* (2) 3: 429-454.

_____ 2012. “(More) Spring of My Discontent: A Reply to Dougherty.” *Logos & Episteme,* 3 (1): 131-137.

Baehr, Jason. 2011. *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford University Press; Reprint edition.

Battaly, Heather. 2008. “Virtue Epistemology.” *Philosophy Compass*, 3: 639–663.

Bergmann, Michael. 2005. “Defeaters and higher-level requirements.” *Philosophical Quarterly,* 55 (220): 419–436.

_____ 2006. *Justification without Awareness*. Oxford University Press.

_____ 2013. “Externalist Justification and the Role of Seemings.” *Philosophical Studies,* 166 (1): 163-184.

Bonjour, Laurence. 1980. “Externalist Theories of Empirical Knowledge.” *Midwest Studies in Philosophy,* 5 (1): 53-73.

Chisolm, Roderick. 1979. “Verstehen: The Epistemological Question.” *Dialectica,* 33: 233-246.

_____ 1982. *The Foundations of Knowing*. Minneapolis: University of Minnesota.

Christensen, David. 2010. “Higher-Order Evidence.” *Philosophy and Phenomenological Research*, 81(1): 185–215.

Code, Lorraine. 1984. “Toward a ‘Responsibilist’ Epistemology.” *Philosophy and Phenomenological Research,* 45 (1): 29–50.

_____ 1987. *Epistemic Responsibility*. Hanover: University Press of New England and Brown University Press.

Cohen, Stewart and Lehrer, Keith. 1983. "Justification, Truth and Coherence." *Synthese,* 55 (2): 191-207.

Comesana, Juan. 2005. “We Are (Almost) All Externalists Now.” *Philosophical Perspectives,* 19: 59-76.

_____ 2010. “Evidentialist Reliabilism.” *Nous,* 44 (4): 571-600.

Conee, Earl. 2004. “The Truth Connection.” In *Evidentialism*, by Earl Conee and Richard Feldman, p.242-58. New York: Oxford University Press.

Conee, Earl and Feldman, Richard. 1985. “Evidentialism.” *Philosophical Studies*, 48: 15-34.

_____ 2004. “Evidentialism.” In *Evidentialism*, by Earl Conee and Richard Feldman, New York: Oxford University Press: 83-107.

_____ 2004a. “Internalism Defended.” In *Evidentialism*, by Earl Conee and Richard Feldman, New York: Oxford University Press: 53-82.

_____ 2008. “Evidence.” In Quentin Smith (ed.), *Epistemology: New Essays*. Oxford University Press: 83-104.

_____ 2011a. "Response to Baehr." In *Evidentialism and its Discontents*. Trent Dougherty (ed). Oxford: Oxford University Press: 310-314.

_____ 2011b. “Response to Goldman.” In *Evidentialism and its Discontents.* Trent Dougherty (ed.). Oxford: Oxford University Press: 302-306.

DeRose, Keith. 2011. “Questioning Evidentialism.” In *Evidentialism and its Discontents*. Trent Dougherty (ed). Oxford: Oxford University Press: 137-146.

Dougherty, Trent. 2010. “Reducing Responsibility: An Evidentialist Account of Epistemic Blame.” *European Journal of Philosophy,* 20 (4): 534-547.

_____ 2011. “Re-Reducing Responsibility: Reply to Axtell.” *Logos & Episteme,* 2 (4): 625-632.

_____ 2012. “Internalist Evidentialism and Epistemic Virtue: Re-Reply to Axtell.” *Logos & Episteme,* 3 (2): pp.281-289.

Engel, Mylan, Jr. 1992. “Personal and Doxastic Justification in Epistemology.” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition ,* 67 (2): 133-150.

Fairweather, Abrol and Zagzebski, Linda. 2001. “Introduction.” In *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility,* Abrol Fairweather and Linda Zagzebski (eds.), Oxford: Oxford University Press: 3-14.

Feldman, Richard. 2000. “The Ethics of Belief.” *Philosophy and Phenomenological Research*, 60: 667–95.

_____ 2003. *Epistemology*. Prentice Hall.

_____ 2004a. “The Ethics of Belief.” In *Evidentialism*, by Earl Conee and Richard Feldman, New York: Oxford University Press: 166-95.

_____ 2004b. “Having Evidence.” In *Evidentialism*, by Earl Conee and Richard Feldman, New York: Oxford University Press: 219-41.

_____ 2005. “Respecting the Evidence.” In *Philosophical Perspectives*, 19 (1): 95–119.

Field, Hartry. 2000. “Apriority as an Evaluative Notion.” In P. Boghossian and C Peacocke (eds.), *New Essays on the A Priori*. New York: Oxford University Press: 117–49.

Foley, Richard. 1987. *The Theory of Epistemic Rationality.* Harvard University Press.

Goldman, Alvin. 1979. “What Is Justified Belief?” In *Justification and Knowledge,* G.S. Pappas (ed.), Dordrecht: Reidel: 1–25

_____ 1992. “Epistemic Folkways and Scientific Epistemology.” *Philosophical Issues,* 3: 271-285.

_____ 2001. “Internalism Exposed.” In *Knowledge, Truth, and Duty,* Matthias Steup (ed.), Oxford: Oxford University Press: 115-133.

_____ 2009. “Internalism, Externalism, and the Architecture of Justification.” *Journal of Philosophy,* 106 (6): 309-338.

_____ 2011. “Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism’s Troubles, Reliabilism’s Rescue Package.” In *Evidentialism and Its Discontents*, Trent Dougherty (ed.), New York: Oxford University Press: 254-280.

Greco, John. 1993. “Virtues and Vices of Virtue Epistemology.” *Canadian Journal of Philosophy*, 23: 413–432.

_____ 1999. "Agent Reliabilism." *Philosophical Perspectives,* 13 (s13): 273-296.

_____ 2000a. *Putting Skeptics in Their Place*. New York: Cambridge University Press.

_____ 2000b. “Two Kinds of Intellectual Virtue.” *Philosophy and Phenomenological Research,* 60 (1): 179-184.

_____ 2003. ``Knowledge as Credit for True Belief.” In *Intellectual Virtue,* Michael DePaul & Linda Zagzebski (eds.), Oxford: Oxford University Press: 111-134.

_____ 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.

Hookway, Christopher. 2000. “Regulating Inquiry: Virtue, Doubt, and Sentiment.” In *Knowledge, Belief, and Inquiry,* Guy Axtell (ed.), Lanham, MD: Rowman and Littlefield: 149-60.

_____ 2003. “How to Be a Virtue Epistemologist.” In *Intellectual Virtue: Perspectives from Ethics and Epistemology,* Michael DePaul and Linda Zagzebski (eds.), Oxford: Oxford University Press: 149-160.

Huemer, Michael. 2007. “Compassionate Phenomenal Conservatism.” *Philosophy and Phenomenological Research,* 74: 30-55.

Kelly, Thomas. 2003. “Instrumental Rationality as Instrumental Rationality: A Critique.” *Philosophy and Phenomenological Research*, 66 (3): 612-640.

_____ 2005. “The Epistemic Significance of Disagreement.” *Oxford Studies in Epistemology,* 1: 167–196.

_____ 2007. “Evidence and Normativity: Reply to Leite.” *Philosophy and Phenomenological Research*, 75 (2): 465-474.

_____ 2008. “Evidence: Fundamental Concepts and the Phenomenal Conception.” *Philosophy Compass*, 3: 933-955.

_____ 2010. “Peer Disagreement and Higher Order Evidence.” In *Disagreement,* Richard Feldman and Ted Warfield (eds.), Oxford: Oxford University Press: 111–174.

_____ 2014. "Evidence", *The Stanford Encyclopedia of Philosophy* (Fall 2014 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2014/entries/evidence/>.

Kornblith, Hilary. 1983. “Justified Belief and Epistemically Responsible Action.” *Philosophical Review*, 92 (1): 33-48.

Kvanvig, Jonathan. 1992. *The Intellectual Virtues and the Life of the Mind*. Savage. Maryland: Rowman and Littlefield.

Lasonen-Aarnio, Maria. 2014. “Higher-Order Evidence and the Limits of Defeat.” *Philosophy and Phenomenological Research*, 88 (2): 314–345.

Lehrer, Keith. 1974. *Knowledge*. Oxford: Oxford University Press.

Leite, Adam. 2007. ”Epistemic Instrumentalism and Reason for Belief: A Reply to Thomas Kelly’s ‘Epistemic Rationality as Instrumental Rationality: A Critique’.” *Philosophy and Phenomenological Research*, 75 (2): 456-464.

Lockhard, Matthew. 2013. ”Epistemic Instrumentalism.” *Synthese,* 190 (9): 1701-1718.

Markie, Peter. 2013. “Searching for True Dogmatism.” In *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism,* Chris Tucker (ed.), New York: Oxford University Press: 248-69.

McCain, Kevin. 2014. *Evidentialism and Epistemic Justification*. Routledge.

Montmarquet, James. 1993. *Epistemic Virtue and Doxastic Responsibility*. Lanham: Rowman and Littlefield.

Plantinga, Alvin. 1993a. *Warrant: The Current Debate*. Oxford: Oxford University Press.

_____ 1993b. *Warrant and Proper Function*. Oxford: Oxford University Press.

Pollock, John. 1986. *Contemporary Theories of Knowledge*. Savage, MD: Rowman and Littlefield.

Poston, Ted. 2014. *Reason and Explanation: A Defense of Explanatory Coherentism*. Palgrave Macmillan.

Rogers, Jason and Matheson, Jonathan. 2011. “Bergmann’s Dilemma: Exit Strategies for Internalists”. *Philosopical Studies: An International Journal for Philosophy in the Analytic Tradition,* 152 (1): 55-80.

Russell, Bruce. 2001. “Epistemic and Moral Duty.” In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, Matthias Steup (ed.), Oxford University Press: 34-38.

Sosa, Ernest. 1985. “The Coherence of Virtue and the Virtue of Coherence: Justification in Epistemology.” *Synthese,* 64: 3–28.

_____ 1991. *Knowledge in Perspective*. Cambridge: Cambridge University Press.

_____ 2007. *A Virtue Epistemology: Volume I: Apt Belief and Reflective Knowledge*. Oxford University Press Uk.

_____ 2015. *Judgment and Agency*. Oxford: Oxford University Press.

Tollhurst, William. 1998. “Seemings.” *American Philosophical Quarterly,* 35: 293-302.

Tucker, Chris. 2013. “Seemings and Justification: An Introduction.” In *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism,* Chris Tucker (ed), New York: Oxford University Press: 1-29.

Turri, John. 2010. “On the Relationship between Propositional and Doxastic Justification.” *Philosophy and Phenomenological Research,* 80 (2): 312-326.

Weslake, Brad. 2011. “Exclusion Excluded.” http://bweslake.s3.amazonaws.com/ research/ papers/welake_exlusion.pdf.

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

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.

_____ 2008a. “Causation and manipulability.” *The Stanford Encyclopedia of Philosophy/(Winter 2008 edition),* edited by E.N. Zalta.

_____ 2008b. “Cause and Explanation in Psychiatry: An Interventionist Perspective.” In *Philosophical Issues in Psychiatry: Explanation, Phenomenology, and Nosology,* Kenneth S. Kendler & Josef Parnas (eds.), Johns Hopkins University Press.

Zagzebski, Linda. 1996. *Virtues of the Mind.* Cambridge: Cambridge University Press.

VITA

Alex Hallam obtained his B.S. in psychology from Colorado State University in 1998. He also earned his M.A. in philosophy from Colorado State University in 2009. Alex worked for Colorado Children's Hospital from 2001-2011 until he moved to Knoxville, TN and attended the University of Tennessee to pursue his Ph.D. in philosophy which he completed in 2016.