

KNOWLEDGE AND INTELLECTUAL SKILL

by

JOSHUE OROZCO

A Dissertation submitted to the  
Graduate School-New Brunswick  
Rutgers, The State University of New Jersey  
in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

written under the direction of

Ernest Sosa

and approved by

---

---

---

---

New Brunswick, New Jersey

October, 2009

## ABSTRACT OF THE DISSERTATION

Knowledge and Intellectual Skill

By JOSHUE OROZCO

Dissertation Director:

Ernest Sosa

This dissertation defends the position that knowledge is best understood as a true belief acquired through the manifestation of intellectually virtuous performance. I argue that intellectually virtuous performance requires intellectual responsibility but not a characteristic motivation. I distinguish my view from other conceptions of intellectual virtues; particularly the virtue reliabilism of Ernest Sosa and John Greco and the virtue responsibilism of Linda Zagzebski. I argue that intellectual virtues are best understood along the lines of Aristotelian skills by looking at various puzzles in epistemology and showing how this view can make more progress in solving these puzzles than its competitors.

## Acknowledgements and Dedication

There are many people to whom I owe many thanks for their continued support and mentorship. I am particularly grateful for the philosophical guidance and instruction that I received over the last five years from the philosophy faculty at Rutgers University; particularly from Ruth Chang, Alvin Goldman, Jeff McMahan, and Dean Zimmerman. I am especially indebted to Ernest Sosa—my supervisor—for his invaluable training and encouragement, and for his remarkable ability to help develop projects and arguments aimed against his own views.

My fellow graduate students, both inside and outside of Rutgers, have had just as big an impact in my growth as a philosopher, providing rich discussion both inside and outside of the seminar room. Geoff Anders, Luvell Anderson, Josh Armstrong, Matthew Benton, Pavel Davydov, Jonathan Ichikawa, Alex Jackson, David Manley, Gary Osmundson, Blake Roeber, Andrew Sepielli, Ryan Wasserman and Stephanie Wykstra have been especially helpful.

Finally, I have received special support from friends and family members who without them this would not have been possible. To my parents who have given me tremendous love and resources. And most importantly to my wonderful wife, Mandy, who has been nothing but supportive and has stood behind me even in the dead of New Jersey winters. I dedicate this to her.

## Table of Contents

Abstract	ii
Acknowledgements and Dedication	iii
Introduction: Knowledge and Intellectual Skill	1
Chapter 1: Epistemic Rationality and Practical/Moral Considerations: The Need for Intellectual Skill	7
Chapter 2: Luck, Lotteries, and Epistemic Responsibility	40
Chapter 3: The Value of Virtuous Behavior: An Answer the Value Problem of Knowledge	72
Chapter 4: Epistemic Responsibility and Doxastic Control	100
Chapter 5: Testimonial Knowledge in Children: I Can Trust You Now ... But Not Later	125
Curriculum Vita	156

## **Introduction: Knowledge and Intellectual Skill**

Gettier's 1963 paper brought the traditional analysis of knowledge—as justified true belief—into question. In response much ink was spilled trying to amend the traditional account of justification. And as the 20<sup>th</sup> century came to a close, some felt an impasse had been reached between competing accounts of justification: in particular between internalist and externalist conceptions of justification.

One characteristic feature of the various modified accounts of justification was their focus on the properties of beliefs. Justified beliefs need to be undefeated, safe, sensitive, rationally entailed by one's evidence, reliably produced, sufficiently coherent, foundationally based, etc. However, in 1980 Ernest Sosa—in “The Raft and the Pyramid”—suggested that focusing on the properties of beliefs was excessive and misplaced, and that we should instead look to the properties of individuals to make epistemic evaluations, in a way suggested by virtue ethics. Sosa proposed that whether a belief qualifies as knowledge depends primarily on the intellectual virtues of the agent. If an agent's intellectual virtues are sufficiently responsible for a belief's truth, then the agent has knowledge.

The set of theories that fall under the sub-discipline “virtue epistemology” center around the thesis that intellectual agents and their “intellectual virtues” are the primary points of epistemic evaluation, with beliefs receiving a derivative evaluation. However, there is serious dispute within virtue epistemology regarding the nature of one's intellectual virtues. If one has knowledge only when the truth of one's belief is credited to one's intellectually virtuous traits, or virtuous performance, then it is important to understand what it is for a trait, or performance, to be intellectually virtuous. For Sosa an

intellectual virtue is a belief forming faculty or process that is reliable at getting at the truth. Things like our perceptual faculties, memory, introspection, and perhaps intuition are considered intellectual virtues. Accordingly, Sosa's virtue-reliabilism is an externalist theory of knowledge—it is not necessary that the conditions for knowledge be accessible to the believer.

In contrast to Sosa, Linda Zagzebski takes intellectual virtues to be perfectly analogous to moral virtues. She argues that one deficiency in the virtue-reliabilist accounts is their inability to accommodate the similarities that epistemology has with ethics. She goes on to defend the view that intellectual virtues are a subspecies of moral virtues. “An intellectual virtue, like a moral virtue,” says Zagzebski “has a motivational component as well as a component of reliable success in reaching the end (if any) of the motivational component. What makes intellectual virtues intellectual is that they (or most of them) include motive dispositions connected with the motive to get truth...”<sup>1</sup> Intellectual virtues are character traits like courage, thoroughness, carefulness, humility, etc., and therefore epistemic evaluations should mirror moral evaluations in virtue ethics ultimately being based on what the intellectually virtuous person would do in a given situation.

The virtue-reliabilism of Sosa and the moral-analogue version of Zagzebski characterize the two main strands of virtue theories in epistemology, and both versions have many attractive features. The reliabilist framework provides many resources for dealing with skepticism and accommodating Gettier-style worries. The moral-analogue framework can easily accommodate parallels between epistemic and ethical evaluations. Our evaluations of beliefs closely resembles our evaluations of actions, and

---

<sup>1</sup> See Zagzebski's introduction in her (2001): 5.

epistemologists ongoing concern over epistemic justification is an example of this resemblance.

My dissertation is work in virtue epistemology. As such, it follows Sosa's suggestion that some epistemic evaluations are best analyzed in terms of properties of individuals—namely her intellectual virtues and/or virtuous activity—rather than properties of beliefs. I focus primarily on evaluations of knowledge and epistemic rationality/justification.

### **Chapter Summaries**

In *Chapters 1-4* I argue that the two main conceptions of intellectual virtues have serious drawbacks, and defend an intermediate understanding of intellectual virtues that retains the attractions of both theories but avoids their shortcomings. I follow Zagzebski in taking intellectual virtues to be more analogous to moral virtues than Sosa's conception. Intellectual virtues *are* character traits like courage, humility, and open-mindedness that require repeated intentional and virtuous behavior for their development. However, I think Zagzebski pushes the analogy too far. Intellectual virtues are not *perfectly* analogous to moral virtues. First, intellectual virtues are not structurally motivational; unlike moral virtues they do not require a characteristic motivation. Second, moral virtues are not necessarily reliable at achieving their aim. One can be benevolent, kind, or generous despite continuously failing to bring about one's intended aim. In contrast, in order for open-mindedness or intellectual courage to be considered intellectual virtues they have to be reliable ways of getting at the truth. Intellectual virtues, on my account, are more akin to Aristotelian skills than to moral virtues. I argue for this intermediate position by looking at some contemporary problems in epistemology.

*Chapter 1* examines the nature of epistemic rationality/justification and its relation to other domains of normativity; particularly, moral/practical norms and evaluations. I argue that although evaluations of epistemic rationality/justification should be distinguished from moral evaluations, it is plausible that moral and practical considerations are relevant to evaluations of epistemic rationality/justification. I argue that an adequate account of epistemic rationality should make use of the notion of intellectually responsible behavior, and that moral and practical considerations are relevant to whether intellectual behavior is responsible. I then amend this account by arguing that epistemic rationality/justification also requires reliability. I propose that epistemic rationality is best conceived as intellectually skillful behavior.

*Chapter 2* looks at the lottery problem, and argues that what explains why one can know an ordinary proposition and not a lottery proposition despite their being formed in a seemingly identical manner (e.g. on the basis of good inductive reasoning), is not issues of reliability but of responsibility. I argue that knowledge requires intellectually virtuous behavior and that intellectually virtuous behavior requires intellectually responsible behavior. The reason we generally fail to know lottery propositions is that in lottery cases we typically violate some norm of intellectual responsibility. I reach this conclusion after examining Sosa's and Greco's virtue-reliabilist accounts, concluding that they are all too impoverished to adequately address the lottery problem.

*Chapter 3* examines the value of knowledge and argues that employing the notions of intellectual virtues and of intellectually virtuous—or skillful—activity *can* provide an adequate explanation of value of knowledge. I first examine a promising first step to solving value problem of knowledge by appealing to intellectual virtues and

credit. This attempt comes from the separate work of Ernest Sosa and Wayne Riggs. I argue that their solution is ultimately inadequate because of an impoverished conception of intellectual virtues. I then attempt to show how a different conception of intellectual virtue and of virtuous activity—one that requires responsible intellectual behavior, and which I call intellectual skill—is better suited to explain the value of knowledge.

*Chapter 4* addresses the connection between intellectual responsibility and doxastic voluntarism. I argue that despite lacking direct voluntary control over our beliefs we do have a significant amount of control to ground attributions of epistemic responsibility. I discuss two main kinds of control that we exhibit over our beliefs. The first is a sort of evaluative control over our beliefs and the second is an indirect control over how we approach our belief forming practices and faculties. I go on to argue that these forms of control are more similar to the sort of control that we exhibit over our actions than might first appear.

*Chapter 5* attempts to make one further application of the previous theory. In particular, I discuss the differences between intellectually skillful behavior in children and in adults. Sandy Goldberg and John Greco have recently tackled the problem of how children acquire testimonial knowledge despite having an indiscriminating character (a kind of local unreliability). Normally, if adults displayed the kind of behavior, and have the kind of intellectual character, as normal children, they would fail to acquire knowledge. So how is it that children come to know things in light of what seems to be unreliable intellectual behavior? I suggest that part of what explains this difference has to do with different standards for intellectually skillful behavior between children and adults. This discussion covers general issues in the epistemology of testimony and

examines how epistemic norms may change depending on an individual's stage in cognitive development.

## **Chapter 1: Epistemic Rationality and Practical/Moral Considerations: The Need for Intellectual Skill**

### **I. Introduction**

Epistemic evaluations strongly resemble evaluations of actions. We require that our beliefs be rational or justified in much the same way we require that our actions be rational or justified. Indeed, many of the same evaluative notions are employed to appraise actions and beliefs. Both can be evaluated as justified or unjustified, right or wrong, permissible or impermissible, obligatory or prohibited, responsible or irresponsible. Moreover, we evaluate *individuals* for their actions and beliefs alike. We judge people as intellectually hasty, overly emotional, stubborn, lazy, dishonest, dogmatic, or cowardly—and they are held responsible in the same way as people who display moral vices. Similarly, we judge people as intellectually thoughtful, objective, open-minded, meticulous, and honest—holding them responsible in the same way as people who display moral virtues.

The purpose of this chapter is to examine the nature of epistemic rationality/justification and its relation to other domains of normativity; in particular, the relationship between epistemic rationality/justification and moral/practical norms and evaluations. Some philosophers have argued that evaluations of epistemic rationality/justification can be reduced into other normative domains—particularly, that epistemic evaluations are a subset of moral evaluations. Others argue that evaluations of epistemic rationality/justification are independent of moral and practical considerations. Indeed, this independence thesis is epistemological orthodoxy, accepted by internalists and externalists alike. I argue that both reduction and independence are problematic. While evaluations of epistemic rationality/justification should be distinguished from

moral evaluations it is plausible that moral and/or practical considerations are relevant to evaluations of epistemic rationality/justification. I argue that an adequate account of epistemic rationality should make use of the notion of intellectually responsible behavior, and that moral and practical considerations are relevant to whether intellectual behavior is responsible. I then amend this account by arguing that epistemic rationality/justification also requires reliability. I propose that epistemic rationality is best conceived as intellectually skillful behavior.

But first a preliminary remark. When one claims that one ought not believe that the Earth is flat one typically makes a claim about *epistemic* rationality/justification. However, even within the epistemic domain there are many different senses of what one ought to believe. For example, it is common to think beliefs aim at truth, and, consequently, to take truth as the fundamental epistemic value. Therefore, it seems natural to evaluate beliefs based on whether or not they achieve their aim. There is a legitimate sense in which one ought to believe the truth. However, I am ultimately interested in knowledge, and hope to illuminate the nature of knowledge by investigating epistemic rationality/justification. For knowledge is also an evaluative notion. When one knows that  $p$  one believes what one ought. This notion of 'ought' goes beyond believing the truth, such that if we believe what we ought, and we reach the truth in virtue of so believing, then we have knowledge. This notion is what I mean to capture by 'epistemic rationality/justification'.

## **II. Distinguishing Epistemic and Moral Evaluations**

When W.K. Clifford claimed “[it] is wrong in all cases to believe on insufficient evidence,”<sup>2</sup> he claimed not only that it is always *epistemically* wrong to believe on insufficient evidence but also that it is always *morally* wrong to believe on insufficient evidence. Failing epistemically entails failing morally. In other words, Clifford took epistemic evaluations as a subclass of moral evaluations. Roderick Chisholm was also open to the “possibility that the epistemic sense of justification can be explicated in purely ethical terms,”<sup>3</sup> and admits his attraction to the view that epistemic justification is a subspecies of ethical justification.<sup>4</sup> Linda Zagzebski, in *Virtues of the Mind*, attempts to “develop ... a virtue theory that is inclusive enough to handle the intellectual as well as the moral virtues within a single theory...;” arguing “that the intellectual virtues are so similar to the moral virtues in Aristotle’s sense of the latter that they ought not to be treated as two different kinds of virtue. Intellectual virtues are, in fact, a kind of moral virtues.”<sup>5</sup>

Clifford, Chisholm, and Zagzebski all seem to suggest the following:

1. Epistemic evaluations are a subclass of ethical/moral evaluations.

However, a proper evaluation of the thesis requires clarifying what is meant by ‘ethical/moral evaluations’. In contemporary moral philosophy the term ‘morality’ refers to two different—though related—domains of evaluation. And how we evaluate 1 may depend on which sense of ‘morality’ one employs.

Under one understanding of ‘morality’ moral evaluations are concerned solely with our duties and responsibilities to other people—for example our duty or

---

<sup>2</sup> Clifford (1874): p. 109.

<sup>3</sup> Chisholm (1980): p. 563.

<sup>4</sup> Chisholm (1991): 119.

<sup>5</sup> Zagzebski (1996): xiv.

responsibility not to harm others, to keep one's promise, to maximize utility, etc.<sup>6</sup>

Clifford often appeals to this kind of moral evaluation and argues that beliefs based on insufficient evidence are morally objectionable, in part, because of the likelihood of such beliefs harming others. Furthermore, Clifford argues that even if no actual harm results we have a responsibility to others not to become intellectually vicious:

“... if I let myself believe anything on insufficient evidence, there may be no great harm done by the mere belief; it may be true after all, or I may never have occasion to exhibit it in outward acts. *But I cannot help doing this great wrong towards Man, that I make myself credulous.*”<sup>7</sup>

Given this understanding of ‘morality’, there are clear cases where epistemic and moral evaluations overlap. Sometimes one is morally required to make sure, to the extent that one's abilities allow, that one's belief fits one's evidence. This is true of Clifford's shipowner who—through self-deception—believes that his ship is seaworthy, and consequently endangers his crew. However, in order to establish that positive/negative epistemic evaluations always correspond with positive/negative moral evaluations one must provide a plausible explanation for those cases in which the two kinds of evaluations seem to come apart—in which epistemic appraisal is positive while moral appraisal is negative, and vice versa.

One might try to explain this divergence by suggesting that epistemic evaluations are *pro tanto* moral evaluations that may be outweighed by *non-epistemic* moral considerations. For example, suppose one has a moral duty to believe that one's spouse is being faithful, even when the preponderance of evidence points to his infidelity. As Susan Haack points out, believing on the basis of the evidence—and thereby receiving positive epistemic appraisal—may be a *pro tanto* moral achievement, which is ultimately

---

<sup>6</sup> Cf. Scanlon (1998): pp. 1-13, 171-77. Williams (1985): pp. 1-21.

<sup>7</sup> Clifford (1874): p. 108 (emphasis added).

outweighed by other moral considerations; in this case, one's obligation to one's spouse.<sup>8</sup> However, it seems unlikely that one will be able to provide a similar explanation for all cases of divergence, since it is implausible that our duties and responsibilities to others apply to all beliefs. Suppose a lone cast away—for reasons of self-preservation—forms the epistemically irrational, and false, belief that he will one day be rescued. It seems implausible that this individual is guilty of some moral infraction—even a *pro tanto* moral infraction—in forming this practically necessary, though epistemically irrational belief.

Clifford suggests that even if some beliefs never *actually* affect another individual, all beliefs are dispositionally related to other individuals. He argues that forming irrational beliefs makes one more susceptible to believing and disseminating falsehoods, and this disposition grounds one's duty to others to uphold one's intellectual character by not believing anything on insufficient evidence.

However, it is not obvious that we always do have a duty—even a *pro tanto* duty—to *others* to maintain an impeccable intellectual character. Unless we have reason to believe that such a character poses some threat to others we typically don't think that anyone's rights have been violated. In the case of the cast away there is no good reason to think that others are threatened in any way by his epistemic irrationality. Consequently, it is implausible to suppose that the resulting belief is a moral failing.

If the scope of the moral evaluations is restricted to our duties and responsibilities to *others*, then it is doubtful that epistemic evaluations are a subclass of moral

---

<sup>8</sup> Haack (2001): pp. 22-3.

evaluations. For it seems that one can be epistemically irrational without violating—or even threatening to violate—the rights of other individuals.<sup>9</sup>

Perhaps appealing to a broader sense of ‘moral evaluations’ can preserve 1. According to this broader notion morality is not just concerned with one’s obligations or responsibilities to others. It is also concerned with answering the more general question, “How ought one live?” A complete answer to this question will make reference to an individual’s character and life as a whole. Under this conception one might be subject to moral criticism even if no interpersonal obligations or responsibilities are flouted—for example, if one fails to have the right motivations for acting, if one fails to develop her skills and talents, or if one is lazy.

Perhaps it is more plausible under this broad construal of ‘morality’ to suppose that epistemic evaluations are a subset of moral evaluations in the sense that epistemic failings are *pro tanto* moral failings. It seems true that, other things being equal, and as much as one can control, one ought to be epistemically rational. But things are not always equal. Sometimes values conflict and it is only possible to realize some, but not all, of the values involved. For example, forming an epistemically rational doxastic attitude might have practical costs. And whether epistemic irrationality is always a moral failing depends on how one understands value conflicts involving epistemic rationality. Clearly,

---

<sup>9</sup> One might suggest that even if believing irrationally or unjustifiedly does not necessarily violate one’s duties to others, it does necessary violate one’s duty to oneself. Perhaps the castaway has a duty to himself to not believe irrationally or maintain an impeccable intellectual character. However, this also seems implausible. For example suppose we follow Scanlon’s contractualist principle that an act (or belief) is wrong if and only if any principle that permitted it would be one that one could reasonably be rejected by others (or oneself). Is it true that any principle permitting an epistemically irrational belief could be reasonably (practically speaking) rejected by others (or oneself)? I don’t think so. At the very least it seems that those proposing reducing epistemic evaluations to moral evaluations should give some reason for thinking that any such principle could be reasonably rejected, but none seems forthcoming. It is no surprise that those defending reduction argue that negative epistemic evaluations entail a negative moral evaluation in the broader sense of morality discussed below, and not in this more narrow sense.

not all cases of value conflict necessitate a *pro tanto* moral failing. In deciding how to get to work I must weigh competing values. I can ride my bike, which is better for the environment, or I can take the car, which would save time. Suppose that, given my options, the best thing for me to do—all things considered—is to ride my bike to work. Even if, other things being equal, I ought to conserve my time, it is implausible to suppose that failing to do so in this case would be a *pro tanto* moral failing.

That said, some think that there are kinds of value conflicts where a *pro tanto* moral failing is unavoidable. These are cases “where there is *decisive* support for two or more incompatible courses of action or inaction.”<sup>10</sup> Consider the following two cases:

2. In good faith and with due caution one makes two solemn and important promises which, as things turn out, conflict.
3. One has to decide either to take part in the killing of one innocent person or to allow many innocent people to be killed.

Some might claim that in cases like these—what I’ll call cases of moral conflict—whatever one decides to do one will be guilty of some moral failing, since for each course of action there seems to be *decisive* reasons that either speak in its favor (as in 2), or count against it (as in 3).<sup>11</sup> It is not important for my purposes whether moral conflicts of this sort really exist, or whether the cases given are examples of such conflict. What is important is if moral conflicts do exist, whether all cases of value conflict involving epistemic rationality qualify as moral conflicts. And this seems doubtful. One thing to note is that the cases typically offered as moral conflicts involve *interpersonal* obligations and responsibilities. In the cases above, we have a duty not to break our promises, a duty not to kill innocent people, and a duty not to allow innocent people to be

---

<sup>10</sup> Nagel (1979): p. 128.

<sup>11</sup> See Stocker (1987).

killed. However, I have already argued that it is unlikely that all cases of epistemic irrationality involve disregarding some obligation or responsibility to other people. If it is essential to cases of moral conflict that they involve conflict between interpersonal obligations and responsibilities, then there is strong reason to doubt that cases of value conflict involving epistemic rationality always involve some *pro tanto* moral failing.

Another way to make this point is by examining Linda Zagzebski's virtue theoretic account of epistemic rationality, and her claim that intellectual virtues are a subclass of moral virtues. Because of life's complexity it is extremely unlikely that any given case of virtuous judgment involves only one sort of consideration. The virtuous life requires some measure of *phronesis*, or practical wisdom, to determine which features in a given situation are most morally salient. Consider the following passage by Julia

Annas:

To be angry in the correct way, based on a correct judgment and not mere feeling, will involve a right grasp of the importance of what it is one is angry about. So getting it right as to how one should act on a particular occasion will involve a correct judgment not merely as to what good-temperedness requires, but as to what temperance, or fairness, requires. So if having the virtue of good-temperedness requires a grasp of the goods achieved in one's life by this virtue, this turns out not to be possible without also having a grasp of the goods achieved by temperance, fairness and so on.<sup>12</sup>

Annas goes on to argue that this sort of reciprocity of virtues is pervasive in moral deliberation. For example, according to Aristotle the characteristic activity of courage involves overcoming one's fears and acting despite threats to one's well-being.<sup>13</sup>

However, a proper manifestation of courage does not require fearlessness. Sometimes it is rash—not courageous—to act in the face of danger; "it is for a noble end that the brave

---

<sup>12</sup> Annas (1993): p. 76.

<sup>13</sup> Cf. Aristotle, NE Book III chapter 6-9.

man endures and acts as courage directs.”<sup>14</sup> A proper manifestation of courage must be receptive to the demands of other virtues and will sometimes require that its characteristic activity not be performed. And this seems true of all moral virtues. Consequently, extending the scope of moral virtues to include intellectual virtues makes the proper manifestation of intellectual virtues *unintuitively* dependent on moral considerations.<sup>15</sup> For it seems that sometimes the characteristic activity of some intellectual virtue will be trumped by other moral considerations. If we suppose that it is *overall* morally virtuous for the cancer patient to believe that she will be healed or for the castaway to believe that he will be rescued (a plausible supposition by my lights), then it will turn out that the proper manifestation of the intellectual virtues in these situations requires believing something despite overwhelming evidence to the contrary. But how can ignoring the evidence be *intellectually* virtuous?

A further problem for assimilating intellectual virtues to moral virtues is their seemingly structural differences. Moral virtues as a class are aimed at attaining the Good, and each moral virtue within the class aims to bring about a particular state of affairs that helps constitute the Good. The aim of intellectual virtues is more theoretical; the general aim being the attainment of truth, or what Zagzebski calls, “cognitive contact with reality.” However, moral virtues unlike intellectual virtues require proper affection or motivation. One cannot act in a morally virtuous manner unless one’s feelings and motivations are properly aligned with one’s aim. However, it seems that one can act in an intellectually virtuous manner without having the proper motivation to attain truth. For example, consider the intellectual virtue open-mindedness. It seems clear that one can be

---

<sup>14</sup> NE 1115b22-24.

<sup>15</sup> I stress “unintuitively” because I go on to argue that moral considerations probably *do* have some relevance for evaluations of epistemic rationality/justification.

open-minded even if one is not motivated by a love of truth. In fact there may be many motivations, some not so noble, for being open-minded. Perhaps being open-minded will help one gain popularity, or perhaps it will help frustrate one's parents who insist their children be like-minded. Perhaps one can't be open-minded if one does not *intend* to reach the truth by fairly considering the merits of differing viewpoints, but one's *motivation* for such intended behavior does not have to be similarly aimed at truth. A pure heart is not needed to be intellectually virtuous.

In summary, it seems that epistemic evaluations should be distinguished from moral evaluations. Although it is clear that there are cases where positive/negative epistemic evaluations correspond to positive/negative moral evaluations, this is not true of all cases of epistemic evaluations.

### **III. Problems for an Independence Thesis**

Given the problems facing 1, one might opt for the following thesis:

4. Evaluations of epistemic rationality/justification are independent of practical and moral considerations.

What one ought to believe will depend on the reasons for believing. Á la Scanlon reasons are considerations that count in favor of a judgment-sensitive attitude.<sup>16</sup> However, there are different kinds of reasons and therefore different kinds of considerations. One kind of consideration that speaks in favor of a belief bears on the question of whether that belief is true. Another kind of consideration that speaks in favor of a belief bears on the question of whether that belief is good to have.<sup>17</sup> These two kinds of considerations can pull in opposite directions—sometimes one kind of reason recommends belief while the

---

<sup>16</sup> Scanlon (1998): p. 17-22.

<sup>17</sup> I am here following Hieronymi (2006): p. 50

other kind of reason recommends either suspension or disbelief. Reasons—or considerations—that bear on the question of whether the belief is true are regarded as *epistemic* reasons for belief. Reasons that bear on the question of whether the belief is good to have are regarded as *non-epistemic* reasons for belief. Accordingly, practical and moral considerations are non-epistemic reasons for belief.

The thought that evaluations of epistemic rationality/justification are independent of practical and moral considerations has been popular among internalists and externalists alike. Internalist versions of independence typically suggest that a belief's epistemic rationality/justification supervenes only on one's evidence.<sup>18</sup> This view—commonly labeled Evidentialism—has *prima facie* appeal. Indeed many find it too obvious to deserve argumentation. The following is Richard Feldman's more precise formulation of Evidentialism:

For any person S, time t, and proposition p, if S has any doxastic attitude at all toward p at t and S's evidence at t supports p, then S epistemically ought to have the attitude toward p supported by S's evidence at t.<sup>19</sup>

And though externalist theories of justification reach beyond what one's evidence supports and appeal to a more general notion of reliability, most seem in agreement with internalists—and Evidentialists—in adopting an independence thesis. The following quotes from Ernest Sosa are representative of how externalists conceive of epistemic justification:

---

<sup>18</sup> There is serious question whether Evidentialism really qualifies as an internalist theory of justification. For my purposes it will not ultimately matter whether Evidentialism is internalist or not. It is enough that internalists typically take Evidentialism to be consistent with their view of rationality/justification and that Evidentialism is a version of the independence thesis. For some references on the relevant debate see Goldman (1999); Steup (2001); and Conee and Feldman, "Internalism Defended" in Conee and Feldman (2004): 53-80.

<sup>19</sup> "Ethics of Belief" in Conee and Feldman (2004): p. 178.

According to epistemic truth monism, truth is the fundamental epistemic value. The epistemic justification of a belief, its epistemically positive status beyond that of being true, is held to involve truth-conducive reliability, however conceived, whether as tracking the truth, or as deriving from a reliable process, or competence, or virtue... Suppose a belief is epistemically justified if and only if it derives from a truth-reliable source, because what matters essentially and distinctively in epistemology is whether and how we are in touch with the truth.<sup>20</sup>

Our subject has been epistemic normativity, a kind of normative status that a belief attains independently of pragmatic concerns such as those of the athlete or hospital patient. Epistemic normativity is a status by having which a true belief constitutes knowledge.<sup>21</sup>

In what follows I present some initial problems for an independent thesis by looking at Evidentialism. I go on to suggest my preferred account of epistemic rationality/justification, and then I turn my attention to reliabilism.

### *III.i Evidentialism and Epistemic Rationality*

Although Feldman's formulation only addresses propositional rationality/justification, Evidentialism can be formulated to address doxastic justification as well:

For any person S, time t, and proposition p, if S has any doxastic attitude at all toward p at t and S's evidence at t supports p, then S epistemically ought to have the attitude toward p supported by S's evidence at t, and appropriately base her doxastic attitude on the evidence.<sup>22</sup>

It seems clear that having this formulation is essential in order to capture the sense of epistemic rationality or justification involved in knowledge.

The viability of Evidentialism as a theory of epistemic rationality/justification depends on what counts as S's evidence, and what it is for evidence to *support* a proposition. Many agree that only one's other beliefs and experiences qualify as one's evidence for a given belief. Feldman claims that "facts which are completely out of one's

---

<sup>20</sup> Sosa (2007): pp. 70-71.

<sup>21</sup> *Ibid.*: pp. 88-89

<sup>22</sup> There are questions regarding what it takes for S to 'appropriately base' her belief on the evidence. I leave these questions concerning the relation between propositional and doxastic justification aside and assume that there is a plausible account available.

cognizance ... are plainly not part of the evidence one has.”<sup>23</sup> If proposition p is evidence for S then S believes p. Additionally, only one’s *justified* beliefs count as evidence. Beliefs formed on the basis of wishful thinking don’t qualify as evidence. There is, however, disagreement about whether there are further restrictions. Timothy Williamson argues that all and only one’s knowledge qualifies as one’s evidence.<sup>24</sup> Richard Feldman argues that only one’s *occurrent* beliefs qualify as one’s evidence.<sup>25</sup> I believe that both of these accounts have manifest problems. For now I assume the least controversial restriction that only one’s experiences and justified beliefs qualify as evidence. I argue that even this most inclusive version of Evidentialism is inadequate as an account of epistemic rationality/justification.

### *III.ii Evidentialism and Cognitive Limitations*

One problem for Evidentialism is that it seems plausible that what one ought to believe depends, in part, on normal human cognitive abilities and limitations. Many restrict one’s evidence to—at least—one’s beliefs because it seems that evidence must—in some sense—be available to an individual in order to have bearing on whether that individual is epistemically justified or rational. Many of our beliefs are evidentially underdetermined, and if we knew more relevant facts many of our beliefs would change. However, being ignorant of relevant facts does not by itself impugn our justification. Some facts are beyond our reach—either cognitively or physically—and therefore never enter into our body of evidence.

Suppose, upon considering whether p, S believes p on the basis of his justified beliefs x, y, and z. Although x, y and z together strongly speak in favor of p there is some

---

<sup>23</sup> Conee and Feldman (2004): p. 226.

<sup>24</sup> Williamson (2000).

<sup>25</sup> Richard Feldman, “Having Evidence” in Conee and Feldman (2004): 219-41.

fact,  $w$ , that, when combined with  $x$ ,  $y$  and  $z$ , supports  $\sim p$ . If—through no fault of his own— $w$  is inaccessible to  $S$ , then surely  $S$  epistemically ought to have based his belief on  $x$ ,  $y$ , and  $z$ . That  $w$  speaks in favor of  $\sim p$  does not, by itself, impugn  $S$ 's justification.<sup>26</sup>

Similarly, suppose when considering whether  $p$   $S$  believes  $p$  on the basis of his justified beliefs  $x$ ,  $y$ , and  $z$ . And though  $x$ ,  $y$ , and  $z$  together sufficiently support  $p$   $S$  has other beliefs  $n_1 \dots n_n$  that when taken together with  $x$ ,  $y$ , and  $z$  support  $\sim p$ . However, if the evidential connection between  $n_1 \dots n_n$  and  $p$  is—through no fault of his own—inaccessible to  $S$  then  $S$  epistemically ought to have based his beliefs on  $x$ ,  $y$ , and  $z$ . That  $n_1 \dots n_n$  support  $\sim p$  does not impugn  $S$ 's justification.<sup>27,28</sup> Just as doxastic justification is not jeopardized if some inaccessible fact in the world tells against our belief, doxastic justification is not jeopardized by some piece of evidence that tells against our belief but whose evidential connection is inaccessible—through no fault of our own—to us.

The above argument presupposes that whether propositions provide evidential support is independent of an individual recognizing an evidential connection, and therefore only speaks against versions of Evidentialism—like Timothy Williamson's—

---

<sup>26</sup> It is not quite right to say that  $w$  speaks in favor of  $\sim p$  since  $w$  is a fact or state of affairs and only propositions stand in an evidential supporting relation. More precisely it is the believed proposition  $\langle w \rangle$  that would speak against  $p$  and in favor of  $\sim p$ . I leave this minor complexity aside since it adds unnecessary verbosity.

<sup>27</sup> Indeed, it seems to me that if  $S$  bases his belief on  $x$ ,  $y$ , and  $z$  and if  $n_1 \dots n_n$  is misleading evidence, then it is possible, in addition to being doxastically justified, for  $S$  to know that  $p$ .

<sup>28</sup> It is important to distinguish two cases. In case 1  $S$  believes  $x$ ,  $y$ ,  $z$ , and  $n_1 \dots n_n$ ; and although  $S$  recognizes that  $n_1 \dots n_n$  has some bearing on the question of whether  $p$ ,  $S$ —and every normal human for that matter—is incapable of seeing whether it supports  $p$  or  $\sim p$ . In case 2  $S$  believes  $x$ ,  $y$ ,  $z$ , and  $n_1 \dots n_n$  but is incapable—through no fault of his own—of recognizing that  $n_1 \dots n_n$  has any bearing on the question of whether  $p$ . I am not arguing that  $S$  is justified in believing  $\sim p$  in both cases. What I am arguing for—and which is sufficient for showing Evidentialism problematic—is that in case 2  $S$  is justified in believing  $\sim p$ . I leave it open whether in case 1  $S$ 's belief that  $n_1 \dots n_n$  has some bearing on the question of whether  $p$  is sufficient to counter-balance  $x$ ,  $y$ , and  $z$ —although I must say I have my doubts that it is. All I need is one case where  $S$  is justified in holding some doxastic attitude despite that attitude not being supported by  $S$ 's total evidence, and case 2 seems to be a clear case of doxastic justification.

that also employ objective supporting relations. Consider Williamson's probabilistic account of evidential support:

EV:  $e$  is evidence for  $h$  for  $S$  if and only if  $S$ 's evidence includes  $e$  and  $P(h/e) > P(h)$ .<sup>29</sup>

What, then, are probabilities on evidence?...The discussion will assume an initial probability distribution  $P$ .  $P$  does not represent actual or hypothetical credences. Rather,  $P$  measures something like the intrinsic plausibility of hypotheses prior to investigation; this notion of intrinsic plausibility can vary in extension between contexts.<sup>30</sup>

The problem I've raised for a Williamsonian brand of Evidentialism is that it's possible that—because of normal cognitive limitations and through no fault of our own—the evidential connection between  $e$  and  $h$  is inaccessible to us. And even though objectively the  $P(h/e) > P(h)$ , we may be justified in not believing  $h$ .

Of course evidentialists can avoid this problem by adopting a subjective supporting relation. Consider the following quote from Feldman and Conee:

There are possible cases in which a person has evidence that implies some proposition, but the connection between that evidence and that consequence is distant and difficult to see. It may be well beyond the talents of the person. I believe that in such cases the person ought not to believe the consequence. Given his failure to see that it is a consequence, to believe it (barring other reasons to believe it) would be rash.... The fact that a person's evidence implies some proposition is not sufficient for the evidence to provide evidential support for the proposition. Roughly, *only those propositions whose connection to the evidence the person apprehends are actually supported by his evidence*. And I think ascertaining this connection is itself an element of the person's evidence.<sup>31</sup>

---

<sup>29</sup> Williamson (2000: p. 187.

<sup>30</sup> Ibid.: p. 211.

<sup>31</sup> "Ethics of Belief" in Conee and Feldman (2001): 181 (emphasis added). This is not an isolated endorsement. The following two quotes repeat their endorsement of a subjective evidential supporting relation.

There is no reason to think that an infinite number of beliefs fit any body of evidence that anyone ever has. The evidence that people have under ordinary circumstances never makes it *evident*, concerning every one of an infinite number of logical consequences of that evidence, that it is a consequence. Thus, believing each consequence will not fit any ordinary evidence. ["Evidentialism" in Conee and Feldman (2001): p. 87]

According to this version of Evidentialism, if S can't recognize the connection between  $n_1 \dots n_n$  and  $\sim p$  then  $n_1 \dots n_n$  does not support  $\sim p$  for S. And given that S does recognize x, y, and z as sufficient evidence for p, then S is doxastically justified in believing p. Unfortunately, this raises new problems for Evidentialism.

Let's distinguish two different interpretations of the thesis Feldman and Conee propose:

ES1: P provides sufficient evidential support for q for S if and only if it seems to S that p sufficiently supports q [and it does not seem to S that there are defeaters for q].

ES2: P provides sufficient evidential support for q for S, if and only if p objectively sufficiently supports q and it seems to S that p sufficiently supports q [and it does not seem to S that there are defeaters for q].

ES1 is more radical than ES2, and clearly makes Evidentialism too permissive. For example it is not uncommon for people to believe that the failure of some event to occur in a random sequence of events supports the proposition that the event will occur in the future. The gambler's fallacy is often employed in test-taking strategies of high school students and college undergraduates. According Evidentialism<sub>ES1</sub> these individuals are doxastically justified in holding their resulting beliefs. In fact very few doxastic attitudes are considered unjustified according Evidentialism<sub>ES1</sub>. Only those doxastic attitudes that are held against an individual's epistemic judgments are deemed unjustified. The sort of doxastic incontinence envisioned is one where an individual has some evidence and

---

A proposition is epistemically justified to someone when *it is evident to the person* that the proposition is true.... Any epistemically support a person has for a proposition is some sort of indication to the person that the proposition is true. ["The Truth Connection" in Conee and Feldman (2001): p. 252]

However, they have not been perfectly consistent. The following seems to be an endorsement of a more objective supporting relation:

But suppose that there were occasions when forming the attitude that best fits a person's evidence was beyond normal cognitive limits. This would still be the attitude *justified* by the person's evidence. If the person had normal abilities, then he would be in the unfortunate position of being unable to do what is justified according to the standard for justification asserted by EJ ["Evidentialism" in Conee and Feldman (2001): p. 87]

apprehends that his evidence supports  $p$ , but fails to believe  $p$ . Some epistemologists doubt that doxastic incontinence is a genuine phenomenon.<sup>32</sup> But even if it is, surely the class of unjustified doxastic attitudes is larger than the class of doxastic incontinent attitudes.

ES2 is more plausible. It states that in order for  $p$  to evidentially support  $q$  for  $S$ ,  $S$  must *correctly* take  $p$  to support  $q$ . If  $p$  fails to objectively support  $q$  then  $p$  fails to evidentially support  $q$  for  $S$ . On the other hand if  $p$  objectively supports  $q$ , then whether  $p$  evidentially supports  $q$  for  $S$  depends on whether  $S$  apprehends such a connection. If  $S$  fails to apprehend that  $p$  supports  $q$ , then  $p$  does not support  $q$  for  $S$ .

Does combining Evidentialism with ES2 (Evidentialism<sub>ES2</sub>) provide an adequate account of doxastic justification? I don't think it does. Any adequate account of doxastic justification must accommodate the following phenomenon: an individual has sufficient evidence and cognitive ability to infer  $q$  from  $p$ , but because of intellectually vicious behavior—e.g. laziness, hastiness, etc.—fails to 'connect the dots', and upon consideration unjustifiably fails to believe  $q$ . However, according to ES2, if  $S$  knows  $p$ , but because of some intellectual vice fails to apprehend that  $p$  is sufficient evidence for  $q$ , then  $p$  does not evidentially support  $q$  for  $S$ . Moreover Evidentialism<sub>ES2</sub> yields that if  $S$  has no other evidence bearing on  $q$ , then  $S$ 's suspending judgment on  $q$  would be doxastically justified—the wrong conclusion. There is a difference between failing to apprehend some evidential connection because the connection is beyond one's cognitive endowments, and failing because of inattention and intellectual sloppiness. Subjective versions of Evidentialism cannot accommodate this difference.

Evidentialism faces a dilemma. If evidential support is objective—à la Williamson—then Evidentialism is overly restrictive since, because of cognitive

---

<sup>32</sup> See Adler (2002)

limitations, it's possible that we ought to have a doxastic attitude not supported by our total evidence. However, if evidential support is subjective—à la Feldman/Conee—then Evidentialism is overly permissive, since sometimes we irresponsibly fail to apprehend an evidential connection.

### *III.iii Intellectual Irresponsibility and Evidence Gathering*

Regardless of one's conception of evidential support, there is another, related, reason why Evidentialism is an inadequate account of epistemic rationality/justification.

Consider the following cases:

**Negligence:** Nancy is suspicious that OJ might have murdered his ex-wife. When investigating the crime scene Nancy finds a glove covered with OJ's blood. However, Nancy is negligent and fails to notice the blood on the glove. Ignoring the blood, Nancy focuses on the fact that the size of the glove is smaller than the size of OJ's hands, and concludes that OJ is innocent. In fact, Nancy got lucky, OJ didn't kill his ex-wife. The blood was strategically placed by the actual killer.

**Forgetful:** Frank is interested in going for a hike. Being a beginner he does not want to embark on a grueling walk. Larry—a good friend who knows Frank's abilities and the relevant hiking trail—tells Frank that the trail to the right of the river is suitable. However, Frank has forgotten that last week Larry lied about another hiking trail. In fact, Frank is in luck: Larry was not lying.

The protagonists in these cases process their evidence perfectly, and appropriately base their beliefs on the evidence. However, there is still something epistemically defective about their doxastic performance, and consequently about their resulting belief. Indeed, although reaching the truth, neither Nancy, nor Frank *knows* what they believe. And their failure to obtain knowledge is the result of not being doxastically justified.

Evidentialists might object that I've conflated two distinct notions of rationality/justification. Diachronic rationality concerns how one should conduct one's rational inquiry in order to satisfy the long-term goal of maximizing truth and minimizing falsehood. Synchronic rationality, by contrast, concerns what one should believe *right*

*now*. They may also claim that only synchronic rationality is relevant for *epistemic* evaluations of rationality/justification. Although admittedly Nancy and Frank are diachronically irrational, they are perfectly synchronically rational.

But now consider one's total evidence for some proposition at a particular time. Thus far I have assumed that one's evidence consists in all and only one's justified beliefs. However, it is unlikely that we often—or ever—immediately bring to mind all of our justified beliefs that have evidential bearing on some proposition under consideration. Indeed, it may often be beyond our cognitive abilities to bring to mind all of our justified beliefs. Typically, only a few of our evidentially relevant justified beliefs immediately come to mind. The rest of our relevant beliefs require some 'digging around'. We weed through irrelevant and unjustified beliefs and bring to mind beliefs that bear on the truth of the proposition considered.

This process of bringing to mind our relevant beliefs takes—to varying degrees—time and energy. Some beliefs require more prodding or prompting than others; depending on things like the content of the proposition and our expertise with the relevant domain. Consequently, bringing to mind our evidence is a diachronic exercise, and therefore we make diachronic judgments about such intellectual behavior. However, if evaluations of diachronic rationality have no bearing on doxastic rationality, then whether one is responsible in bringing to mind possessed evidence has no bearing on doxastic justification. Evidentialism must claim that only one's *occurrent* justified beliefs are relevant for doxastic justification, but this seems unacceptably restrictive. After all consider forgetful Frank, but suppose he *retained* the belief that Larry lied to him last week. If given the slightest bit of thought or prompting this belief would be available.

However, if Frank—because of hastiness—fails to recall that Larry lied, then surely his resulting belief is unjustified. Indeed he seems less justified than in the original forgetting case.<sup>33</sup>

Either considerations of diachronic rationality are relevant for evaluations of doxastic rationality/justification or their not. If they are relevant then considerations of how well we gather evidence are relevant to doxastic justification. This is tantamount to a rejection of Evidentialism. If considerations of diachronic rationality are not relevant then Evidentialism is preserved but becomes unacceptably restrictive.

#### **IV. Epistemic Rationality/Justification and Practical/Moral Considerations**

Evidentialism is inadequate as a theory of epistemic rationality/justification. In the cases above Nancy and Frank fail to believe as they ought and consequently fail to obtain knowledge. I suggest that this sense of ‘ought’ is partially explicated in terms of intellectually responsible behavior. What intellectually responsible behavior requires in a given situation is a complicated issue, and I doubt a general formula exists for making

---

<sup>33</sup> Not surprisingly, Feldman recognizes this conclusion but surprisingly defends it in a paper titled “Having Evidence.” Although I will not go through the details of that paper I will mention one interesting dialectical point. Feldman defends his view that only beliefs one is currently thinking about qualify as available evidence for an individual by showing that there is no epistemically relevant difference between stored but not occurrent beliefs and facts that are not believed. He then argues that if the former make a difference to whether one is doxastically justified then so do the latter. But since the latter clearly don’t make a difference, then the former don’t either.

I agree with Feldman that there is no epistemically relevant difference between stored but not occurrent beliefs and facts not believed, but I use this premise to argue in the opposite direction. I take it as obvious that some of my stored beliefs do affect whether I am doxastically justified. Therefore, I conclude that some facts that are not occurrently believed—and perhaps not believed at all—are also relevant to whether I am doxastically justified. Of course some of my beliefs are buried deep in my consciousness and may be practically inaccessible to me. However, others are much closer to the surface and if relevant to the question I’m considering will have justificatory relevance. Similarly, there are some facts that although have a bearing on the question I’m considering, are inaccessible to me. However, sometimes there are facts that are very easily accessible and depending on the situation might have justificatory relevance to my resulting belief.

I am unsure how to solve this dialectical stand-off with Feldman. Perhaps it comes down to a sort of Moorean evaluation in which case I argue that I am more certain that some stored beliefs have justificatory relevance than that no un-believed fact has justificatory relevance.

such evaluations. However, I've argued that non-evidential considerations are relevant for determining whether one is intellectually responsible: particularly, (i) our cognitive abilities/limitations and (ii) how well/poorly we gather evidence, or bring our evidence to bear on a given question. Questions arise regarding *how much* evidence we should gather, or *how many* of our beliefs we must bring to mind, in a given situation. Although Nancy should have recognized the blood on the glove, and Frank remembered his friend's false testimony, there are limitations to how much gathering they are expected to do. If Larry had lied to Frank 20 years ago, then perhaps not remembering Larry's testimony does not impugn Frank's justificatory status. There comes a point when intellectual inquiry comes to an end and evidential processing begins, but what determines this point? Moral and practical considerations are likely candidates. Perhaps how much investigation is required depends on whether the considered proposition has moral or practical significance.

Jeremy Fantl and Matthew McGrath have offered cases lending support to the claim that the amount of evidence required for doxastic justification shifts depending on the practical significance of the considered proposition.<sup>34</sup> Whether accepting someone's testimony that the train stops in Foxboro is sufficient for doxastic justification depends on things like whether one's career depends on not being late.<sup>35</sup> I believe that similar cases

---

<sup>34</sup> They are not alone in arguing for this position. David Owens (2000) similarly suggests that the needs and interests of the believer may be relevant to whether the believer is justified. John Hawthorne and Jason Stanley (Hawthorne (2004) and Stanley (2005)) defend the view that whether S knows that p will depend on whether  $\sim p$  is a salient possibility for S, and whether it is acceptable to use p as a premise in practical reasoning. Although Hawthorne and Stanley don't focus on epistemic rationality/justification, my arguments against Evidentialism—in addition to the arguments presented by Fantl, McGrath, and Owens—should be taken in the same spirit and further support Hawthorne's and Stanley's general conclusion about pragmatic encroachment.

<sup>35</sup> Fantl and McGrath (2002): p. 67.

support the relevance of moral considerations to evaluations of epistemic rationality/justification.<sup>36</sup>

Train Case 1\*: You're at the train station hoping to get to Foxboro. It doesn't matter to you what time you get there. While waiting on the platform you overhear someone say that the next train stops at Foxboro.

Train Case 2\*: You're at the train station hoping to get to Foxboro. It's important that you get there at 2:00 pm because you promised your friend to be at her show that starts at 2:30 pm. While waiting on the platform you overhear someone say that the next train stops at Foxboro. If you get on the train and he's wrong you're sure to be late and break your promise.

Just like in Fantl's and McGrath's original cases, intuitively you do have sufficient evidence for doxastic justification in case 1\*, but you don't in case 2\*. The moral stakes in case 2\* seem to require that you find evidence corroborating the stranger's testimony.

Here's another case. Sally one day starts acquiring evidence suggesting that her husband is having an affair. At what point has Sally gathered enough evidence to justify belief in her husband's infidelity? If Sally has a moral reason to give her friends and loved ones the 'benefit of doubt', perhaps the amount of evidence Sally needs in order to make suspension or disbelief unjustified is more than were she an indifferent third-party observer.

Simon Keller and Sarah Stroud have recently argued that friendship does require epistemic partiality.<sup>37</sup> They claim that one ought to spend more time and energy investigating evidence that seems to impugn a friend's character, either by attempting to find defeaters for such information, or by constructing and entertaining alternative interpretations of the evidence presented.<sup>38</sup> As Stroud points out, "[this] need not be a

---

<sup>36</sup> Heil (1992) also suggests this possibility.

<sup>37</sup> Keller (2004). Stroud (2006).

<sup>38</sup> Stroud (2006): p. 505-506.

matter of flatly denying the obvious. It is rather a matter of extending more interpretive charity to your friends than you naturally would to strangers—of offering your friends more leeway.”<sup>39</sup>

According to Stroud having these differential epistemic practices as *constitutive* of good friendship helps explain: first, the moral psychology of friendship; second, the intuition that friendship requires that one maintain a favorable opinion of one’s friend’s character; and third the special commitment involved in friendship from which we approach new situations and questions.<sup>40</sup> Keller argues that without such differential practices we may be prevented from a sort of mutual directing and interpreting that is characteristic of good friends.<sup>41</sup> Since I do not have the space to carefully examine their arguments, and since I find their conclusion *prima facie* appealing, I will assume that they have correctly identified a norm of friendship. What’s relevant for my purposes is the relationship that this purported norm of friendship has with the norms of epistemic rationality.

Both Keller and Stroud argue that the two norms conflict: performing differential epistemic practices seems epistemically irrational. Stroud explicitly states,

The beliefs associated with friendship thus appear to contravene the general standards articulated by mainstream epistemological theories.... [They] do not propose different epistemological standards depending on the subject matter at issue. Rather, the canons of epistemically responsible belief formation are content invariable.<sup>42</sup>

However, if what I have said above is correct then the notion of a ‘purely epistemic point of view’ when concerned with knowledge or epistemic rationality/justification is

---

<sup>39</sup> Ibid.: p. 507.

<sup>40</sup> Ibid.: pp. 511-512.

<sup>41</sup> Keller (2004): pp. 339-344.

<sup>42</sup> Stroud (2006): p. 514.

problematic. If moral considerations are relevant for evaluations of epistemic rationality/justification, and if the norms of friendship qualify as moral norms, then perhaps the norms of friendship and the norms of epistemic rationality do not conflict as much as Keller and Stroud suppose. For whether I have responsibly gathered evidence may depend on whether one's friend is concerned. Perhaps epistemic evaluations do permit a sort of partiality.

This does not mean that the norms of friendship never conflict with the norms of epistemic rationality. Perhaps there are cases in which a good friend ought to form epistemically irrational beliefs. I merely suggest that it may be false that insofar as the norms of friendship require epistemic partiality they, *ipso facto*, conflict with the norms of epistemic rationality. The intellectually virtuous performance requires that one recognize the practical or moral import of the proposition in question and adjust the extent of one's investigation accordingly. Furthermore, I am not pretending to offer a precise characterization of how practical and moral considerations affect evaluations of epistemic rationality and justification—I doubt such a characterization is possible.<sup>43</sup>

---

<sup>43</sup> With the recent revival of intellectual virtues to epistemology came a renewed interest in other previously neglected intellectual qualities and states. For example Wayne Riggs (2003) attempts to give an account of the nature of wisdom (and understanding) and identify its role in our intellectual lives. The following is his proposed definition

- (W1) S has wisdom only if
- (i) S has a grasp of the truth about the subjects that are most important, and
  - (ii) S has understanding of these subjects as well.

It should be clear that on this account wisdom is not necessary for knowledge. One can have knowledge without being wise. The person who spends his days memorizing the phone book or counting the blades of grass in his yard fails to focus on those subjects that are most important in our lives and therefore lack wisdom. However, there is a different sense of wisdom that I have been discussing in this chapter that, if I am right, is necessary for knowledge.

- (W2) S has wisdom only if
- (i) S is able to see the moral/practical importance of a given inquiry, and
  - (ii) S can adjust the level of inquiry appropriate for knowledge.

Like (W1), (W2) comes in degrees. However, one can have (W2) without having (W1). I may have a propensity to collect trivia—opting to spend all my time watching ‘Who Wants to be a Millionaire’ instead of CNN—but when presented with a question of great moral or practical relevance, I may still be able to adjust the level of inquiry needed for knowledge.

## V. From Intellectual Responsibility to Intellectual Skill

### *V.i Reliability*

Up to this point I've argued that an adequate account of epistemic rationality/justification must employ the notion of intellectual responsibility. However, this can't be the whole story, since there are clear cases of intellectually responsible individuals who, because of some cognitive malfunction or deficiency, have wildly unreliable belief forming practices, and consequently lack knowledge.<sup>44</sup> William Alston gives an example of a primitive tribesman who faultlessly follows the unreliable epistemic traditions of his tribe.

The foregoing discussion of intellectual responsibility can be couched in a virtue theoretic framework. One cannot be skillful if one is not successful in bringing about one's intended aim. If our intellectual aim is truth, then one cannot be *intellectually* skillful unless the methods employed are truth conducive. Aristotle states, "a craft [or skill] ... is essentially a certain state involving reason *concerned with production*."<sup>45</sup> Consider a baseball analogy: suppose my wife is perfectly responsible in her approach to hitting, giving undivided attention to the pitch, and trying her best to follow the appropriate instructions. Unless she employs reliable methods—unless she keeps her weight back, her front elbow in, hands loose and head down, she will not be acting skillfully. Similarly, as intellectual agents we can be perfectly responsible in our belief forming processes, but unless those methods are truth-conducive, we will not be intellectually skillful. Understanding intellectual virtues along the lines of intellectual skills accommodates the epistemic desideratum that our belief forming processes be reliable. Intellectually virtuous (or skillful) activity requires responsible *and* reliable

---

<sup>44</sup> See Plantinga (1988): pp. 3-12 and Alston (1988): pp. 284-294.

<sup>45</sup> NE 1140a6-11: emphasis added.

belief formation. I suggest that my account of intellectually virtuous activity does provide the resources for an adequate theory of knowledge. Knowledge of *p* is had when one arrives at the truth because of one's intellectually virtuous (or skillful) activity.

The need for reliability raises further problems for those claiming that intellectual virtues are structurally identical to moral virtues. Morally virtuous activity is judged to a large extent by the intentions and motivations of the individual and not by the success of the intended action. Julia Annas claims that success is part of the morally virtuous life but that that sort of success involves having the right motivations:

Virtue ethics is concerned with the person's life as a whole, with character and the kind of person you are. The right perspective on an action, therefore, will for virtue ethics be the one which asks about success in achieving the overall goal, rather than success in achieving the immediate target. *What matters is what the person's motivation was, and how this relates to her developed character and life as a whole; for this is her achievement, what she has made of her life. To the extent that success in achieving the immediate target depends on factors over which the person has no control—moral luck of various kinds—it will be of less interest to virtue ethics.*<sup>46</sup>

Take benevolence for example. Consider a man who—out of love for his fellow men and women—intends to bring about good to others, and responsibly gives his money to organizations that have a good reputation. If through no fault of his own his efforts systematically fail to achieve their aim he should still be judged as benevolent and generous. By conceiving of intellectual virtues as skills, rather than structurally identical to moral virtues, we avoid the problem of reliability.

### *V.ii Responsibility Revisited*

Before moving on, more precision is needed in characterizing what qualifies as intellectually responsible behavior. Suppose Sally, in the face of all the evidence, tries her hardest to believe that her husband is faithful. She constructs implausible

---

<sup>46</sup> Annas (2003): p. 25 (emphasis added).

interpretations of the evidence and continues to irrationally suspend judgment with the hopes of finding defeaters. One day, however, the evidence becomes overwhelming—perhaps she catches her husband kissing another woman. As a result, she comes to know that her husband is having an affair. It seems that although acquiring knowledge Sally failed to be intellectually responsible. Of course at the time of belief formation she is being intellectually responsible—she trusts her eyesight, which she has good reason to believe delivers the truth. But her intellectual performance across time is abysmal.

In the last section I argued that because Nancy and Frank are diachronically irrational—although perfectly synchronically rational—they fail to be doxastically justified. It might appear that I am in the awkward position of having to make a similar evaluation regarding Sally's belief. However, there is an important difference between Nancy and Frank, and Sally. In Sally's case although she is diachronically irrational, there is a sense in which her final doxastic attitude is not held in virtue of her irresponsibility. By contrast, Nancy's and Frank's final doxastic attitude *is* held in virtue of their irresponsibility. What best explains Sally's belief in her husband's infidelity is that she saw him kissing another woman, whereas what best explains Nancy's belief that O.J. is innocent and Frank's belief that the trial is suitable is their failing to take into account relevant and easily available evidence. Therefore, what I argue in section IV is that even if one is perfectly synchronically rational, if one's doxastic attitude is held in virtue of being diachronically irresponsible then one fails to be doxastically rational/justified, and is precluded from acquiring knowledge.

The claim that the epistemic standing of one's doxastic attitude depends on whether it is being held in virtue of one's irresponsible performance is admittedly

imprecise. However, I do think that a similar phenomenon can be found in moral evaluations of actions. Consider the following case:

Thomas is wondering what to do with his money. After leaving his ethics class he begins to wonder whether he is obligated to give some of it to Oxfam. When he gets home Thomas tries to push the arguments he heard in class out of his mind, and turns on the television. After channel surfing for a couple of minutes he lands on a commercial for Oxfam, and sees first hand the suffering and pain that his money could help alleviate. Convicted, Thomas immediately picks up the phone and donates his money.

Thomas clearly behaves irresponsibly by trying to ignore seemingly strong arguments for supporting Oxfam. And we can suppose that were it not for his irresponsible behavior he never would have seen the commercial. However, there is still a sense in which his donation is not *in virtue* of his being irresponsible. Thomas still deserves credit for doing the right thing. I take this to be structurally analogous to the epistemic case involving Sally. Despite behaving intellectually irresponsibly Sally still deserves credit for believing as she ought, and as a result acquires knowledge. Therefore, knowledge requires intellectual responsibility in the sense that one's belief cannot be formed in virtue of intellectually irresponsible behavior.

## **VI. Reliabilism and Moral/Practical Considerations**

Given that I've included reliability in my account of epistemic rationality/justification, some might wonder whether I was too quick in concluding from my arguments against Evidentialism that moral or practical considerations are relevant to evaluations of epistemic rationality/justification. Recall that practical/moral considerations were marshaled in to help determine how much evidence one must gather or how much evidence one needs to bring to mind for doxastic justification. However, perhaps reliability can do this job. Whether or not one's intellectual behavior confers doxastic

justification depends on whether such behavior is sufficiently reliable to produce true, rather than false, beliefs.

Unfortunately, appealing to reliability does not provide an adequate solution to our questions about intellectual investigation. For we are still left with the question, what sort of intellectual behavior qualifies as 'sufficiently reliable'. It is widely accepted that although doxastic justification does not require perfect reliability it does require that one's intellectual performance be more likely to reach truth than falsehood. But how much more likely? The level of reliability seems to change depending on context and content. Intuitively, believing on the basis of someone else's testimony is sufficiently reliable in train case 1\*, but not in train case 2\*. We ask for second opinions when the doctor's diagnosis is life-threatening but not when it is a minor cold. Or consider the general phenomenon of gathering evidence or bringing evidence to bear on a question. For most of us the more time and energy we spend investigating some question the more likely we are to get at the truth. But again we don't require perfect reliability so what determines when our investigations should end?

One might suggest that a performance is sufficiently reliable if its likelihood of success is sufficiently close to the likelihood of performances of normal individuals in similar situations. Consider another sports analogy. The average professional basketball player makes about 45% of his shots. That means in order to be credited for a successful basketball performance a professional player's likelihood of success cannot be significantly lower than the mean. However, suppose that the reason players typically shoot only 45% is because there is a 24 second shot clock. This means they have to take a shot within 24 seconds of gaining possession otherwise the other team gets the ball. If

given an indefinite amount of time, one would be able to wear down one's opponent and have a higher shooting percentage. Of course, there are other times in the game where a 45% shot will not do. Shots directly after a time-out are expected to be more reliable than 45%. One shouldn't settle for a jump shot but instead look for something closer to the basket, where the likelihood for success is higher. Therefore, different features of the situation will affect the relevant level of reliability for a given shot. The shot-clock lowers the reliability threshold, while time-outs raise the threshold. I suspect that when we turn our attention to our intellectual performances we will find that relevant practical and moral situational features will have an analogous affect on setting the reliability threshold for beliefs. Even if the likelihood of acquiring true beliefs is proportionate with the length of investigation, we are not expected to engage in an exhaustive investigation for every question considered. There are other pressures in life that require we end our intellectual investigations and take a doxastic shot. This means that the reliability threshold is lower than it would be if we could investigate the question indefinitely—but such is life. However, there are other situations that require a more exhaustive investigation. Because we are beings with moral and practical sensitivities these situations will typically be ones where the practical and moral stakes are higher than normal. As I noted above, we tend to look for more evidence when our well-being or the satisfaction of our moral duties depend on the truth of the proposition considered. And as a result the reliability threshold is raised in these situations.

The challenge for reliabilists is this: if we are fallibilist with respect to epistemic rationality/justification and knowledge then there should be some principled way—however imprecise—of setting the reliability threshold. Even if we look to the normal

practices and performances of normal human beings as the standard, we are still faced with the undeniable fact that moral and practical considerations affect our intellectual practices. Consequently, there will not be *one* standard to which all of our intellectual performances are measured. Rather, it seems likely that the standard will shift depending on the moral or practical features of the situation. Of course, one way to avoid this fact, and keep epistemic evaluations independent from moral and practical considerations, is by requiring perfect reliability. But this move will be especially unattractive for those hoping to avoid widespread skepticism.

## **VII. Conclusion**

I've argued that the relationship between epistemic rationality and other normative domains is not as clean as some have supposed. Although evaluations of epistemic rationality/justification should be distinguished from moral evaluations, it is unlikely that evaluations of epistemic rationality/justification will be completely independent of moral and practical considerations. I've suggested that the sense of epistemic rationality/justification most closely connected to knowledge should employ the notion of intellectual responsibility, and should be understood specifically as intellectually skillful performance. Whether one's intellectual performance is skillful requires that it be sufficiently reliable and responsible, and this requires sensitivity to moral and practical considerations. Before closing I'll briefly discuss one more connection between my understanding of epistemic rationality/justification (understood along the lines of intellectually skillful performance) and moral evaluations.

Many ethical theories make a distinction between objective and subjective moral evaluations. An action's objective moral evaluation is typically taken to be independent of the agent's beliefs regarding the rightness or wrongness of an action. Whether or not

an action is *objectively* right (or wrong) depends on the facts. If an action is the best (or a sufficiently good option) of the available alternatives then it receives positive objective moral evaluation. On the other hand, an action's *subjective* moral evaluation is dependent on the agent's beliefs, or at least on what the agent has good reason to believe. If an agent performs an action for which he has good reason to believe is best (or sufficiently good), then the action receives a positive subjective moral evaluation. Note that an action's subjective and objective moral evaluations may come apart. S may have good reason to believe that action p is the best action, when in fact q, and not p, is the best action. Accordingly, S's act of p will receive positive subjective moral evaluation but negative objective moral evaluation. This distinction highlights that one may not always be blameworthy for performing an objectively wrong action, and that one may not always be praiseworthy for performing an objectively right action.

Feldman has argued that epistemic evaluations do not accommodate the distinction between objective and subjective evaluations in the way that moral evaluations do.<sup>47</sup> I do not have the space to examine his arguments. However, I will note that my account of intellectually virtuous activity does seem to fit nicely with this feature of ethical evaluations. To repeat, I've suggested that intellectually virtuous behavior requires not only that one be intellectually responsible, but also that intellectual performance be reliable. In other words, in order to exhibit intellectually virtuous (or skillful) behavior one must receive positive subjective and objective epistemic evaluations. One may have good reasons to think that trusting the community oracle is a reliable method for attaining the truth. However, if consulting the oracle is an unreliable guide for truth, then such behavior will receive positive subjective epistemic evaluation

---

<sup>47</sup> Feldman (1988).

but negative objective epistemic evaluation, and, consequently, fall short of being intellectually virtuous (or skillful).

## Chapter 2: Luck, Lotteries, and Epistemic Responsibility

### I. Introduction: Intellectual Virtues and Epistemic Luck

It is widely accepted that knowledge precludes epistemic luck. What is deeply contested—and has been since Gettier’s famous paper—is the type of luck incompatible with knowledge. Even within the subfield of virtue epistemology there is major disagreement regarding what kind of luck is pernicious to knowing. Virtue epistemologists agree that the locus of epistemic evaluations should primarily be the properties of persons rather than the properties of beliefs. Key epistemic notions, like knowledge and justification, are analyzed in terms of the intellectual virtues of *individuals*. However, there is serious disagreement regarding the proper understanding of intellectual virtues. On one side are those who define intellectual virtues along reliabilist lines. An intellectual virtue is a stable disposition within an intellectual agent to reliably form true beliefs. Our most deeply entrenched belief forming processes such as perception, memory, introspection, and intuitions are paradigm examples of intellectual virtues, since when appropriately situated they are disposed to yield true beliefs.<sup>48</sup> On the other side are those who define intellectual virtues analogous to a classical, or Aristotelian understanding of moral virtues; they are character traits resulting from habitual intellectual agency or practice. Here emphasis is placed on intellectually responsible behavior—displaying intellectual wisdom in belief formation by being appropriately courageous, open-minded, careful, humble, etc.

---

<sup>48</sup> Ernest Sosa claims that intellectual virtues are “psychological mechanism(s) that would deliver a high enough preponderance of true beliefs.” See Sosa (2003): p. 163. Similarly, John Greco states that “[a]n intellectual virtue or faculty...is a power or ability or competence to arrive at truths in a particular field and to avoid believing falsehoods in that field. Examples of human intellectual virtues are sight, hearing, introspection, memory, deduction and induction.” See Greco (1992): p. 520.

Not surprisingly, this difference in characterizing intellectual virtues is closely correlated with the issue of epistemic luck—each group emphasizing a different kind of luck as incompatible with knowledge. Virtue reliabilist focus on what I will call *veritic epistemic luck*.<sup>49</sup> One way to understand this form of epistemic luck is as follows:

For any belief forming process R, S's true belief that p is the result of veritic luck to the extent that reaching the truth is unlikely, given that the belief was produced by R.<sup>50</sup>

Since intellectual virtues are characterized as stable dispositions that reliably produce true beliefs, then an epistemically lucky belief—in the above sense—is intellectually vicious.

Virtue responsibilists focus on epistemic norms governing our practices as intellectual agents, censuring those beliefs formed by practices violating such norms. And because intellectual virtues are understood as character traits analogous to classical moral virtues, the norms that are taken to govern our intellectual behavior extend to our motives. In order to be epistemically responsible one's belief has to be subjectively appropriate. As Guy Axtell points out, it isn't enough for our beliefs to be in conformity with our epistemic norms, they also need to be held "*from or in light of* such norms."<sup>51</sup> In chapter 1 I argued that intellectual virtues do not require proper motivation in the way that moral virtues do. Instead a responsibilist need only require that in order for a true belief to count as knowledge, an agent's intellectual performance must have appropriately

---

<sup>49</sup> Although I follow Mylan Engle and Duncan Pritchard in using this term to capture this notion of epistemic luck, my characterization is borrowed from Wayne Riggs (1998). See Engle (2002), and Pritchard (2003). Although Engle and Pritchard both present different characterizations of epistemic luck Riggs's formulation is sufficiently similar for my present purposes. I will address Pritchard's account of epistemic luck later in this chapter.

<sup>50</sup> Riggs (1998): p. 463. This characterization will have to be altered to account for Sosa's view that a belief qualifies as knowledge if true because of an intellectual competence seated in the intellectual agent. For Sosa, a belief need not be safe in the way specified above in order to be apt (or competent), and aptness is all that is necessary for knowledge. However, given that aptness is determined in part by a true belief formed by a reliable disposition, the general point still applies that a belief p is veritic lucky if it is unreliably true, whether that be because the belief is unsafe, or because it is not the result of an epistemic competence.

<sup>51</sup> Axtell (2001): p. 165.

contributed to an agent's belief formation by being intellectually responsible and reliable. Including a notion of epistemic responsibility in one's account of virtuous performance suggests a different way that a belief can be epistemically lucky from the veritic account above. According to what I'll call subjective luck:

S's true belief that p is the result of subjective luck to the extent that S's belief that p is not responsibly formed.<sup>52</sup>

My main aim in this chapter is to extend the thoughts of the first chapter and further motivate an account of knowledge that takes subjective luck seriously by adopting a responsibilist understanding of intellectual virtues. This motivation will be indirect: via an examination of the lottery problem. I will examine some attempts within virtue epistemology to solve the lottery problem and suggest that all are unsatisfactory. I argue that their failure results from mistakenly addressing only veritic luck, and failing to give adequate attention to subjective luck. I suggest that their myopia results from holding too tightly to a reliabilist understanding of intellectual virtues. I argue that by shifting our focus toward subjective luck, and adopting an account of knowledge that incorporates intellectual virtues in the responsibilist mold, the lottery problem becomes more tractable. This does not show that attempts to deal with veritic luck are otiose, or even that veritic luck should be ignored. As I made clear in Chapter 1, my view of intellectual virtues aims to eliminate this form of luck. However, in this chapter I give a reason to think that a condition of epistemic responsibility is necessary for knowledge.<sup>53</sup>

## II. The Lottery Problem

---

<sup>52</sup> See Riggs (1998): p. 467, for a similar understanding of subjective luck.

<sup>53</sup> I should note that the scope in this chapter is restricted to only virtue epistemology. I do not address proposed solutions to the lottery problem that fall outside the scope of virtue epistemology. The standard contextualist responses of Cohen, DeRose, Lewis, and others are not mentioned. I'm making the modest claim, that virtue reliabilists are unable—and virtue responsibilists are able—to adequately deal with the lottery problem. I make no claims about the relative advantages (disadvantages) of contextualist accounts.

Consider the following valid argument:

1. S knows that S won't have enough money to go on a safari this year.
2. If S knows that S won't have enough money to go on a safari this year, then S is in a position to know that S will not win a major prize in a lottery this year.
3. Hence, S is in a position to know that S will not win a major prize in a lottery this year.<sup>54</sup>

Most agree that, given the right background conditions that can be uncontroversially stipulated, 1 is true. If S is not extremely wealthy, doesn't have a large income, has no rich relatives who are on their deathbed and about to bequeath him a large sum of money, etc., then S knows that he will not have enough money to go on a safari this year. 2 is also very plausible. *Pace* Dretske and Nozick, most epistemologists find some formulation of the closure principle attractive.<sup>55</sup> However, given 1 and 2, 3 seems logically entailed. But according to Hawthorne, "were (S) to announce that he knew that he would not win a major prize in a lottery this year, we would be less inclined to accept his judgment as true. We do not suppose that people know in advance of a lottery drawing whether they will win or lose."<sup>56</sup> Moreover, in light of accepting closure and rejecting 3, we are now in the uncomfortable position of needing to reject 1, but 1 is intuitively plausible. Therefore, recognizing that one fails to know that one's lottery ticket will lose seems to preclude all beliefs whose content are entailed by such a "lottery proposition"<sup>57</sup> from qualifying as knowledge.

---

<sup>54</sup> Taken from Hawthorne (2004); pp. 2-3.

<sup>55</sup> This is definitely true of single-premise closure principles. Multi-premise closure is more controversial. However, the formulation of the lottery problem above only assumes the truth of the more widely accepted single-premise principle.

<sup>56</sup> *Ibid.* p. 2. It is important to note that Hawthorne here draws a connection between knowledge and the norm of assertion. I will come back to this later, but for now I simply focus on his intuition about whether the belief qualifies as knowledge.

<sup>57</sup> The term "lottery proposition" was introduced by Vogel (1989).

Things get worse. Our inability to know “lottery propositions” threatens the epistemic status of all our pre-theoretically justified, but fallible, beliefs. Consider the following two cases of inductive reasoning.<sup>58</sup>

Case 1. On the way to the elevator, S drops a trash bag down the garbage chute of her apartment building. A few minutes later, reasoning on the basis of past experience and relevant background knowledge, S forms the true belief that the bag is in the basement garbage room. Of course her grounds for so believing are merely inductive: it is possible that the trash bag somehow gets hung up in the chute, although this is extremely unlikely.

Case 2. S buys a ticket for a lottery in which the chances of winning are ten million to one. A few minutes later, reasoning on the basis of past experience and relevant background knowledge, S forms the true belief that she will lose the lottery. Of course her grounds for so believing are merely inductive: it is possible that she buys the winning ticket, although this is extremely unlikely.

In both cases S uses excellent inductive reasoning. However, intuitively S knows that her bag is in the garbage room, but not that her lottery ticket will lose. Something has to give. Unless one is willing to either concede to skepticism or betray one’s intuitions, something should explain why our evaluations of S differ in the two cases.

In the next section I will survey three virtue theoretic attempts to deal with the lottery problem. I argue that all three fail to adequately account for our conflicting intuitions. The three views I consider are: Duncan Pritchard’s safety reliabilism, John Greco’s agent reliabilism, and Ernest Sosa’s performance reliabilism. Following this evaluation, I discuss how a responsibilist virtue account is better suited to address the lottery problem.

### *III.i Attempted Solution: Pritchard’s Safety Reliabilism*

---

<sup>58</sup> Taken from Greco (2003): p. 112. Greco takes Case 1 from Sosa (2000).

It is somewhat misleading to include Pritchard's account of safety reliabilism into the class of virtue reliabilist theories. Pritchard argues that a safety condition is necessary in account of knowledge, but explicitly rejects virtue theoretic accounts of knowledge.<sup>59</sup> However, many virtue theories, especially early formulations, make use of the safety condition.<sup>60</sup> Since Pritchard's account is, by his own lights, consistent with such accounts, I use Pritchard's attempted solution of the lottery problem as an available approach to virtue epistemologists characterizing intellectually virtues in terms of a safety condition.

Pritchard clearly connects the lottery problem with epistemic luck. He states that "the moral of [the lottery] cases seems to be that what is primarily important to knowledge possession is not the quality of one's evidence in favour of one's belief, but more specifically whether or not one's belief is epistemically lucky."<sup>61</sup> For Pritchard, the relevant form of epistemic luck in play in the lottery problem is veritic luck, and the lesson of the lottery problem is that lottery beliefs fails to meet the following safety condition:

S's true belief is safe iff in most near-by worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, and in all very close near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, the belief continues to be true.<sup>62</sup>

The reason why in Case 2 S does not know that her lottery ticket will lose is that there are some 'very close near-by' possible worlds in which S forms the belief that her lottery ticket will lose in the same way that she does in the actual world, but her belief is false.

---

<sup>59</sup> He claims that once a safety condition is admitted in one's account the motivation for a virtue theoretic account is eliminated. See Pritchard (2003).

<sup>60</sup> See Sosa (1999), (2000) and (2002).

<sup>61</sup> Pritchard, Duncan (2007): p. 9 (manuscript pagination)

<sup>62</sup> Ibid.: p. 12. (manuscript pagination)

The challenge remains for Pritchard, and for any virtue epistemologist hoping to reply to the lottery problem along these lines, to explain what constitutes a ‘very close near-by’ world. After all, isn’t it also true in Case 1 that in some ‘very close near-by’ world S’s trash bag gets caught in the chute, making her belief that her bag is in the garbage room false? Pritchard claims that the way Case 1 is specified, the world in which the bag gets caught in the trash chute is sufficiently different from the actual world to remove it from the class of ‘very near-by worlds.’ In order to include this world into the relevant class the case would have to be restated such that the possibility of the bag getting caught is significantly increased—“for example, [if] there is a snag in the chute that the bag is *almost* snagging on each time.”<sup>63</sup> In contrast, the world in which S’s lottery ticket wins is sufficiently similar to the actual world to be considered a ‘very close’ world—“[after] all, the possible world in which I win the lottery is a world just like this one, where all that need be different is that a few coloured balls fall in a slightly different configuration.”<sup>64</sup> According to Pritchard, the world in which S wins the lottery is sufficiently similar to the actual world to preclude S from knowing that her lottery ticket will lose.

However, distinguishing the lottery case from the trash chute case on account of S’s belief being unsafe in the former, but safe in the latter, seems unsatisfactory and misguided. My main complaint is that his descriptions of what qualifies as a very close near-by world for each case seem unacceptably ad hoc. For example, consider a trash chute that has a snag, making it possible, though extremely unlikely, that a garbage bag gets caught on the way down. Pritchard claims that unless trash bags are always *almost*

---

<sup>63</sup> Ibid.: p. 11. (manuscript pagination)

<sup>64</sup> Ibid.: p. 12. (manuscript pagination)

getting snagged, the world in which a trash bag gets snagged is not sufficiently similar to be a ‘very close near-by’ world. But why think that? It seems that Pritchard would not make the analogous claim that unless S’s lottery ticket is always *almost* the winning ticket, the world in which S’s ticket wins is not a ‘very close near-by’ world. Surely her ticket isn’t always almost the winning ticket—if she plays frequently enough, there will be plenty of days when she gets only one, or no numbers, drawn. Consistency requires Pritchard to claim that S’s lottery belief is safe, and qualifies as knowledge.

Pritchard exaggerates the similarities between the world in which S’s lottery ticket loses, and the world in which it wins. He states “all that need be different is that *a few* coloured balls fall in a slightly different configuration.” But if the lottery involves more than a million possible outcomes (as is usually the case), then in some nearby worlds in order for S’s ticket to have won, much more would be required than a slight change in direction of a few balls. Furthermore, even if we grant Pritchard this description of the lottery case, there seems no justification for withholding a similar description in the trash chute case—all that need be different is a slightly different angle and velocity of the trash bag.

Even if the lottery is an extremely complex process, we typically don’t think one can know in advance of the drawing whether one’s ticket will lose. However, many of one’s other beliefs qualify as knowledge despite being equally or more likely to be false than the lottery belief.

If we are to look at Pritchard’s approach through virtue theoretic lenses, then intellectual virtues are the belief forming processes that yield safe beliefs. I, however, have argued that merely appealing to safety does not adequately distinguish the lottery

case from other cases of fallible knowledge. Therefore, if veritic luck is understood in terms of a safety condition then giving an account of knowledge that addresses only veritic luck will not solve the lottery problem. With regards to this understanding of veritic luck S's belief that the trash bag is in the garbage room and S's belief that her lottery ticket will lose are on a par. If one maintains that in the lottery case, S's belief fails to qualify as knowledge because the belief is epistemically lucky, then either one will have to give a different account of veritic luck or one will have to recognize a different type of epistemic luck in play.

### *III.ii Attempted Solution: Greco's Agent Reliabilism*

John Greco, like Pritchard, has a modal understanding of veritic luck. Our beliefs are subject to veritic luck to the extent that such beliefs would turn out false in nearby possible worlds where the actual belief forming process is employed. Therefore, all of our fallible beliefs are subject to some degree of veritic luck. But Greco claims that knowledge is not inconsistent with veritic luck, as illustrated in the garbage chute case. Only some forms of veritic luck preclude knowledge, as illustrated in the lottery case. How do we distinguish the pernicious from the innocuous forms? Pritchard's attempt is unsatisfying. Greco appeals to the idea that knowledge requires credit for true belief to explain the distinction.

In order for S to know that p, it is not enough that her belief is formed from an intellectually virtuous disposition. Her belief that p must be true *because* of her intellectual virtue(s). Only then does S know that p, since only then does she receive credit for the true belief. Greco argues that the difference between pernicious and innocuous veritic luck is that the former, but not the latter, precludes credit for true belief.

At first, this move seems no more promising than Pritchard's attempted solution. We can stipulate that in the lottery case S uses great inductive reasoning, perhaps even greater than in the garbage chute case, to arrive at her true belief that her lottery ticket will lose. But even with this stipulation, the intuition that S does not know remains unaffected. Why would S be credited for her true belief in the garbage chute case, and not in the lottery case, when the reasoning employed in the lottery case is just as good as—and perhaps better than—her reasoning in the garbage chute case?

To answer this question Greco turns his attention to our practices of credit attribution. Drawing from Joel Feinberg's work on moral blaming, Greco offers the following general theory of credit:

- A person S deserves credit of kind K for action A only if:
- a. A has value of kind K,
  - b. A can be ascribed to S, and
  - c. A reveals S's K-relevant character. Alternatively: S's K-relevant character is the most important necessary part of the total set of causal factors that give rise to S doing A.<sup>65</sup>

Clearly much more needs to be said about what makes one's character the *most important* part of the causal story. Greco does not attempt to give a detailed list of all the factors determining when something is an *important* or a *salient* feature in a set of causal factors. He does, however, provide two factors, again borrowed from Feinberg, that he takes to be relevant for determining a proper causal explanation. First, a proper causal explanation will often pick out causal features that are abnormal or contrary to our expectations. Greco offers the following example: "we will say that sparks caused the fire if the

---

<sup>65</sup> Greco (2003): p. 121. In this formulation S's K-relevant character is only *an* important part, not *the most* important part, of the total set of causal factors. However, later in the article Greco adjusts to the stronger version in order to avoid making Gettier cases instances of knowledge. See pp. 127-32.

presence of sparks in the area is not normal.”<sup>66</sup> Second, our interests and purposes often pick out the important causal features. Greco claims that there are various legitimate explanations for New York’s drop in crime. Some explanations highlight legislative policies, others highlight good policing, and others might highlight the welfare of the national economy. Which explanation one adopts will partially depend on one’s interests and purposes.

Greco then applies this general account of credit to intellectual credit:

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 character is the most important necessary part of the total set of causal factors that give rise to S’s believing the truth regarding p.

Greco claims that in the lottery case—but not in the garbage chute case—S does not receive intellectual credit for her true belief. Only in the former is S’s reliable character the most important necessary part of the causal factors explaining why S’s belief is true.

He gives the following explanation for this difference:

In both cases ... S’s abilities make up a necessary part, but only a part, of the whole story regarding S’s believing the truth. [However] it is only in the garbage chute case that S’s abilities are a salient part of the story. In the lottery case, what is most salient is the element of chance. Why does the element of chance become salient in the lottery case? I would suggest that the very idea of a lottery has the idea of chance built right into it.<sup>67</sup>

Elsewhere he repeats this thought:

The very idea of a lottery involves the idea of chance, and so we have an explanation why chance is salient in cases where the lottery is salient. We can then apply a familiar general principle of credit attribution: that salient chance undermines credit.<sup>68</sup>

---

<sup>66</sup> Ibid.: p. 118.

<sup>67</sup> Ibid.: p. 124.

<sup>68</sup> Ibid.: p. 127.

Unfortunately, Greco does not explicitly connect these claims with his earlier discussion of how abnormal causal features and our interests and desires affect salience. One possible connection may be that the presence of a lottery is an abnormal causal feature, and this, perhaps along with our interests and purposes, makes it the most salient feature. Any correct explanation for why S arrived at the truth does not focus on S's intellectual character, but on the lottery and the chance it involves. Therefore, S cannot be properly credited for her true belief in the lottery case. Her belief is true, not because of her intellectual character, but because of the chance involved in the lottery. Given that Greco takes intellectual credit as a necessary condition for knowledge, he seems to preserve the intuition that S does not know that her ticket will lose.

Although I agree with Greco that knowledge requires intellectual credit, I think analyzing credit in terms of “the content and pragmatics of causal explanations”<sup>69</sup> is problematic. Even if his additional condition on knowledge comports with our intuition in the lottery case, it has many unwelcome consequences. Specifically, there are many cases in which an individual does not deserve the sort of intellectual credit that Greco defines for her true belief, though her belief intuitively qualifies as knowledge.

Take the following example from Jennifer Lackey's paper “Why We Don't Deserve Credit for Everything We Know”:

Case 3: Having just arrived at the train station in Chicago, Morris wishes to obtain directions to the Sears Tower. He looks around, approaches the first adult passerby that he sees, and asks how to get to his desired destination. The passerby, who happens to be a Chicago resident who knows the city extraordinarily well, provides Morris with impeccable directions to the Sears Tower by telling him that it is located two blocks east of the train station. Morris unhesitatingly forms the corresponding true belief.<sup>70</sup>

---

<sup>69</sup> Ibid.: p. 132.

<sup>70</sup> Lackey (2006): p. 352.

We may even add that the passerby has in the past taken extraordinary care in figuring out the directions from the train station to the Sears Tower. Lackey correctly points out that according to the guidelines that Greco presents for proper attribution of intellectual credit, Morris does not deserve credit for his true belief. After all, the most salient part of the explanation for why Morris' belief is true is not Morris' intellectual character, but the intellectual character of the passerby. This case is problematic for Greco because it describes a paradigmatic case of testimonial knowledge. However, according to Greco's analysis of intellectual credit Morris can't be given intellectual credit for his true belief. Either we reject the claim that Morris' belief qualifies as knowledge or we reject Greco's analysis.

Furthermore, Lackey's point regarding Greco's understanding of intellectual credit generalizes beyond testimonial knowledge.<sup>71</sup> Consider the following two cases:

Case 4: For the past 100 years scientists have thought that bird A and bird B sing the same bird song. Their songs are quick and somewhat faint, and even recordings suggested they are same. Last Tuesday, however, Jason heard bird B sing its song much slower, louder, and clearer than usual. To Jason's surprise he discovers that B's bird song is slightly different from A's, and forms the corresponding true belief.

Case 5: For years Bob has admired the painting in his house. He can almost perfectly describe the character of the painting. However, last Saturday while at Home Depot, Bob bought a state of the art light bulb that shines brighter than all other bulbs. Upon installing the new bulb in his living-room Bob discovers that his favorite painting has a subtle shade of blue which he had heretofore not recognized, and forms the corresponding true belief.

In both cases the beliefs of Jason and Bob seem to qualify as knowledge. However, according to Greco's analysis neither Jason's nor Bob's intellectual character is the most salient feature. The most salient feature in each case is the abnormally high quality of

---

<sup>71</sup> Lackey thinks the case above shows that any attempt to distinguish knowledge from non-knowledge true belief by appealing to intellectual credit is misguided. I am more optimistic. Her cases only suggest that an account of intellectual credit cannot be analyzed in the way Greco suggests.

evidence they receive for their true belief. Without such evidence they would have remained ignorant. And since the correct explanation for their true beliefs does not involve their intellectual character, neither deserves intellectual credit—on Greco’s account—for reaching the truth. Greco seems forced to make the unintuitive judgment that Jason and Bob fail to know.

Greco attempts to account for these cases in the following passage:

Not all abnormalities undermine the salience of cognitive character, however. For example, an unlikely coincidence reminds the detective of evidence he has neglected, and this missing piece of the puzzle allows him to solve the crime... We would like an account of which abnormalities undermine the salience of character in the explanation of true belief and which do not. Here is a suggestion...: abnormalities in the way one *gets* one’s evidence do not undermine credit, whereas abnormalities regarding the way one gets a true belief, given that one has the evidence that one does, do undermine credit. Put another way, in cases where something unusual does take away the salience of character, it seems just a matter of good luck that S ends up with a true belief, *even given that she has the evidence that she does.*<sup>72</sup>

First, it is not clear to me how this amendment is not ad hoc. If intellectual credit is analyzed in terms of the correct explanation for why S’s belief is true, then surely there are cases where the correct explanation focuses on the way S gets her evidence for her belief, and not on S’s intellectual character. Greco simply states that these explanations do not undermine credit. But why not? There does not seem to be any justification forthcoming.

Second, in the above examples what is abnormal is not *the way* that Jason or Bob get their evidence. What is abnormal is *the quality* of the evidence. For Jason the bird sang slower, louder, and clearer than before. For Bob the lighting was significantly better than it had previously been. Perhaps, Greco will also preclude such forms of salience from undermining intellectual credit, but again I see no principled reason for doing so.

---

<sup>72</sup> Greco (2003): p. 131.

Third, there are still clear cases of knowledge where, because of “an abnormality in the way one gets a true belief, given that one has the evidence that one does,” the most salient feature is not an individual’s intellectual character. Consider the following case:

Case 6: Jason is taking an introductory forensics class. He and the rest of the students are given case studies that present them with a set of evidence. Their job is to determine how the victim of some crime died. Unfortunately for Jason, his peers are uninterested in the class and spend all their time talking about their social lives. This is such a distraction to Jason that he is never able to give the evidence adequate attention to make the correct inference and form the true belief. He finds himself constantly uncertain about how the victim in the given case died. Last Friday, however, Jason’s classmates—because of a late party the night before—were particularly tired and quiet. As a result Jason is able to carefully examine the evidence and form the correct belief about how the victim died.

In this case the most salient feature explaining why Jason formed a true belief is not Jason’s intellectual character but rather his classmates’ abnormally somnolent behavior. Moreover, this abnormality concerns, not how Jason gets his evidence, but how he forms his true belief given the evidence he possesses. This means that according to Greco’s account—even with the added amendment—Jason does not know how the victim in Thursday’s case study died. However, this seems like the wrong conclusion.

Finally, Greco’s recognition that our interests and purposes affect what qualifies as a legitimate explanation presents further complications to his view. He suggests that what is most salient often varies depending on one’s interests and purposes:

Another major factor governing salience is our interests and purposes. For example, often when we are looking for something’s cause we are looking for something that we can manipulate to good effects.... [Witness] the various explanations of New York City’s plunging crime rate. The police attribute it to good policing, the mayor attributes it to broader social policy, and opposing politicians attribute it to things over which the mayor has no control, such as the upturn in the national economy. Of course...[different] people have different interests, and so highlight different parts of the causal conditions that were together sufficient for the drop in crime.<sup>73</sup>

---

<sup>73</sup> Greco (2003): p. 119.

But now suppose that Jason has just had his philosophy class and formed the true belief that the J-T-B account of knowledge is false. Jason was attentive in class and understood all of the reasoning behind the refutation. But Jason's teacher was also in rare form explaining the material with unusual skill and clarity. Furthermore, Jason's peers were exceptionally quiet, not creating any of the usual distractions. Surely Jason's true belief qualifies as knowledge. However, Greco seems committed to allow that the most salient feature of the story may depend on the various interests in play. Perhaps according to the professor's interests the most salient feature was his masterful explanations, or perhaps according to the interests of Jason's peers the most salient feature was their comatose behavior during class. In these contexts the correct explanation for why Jason formed his true belief will not involve Jason's intellectual character, therefore, Jason will not be credited for his true belief, and his true belief will not qualify as knowledge—again, an unacceptable consequence.

Greco acknowledges that with regards to veritic luck the lottery case is on equal footing with the garbage chute case. The belief forming processes that formed S's true beliefs are equally reliable. This makes any distinction between knowledge and non-knowledge true beliefs solely on the basis of veritic luck too permissive. Attempting to be more restrictive, Greco appeals to a form of intellectual credit that is analyzed in terms of the content and pragmatics of causal explanations. Not only must those dispositions that constitute S's intellectual character, and which formed S's true belief, be reliable, but they must also be the most salient part of the explanation for why S's belief is true. Unfortunately, this amendment makes Greco's account counterintuitive.

Greco may settle for a weaker analysis of intellectual credit in which all that is needed is that one's intellectual character is *a* relevantly salient feature, rather than *the most* salient feature in the total set of causal factors. However this analysis will not be able to distinguish between the lottery case and the garbage chute case since in both the inductive reasoning employed is excellent. The reasoning in the lottery case resembles nothing like a lucky guess. In fact it may even be more skillful than that employed in the garbage chute case.

### *III.iii Sosa's Performance Reliabilism*

Ernest Sosa, in his most recent virtue theoretic formulation, rejects safety as a condition on knowledge. Instead Sosa claims that a true belief qualifies as knowledge only if it is apt. And a belief is apt if and only if its being true is properly credited to intellectual competences exercised by the agent. Although appealing to credit resembles Greco's account, Sosa does not analyze intellectual credit in terms of the pragmatics of causal explanations. In fact, Sosa does not give any detailed analysis of what is required for intellectual credit; he simply gives an intuitive picture of what credit, or aptness, amounts to. Consider his discussion of the various assessments one can give to an archer:

First, we can assess whether it succeeds in its aim, in hitting the target....  
 Second, we can assess whether it is adroit, whether it manifests skill on the part of the archer... A shot can be both accurate and adroit, however, without being a success creditable to its author. Take a shot that in normal conditions would have hit the bull's eye. The wind may be abnormally strong, and just strong enough to divert the arrow so that, in conditions thereafter normal, it would miss the target altogether. However, shifting winds may next guide it gently to the bull's-eye after all. The shot is then accurate and adroit, but not accurate because adroit (not sufficiently). So it is not apt, and not creditable to the archer.<sup>74</sup>

---

<sup>74</sup> Sosa (2007): p. 22

An agent's beliefs, like an archer's shot, can be assessed according to its accuracy, adroitness, and aptness. Knowledge requires that the true belief be apt; true because adroit (or intellectually skillful).

Sosa endorses a reliabilist understanding of intellectual virtues—they are belief forming processes that reliably yield true beliefs under appropriate conditions. If a true belief is formed by one of these processes then it is considered adroit, and if it is true in virtue of being so formed then it is apt, and qualifies as knowledge.

Although Sosa does not explicitly apply his theory of intellectual credit and aptness to the lottery problem, it is instructive to see whether an explanation to the lottery problem can be made on his behalf. And given the characterization above, it seems Sosa's theory is ill-equipped to provide an adequate explanation. It is stipulated that the inductive reasoning employed in the lottery case is no less reliable—and therefore, no less adroit—than in the garbage chute case. Furthermore, in both cases S's belief seems apt. For example, there are no Gettier-style conditions suggesting that the belief's truth is not properly credited to S's intellectual virtues. Sosa's theory seems committed to yielding that if S's belief in the garbage chute case qualifies as knowledge, then so does S's belief in the lottery case.

However, this does not exhaust Sosa's resources. Sosa distinguishes between animal knowledge and reflective knowledge. Of the two types of knowledge an individual can attain, reflective knowledge is the more demanding. Animal knowledge requires "apt belief without requiring defensibly apt belief, i.e., apt belief that the subject aptly believes to be apt, and whose aptness the subject can therefore defend against relevant skeptical doubts." By contrast, reflective knowledge requires "not only apt belief

but also defensibly apt belief.”<sup>75</sup> In other words, animal knowledge requires that one’s true belief be credited to one’s reliable cognitive processes, but not that one also believe that one’s belief was aptly formed. In contrast, reflective knowledge requires that “one’s judgment or belief manifests not only such direct response to the fact known but also understanding of its place in a wider whole that includes one’s belief and knowledge of it and how these come about.”<sup>76</sup> Reflective knowledge requires—well—reflection, and the second-order belief that one’s first-order belief was formed in a truth-conducive way.

Can this help with the lottery problem? Sosa might concede that in the garbage chute case *and* in the lottery case S’s belief is apt, and qualifies as animal knowledge, but maintain that *only* in the garbage chute case can S’s belief reach the level of reflective knowledge. Perhaps there is some feature in the lottery case, absent in the garbage chute case, that precludes S from taking a reflective position toward her apt belief. But what might this feature be? Consider the following kaleidoscope case:

Case 7: Upon entering a room S sees a red surface in good lighting and forms the corresponding belief that the surface is red. However, unbeknownst to S the surface is a kaleidoscope surface that is being controlled by a jokester. The jokester could have very easily set up the room so that a white surface appeared red because of the lighting. In fact, if S had come in slightly earlier or later than she did, her true belief would have been mistaken.<sup>77</sup>

Sosa denies that knowledge requires safety and grants that in Case 7 S knows that the surface is red. Even though S’s conditions could have easily gone bad, her perceptual faculties in fact were operating under suitable conditions and S’s belief was aptly formed. However, Sosa also can accommodate the intuition that S belief does not qualify as knowledge. In this case the presence of the jokester prevents S from being

---

<sup>75</sup> Ibid.: p. 24.

<sup>76</sup> Sosa (1991): p. 292.

<sup>77</sup> This case is taken from Sosa’s (2007) discussion of animal and reflective knowledge: pp. 92-112.

reflectively situated to her belief. Although she aptly believes that the surface is red, she cannot aptly believe that her belief is apt. Any second-order belief regarding the reliability of her first-order belief will be grossly unreliable. The presence of the jokester, and the fragility of S's surroundings, destroys S's prospects for reflective knowledge.

Unfortunately, nothing analogous to the jokester seems operative in the lottery case. The conditions under which S employs her inductive reasoning are extremely stable. And there is nothing about the lottery case that prevents S from recognizing this fact. Therefore, if Sosa is willing to concede that S has animal knowledge in the lottery case, there seems no reason to deny the possibility of S having reflective knowledge as well. There is nothing in the case that precludes S from aptly believing that her belief that her ticket will lose is apt. I am doubtful that appealing to reflective knowledge will help Sosa account for the difference between the garbage chute case and the lottery case.

### *III.iv Summary*

I have considered three different possible solutions to the lottery problem, and have found them all inadequate. The accounts have substantive differences, but they all are committed to virtue reliabilism, and understanding intellectual virtues as stable cognitive processes that reliably lead one to the truth. This conception of intellectual virtues seems motivated by the thought that knowledge requires the elimination (or significant reduction) of veritic luck. Pritchard avoids veritic luck by suggesting that knowledge requires safety. Greco strengthens the safety condition by also requiring that the reliable cognitive processes are the most salient feature in explaining why one's belief

is true. Sosa replaces safety with aptness, but still claims that in order to have knowledge one's true belief be true *because* of one's reliable cognitive processes.

I agree that veritic luck is incompatible with knowledge. However, I think that there is another type of epistemic luck, what I have called subjective luck, which is also incompatible with knowledge. Recall that S's true belief that p is the result of subjective luck to the extent that belief that p is not responsibly formed. The three views I have considered do not address this type of epistemic luck, and I take this as deficiencies in their theories, highlighted by their shortcomings in attempting to answer the lottery problem. I suggest that by focusing this form of epistemic luck the lottery problem becomes more tractable.

In what follows I discuss how an account of epistemic responsibility deals with subjective luck and might help with the lottery problem. If successful, this should help motivate having a condition of epistemic responsibility in one's analysis of knowledge.

#### **IV. Epistemic Responsibility**

I have already mentioned that virtue responsibilists differ from virtue reliabilists by understanding intellectual virtues as analogous to moral virtues. In Chapter 1 I argued that this analogy is not perfect. First, intellectual virtues are not a subclass of moral virtues. Although there may be cases where moral and epistemic evaluations overlap, there are cases where the two come apart. Second, intellectual virtues do not require that an individual be properly motivated in the way that moral virtues do. One cannot be morally courageous without one's affections and motives being aligned with the aim of one's intended action. However, no such alignment is necessary for intellectual courage. Third, intellectual virtues have a success component: one cannot be

intellectually virtuous if one is not reliable at reaching the truth, whereas one can be morally virtuous without being reliable in achieving one's aim. What's essential for moral virtues is that one's intentions and motivations are properly oriented, not that one is successful in bringing about what one's intentions and motives aim at. These differences aside, intellectual virtues are character traits—like intellectual courage, humility, open-mindedness, etc.—for which we are responsible and which are developed through intentional intellectual behavior aimed at acquiring truth and avoiding falsehood. Finally I argued that one can behave responsibly with regards to one's intellectual inquiries and development and fail to behave virtuously (or skillfully). This happens when one's behavior is unreliable at yielding true beliefs. In order to behave intellectually virtuously one must behave responsibly and reliably.

I'll adopt Sosa's formulation that a true belief qualifies as knowledge only if it is apt. But I take a belief as apt if and only if it is true because of the intellectually virtuous behavior of the agent. Requiring that virtuous behavior reliably yield true beliefs precludes that one's belief is veritically lucky. However, since focusing on veritic luck has been shown unfruitful in addressing the lottery problem, I suggest we look to the feature of virtuous behavior that eliminates subjective luck—namely, epistemic responsibility. If the aim of intellectual inquiry is to acquire true beliefs avoid false beliefs, then one behaves in an intellectually responsible manner only if one engages in the intellectual practices that one has good reason to believe will reliably yield true beliefs and avoid false beliefs. But before this can be useful in addressing the lottery problem I will need to discuss some factors that help determine which intellectual practices are intellectually responsible in a given situation. Although I do not intend on

providing any sort of algorithm—indeed, I doubt that such an algorithm is possible—I discuss some principles that place responsibilist restrictions on our intellectual practices.

It is widely recognized that morally virtuous action is not the product of mechanical rule following. Virtuous, and responsible, behavior requires sensitivity to various situational features. Aristotle's *doctrine of the mean* highlights this point. According to Aristotle, the virtuous action falls intermediate between two vices, one of excess and one of deficiency. But this intermediate act should not be seen simply as a midpoint between two distinctive act types. In certain situations the virtuous act will resemble one extreme and in other situations it will resemble the other extreme. Terrence Irwin notes:

Aristotle warns us not to suppose that his appeal to a mean is intended to offer a precise, quantitative test for virtuous action that we can really apply to particular cases—as though, e.g., we could decide that there is a proper, moderate degree of anger to be displayed in all conditions, or in all conditions of a certain precisely described type... [Correct] reason may require extreme anger at extreme injuries and slight anger at trivial offenses; in both cases moderate anger would be wrong. To find the mean relative to us is to find the state of character that correct reason requires, neither suppressing nor totally indulging nonrational desires.<sup>78</sup>

And here Irwin seems to be drawing from statements of Aristotle like the following:

... it is not easy to define the way we should be angry, with whom, about what, for how long. For sometimes, indeed, we ourselves praise deficient people and call them mild, and sometimes praise quarrelsome people and call them manly... [In] every case the intermediate state is praised, but we must sometimes incline toward the excess, and sometimes toward the deficiency; for that is the easiest way to hit the intermediate state.<sup>79</sup>

Acting morally virtuously depends on various contingent features of one's situation; things like the capacities and abilities of the individual, the available alternatives open to

---

<sup>78</sup> See Aristotle (1999): p. 197.

<sup>79</sup> Ibid.: pp. 29-30 (1109b20-25).

the individual and the various relationships that an individual has with the other individuals involved.

Similarly, the requirements for intellectually virtuous behavior will depend on various contingent features of a situation. For example, in some situations intellectual courage may resemble dogmatism, in other situations it may resemble fickleness. Sometimes different virtues may appear in conflict and wisdom is needed to determine how to balance the various values. The main point is that the type, and extent, of the intellectual practices required for intellectually virtuous behavior will also depend on various contingencies. For example, it may depend on the content of the proposition considered. Intellectually virtuous behavior with regards to a simple perceptual belief will generally be radically different than virtuous behavior for a belief concerning some philosophical theory. More care, caution, thoroughness, etc., will be required for the latter. With most of our perceptual beliefs, virtuous behavior intellectual trust of one's perceptual faculties suffices.

Furthermore, our cognitive abilities along with our other doxastic states will help determine how much intellectual labor is required. A normally functioning adult can responsibly form the belief that there are three items on a surface simply by looking at the surface, and trusting one's appearance. Not so when the number of items on the surface is sufficiently large. I would need to do more intellectual labor to responsibly form the true belief that there are twenty objects on the table. Similarly, how much work I have to do with regards to responsibly forming a true belief will often depend on the other beliefs I happen to hold. For example, the amount of intellectual labor I must do to responsibly

form the belief that Tom stole the book from the library will depend on whether or not I believe Tom has a twin brother.

Given that our epistemic standards are sensitive to the above features, in order to perform epistemically responsibly one must adjust one's intellectual labor accordingly whenever one has good reason to believe that such features are present. If one fails to adjust then one fails to believe responsibly and—given the responsibilist understanding of intellectual virtues—one's belief is not virtuously formed and can't qualify as knowledge. I suggest that what distinguishes the garbage chute case (and other normal cases of fallible knowledge) from the lottery case is that generally, in the garbage chute case S's belief is responsibly formed, while in the lottery case S's belief is not responsibly formed. In both cases S's true belief is formed on the basis of excellent inductive reasoning, but only in the garbage chute case does such reasoning satisfy the demands for responsible behavior.

Why is S's inductive reasoning insufficient for responsible belief formation in the lottery case? Consider the following normative principle governing our intellectual behavior:

(Anti-Luck): Generally, one should not form a belief if one thinks that the reason why that belief is true is a matter of luck.<sup>80</sup>

If responsible behavior requires that one engage in those intellectual practices that one has good reason to think will reliably yield true beliefs, and if in a given situation one thinks that despite one's intellectual behavior the truth of one's belief will be the result of luck, then in that situation one's intellectual behavior is insufficient for responsible belief

---

<sup>80</sup> One might wonder what kind of luck is relevant here. I think that an Anti-Luck normative principle is plausible for both forms of luck under consideration in this paper. However, the form of luck that I will consider in the lottery case is veritic luck. Therefore, if one thinks that the reason why one's belief is true is because of veritic luck then one should not form that belief.

formation. This normative principle helps explain why Truetemp-like beliefs fail to qualify as knowledge.<sup>81</sup> In the original case Truetemp, due to surgical enhancement, is able to reliably detect the temperature. Unfortunately, he does not know he has this ability. He takes his beliefs about the temperature to be mere hunches and is unsure why he can't shake his conviction. Intuitively, despite reaching the truth because of reliable cognitive mechanism Truetemp fails to know. By thinking that his temperature is the result of a lucky hunch, the only way his belief can be rendered epistemically responsible—and suitable for knowledge—is through more intellectual labor. He would need to confirm one's belief with a thermometer, or with a weather report—until then one's belief is not responsibly formed and falls short of knowledge.

I suggest that in the lottery case S's belief fails to be responsibly formed because something like (Anti-Luck) is being violated. The reason why we are generally reluctant to credit S with knowledge that her lottery ticket will lose, even after having performed seemingly great inductive reasoning, is that we think that if our beliefs regarding the outcome of a fair lottery are true, their truth is the result of luck. Therefore, if S's doxastic makeup is similar to ours then in order to responsibly form the lottery belief, then she must do more intellectual labor than what is needed in the garbage chute case where there seems to be no violation of (Anti-Luck).<sup>82</sup>

Why is there a violation of something like (Anti-Luck) in the lottery case? I have two reasons for thinking this. First, I borrow a point made by Greco in his analysis of the lottery problem:

---

<sup>81</sup> See Lehrer (1990): p. 162ff.

<sup>82</sup> If S's doxastic state is not similar to ours then my intuitions change and I am less reluctant to deny that she knows that her ticket will lose. I say more on this in what follows.

Why does the element of chance become salient in the lottery case? I would suggest that the very idea of a lottery has the idea of chance built right into it. Here is the way we think of the lottery case: first, S reasons on the basis of excellent grounds that she will lose the lottery. Second, the lottery is held and reality either does or does not match up with S's belief—it's just a matter of chance.<sup>83</sup>

Contra Greco, I claim that what is relevant about this fact for an account of knowledge is not that S's intellectual faculties are no longer the most salient feature in a causal explanation for why S formed a true belief. Rather, I claim that this fact points to why the requirements for responsible belief formation with respect to the lottery proposition are more demanding than with respect to the garbage chute case. Because "the very idea of a lottery has the idea of chance built right into it," we tend to think that even if one is right in believing that her lottery ticket will lose, reaching the truth is matter of luck. And as a result the standards for epistemically responsible behavior increase. Forming a belief on past experiences and probability calculations is not enough. One must attempt to remove the notion of chance unique to the lottery case.

This understanding of responsible behavior also accounts for why our evaluation of S changes if she reads the newspaper confirming her belief regarding her losing ticket, or if she has good evidence that the lottery is rigged. Now her inductive reasoning is based on the reliability of the newspaper, or on the evidence of the testifier, and the role of chance involved in the lottery is no longer built into her reasoning process.

Greco also presents a case in which S visits another culture and observes lottery-like behavior, but has no knowledge of lotteries. According to my account of epistemic responsibility, if "over a long period of time S observes many people exchanging dollar bills for small slips of paper, which they invariably discard after brief examination, [and]

---

<sup>83</sup> Greco (2003): p. 124.

on this basis of excellent inductive reasoning, S concludes that the next person in line will soon discard his slip of paper,” then S’s belief would qualify as knowledge. In this situation the idea of chance that elevates the standards of epistemic responsibility in the typical lottery case, is no longer present. Therefore, in this case the standards for S more closely resemble her standards in the garbage chute case.

Greco finds this result counterintuitive. I on the other hand find my intuitions pulled in both directions, which my account of epistemic responsibility can account for. The intuition that S does not know—in the modified case in which he has no concept of the lottery—that the next person in line will discard his slip of paper is accounted for by our projection of our standards for epistemic responsibility on S. We might think that S is still responsible to do more since after all her belief still concerns the lottery. However, according to this account such projection would be mistaken.

Second, whether we think chance is a relevant feature in our epistemic situation is often influenced by practical considerations in the situation. How we evaluate the chances of some outcome is sometimes influenced by how much we stand to gain or lose by the outcome being realized. We tend to elevate the chances of some state being realized—and to some degree, rightly so—if we stand to gain or lose a lot by its realization. In fact, psychological studies have shown that people can be manipulated to overestimate the likelihood of some event merely by changing the description of what we stand to gain or lose from the event’s realization—in particular, when such potential gains or losses are made salient.<sup>84</sup> In the normal lottery case, that S stands to win a lot of money if her ticket is selected may have a psychological impact on how we evaluate the possibility of her winning—the possibility becomes more relevant than if there were no

---

<sup>84</sup> See Johnson, Hershey, Meszaros and Kunreuther (1993).

practical considerations in play. Because of its relevance we tend to think that if S believes truthfully that her ticket will lose, her reaching the truth was a matter of luck. And in so far as S's psychology resembles ours she will be unable to responsibly believe that her ticket will lose purely based on inductive reasoning.<sup>85</sup>

I have argued that what distinguishes the lottery case from the garbage chute case is the psychology of the subject forming the belief. Typically when forming the relevant lottery belief we think that chance is playing a major role in our getting at the truth. This doxastic state prevents us from satisfying the (Anti-Luck) condition for intellectually responsible behavior. By failing to form our belief in a responsible manner then we fall short of knowledge. Of course, nothing that I have said *precludes* one from knowing lottery propositions. In fact I think our intuitions tell us that there are cases in which a subject does in fact know that that her lottery ticket will lose. These intuitions are confirmed in various "come-on" cases in which an individual successfully removes the idea of luck from the lottery proposition. Consider the following passages from John Hawthorne's *Knowledge and Lotteries*:

Try raising the possibility of lottery success to people who are planning out their lives. Very often, they will respond with 'You know that's not going to happen' or 'I know full well I'm not going to get that lucky.' Similarly, when someone is deliberating about whether to buy a lottery ticket, ordinary people will often say 'You know you are wasting your money.' Granted we sometimes make knowledge claims using a tone indicating that we are not to be taken literally. But I see no good evidence that this is always going on in these cases.<sup>86</sup>

As we have seen, there do seem to be circumstances in which people flat out assert that they will not win the lottery: I'm busy trying to figure out how to save my

---

<sup>85</sup> Isn't the fact that S forms her belief that her ticket will lose via inductive reason supposed to preclude that she is distorting the likelihood of her ticket winning? Not necessarily. It is not uncommon for us to engage in a sort of double thinking in which we are unable to shake the belief that p, and the belief that not p. Of course the case can be constructed so that S is not vulnerable to practical considerations. However, once again, then my intuitions that S does not know starts to weaken.

<sup>86</sup> Hawthorne (2004): p. 18.

business. You say ‘You may win the lottery’. I say ‘You know that’s not going to happen. There are over 10 million tickets sold and only one big winner’.<sup>87</sup>

Although Hawthorne in these passages is primarily concerned with the connection between knowledge and norms of assertion, I think he touches on an intuition many of us have—that in some situations we can come to know lottery propositions. It should not be surprising that in the cases Hawthorne describes the practical considerations of winning the lottery are either out-weighed or cancelled out by other practical considerations of the subject. Given their situations the subjects are able to form the relevant lottery belief by means of great inductive reasoning and without violating (Anti-Luck). As a result their beliefs are formed responsibly and qualify as knowledge. However, such situations are not ubiquitous, and all too often we are unable to remove the notion of luck from lottery propositions—either because the idea of luck is a part of our concept of a lottery and the concept of the lottery is salient in our reasoning process, or because we bring the chance of winning to salience because of practical considerations. In these, more typical lottery cases, we fail to behave responsibly in virtue of our violating something like (Anti-Luck).

Before concluding let me return to the following argument:

1. S knows that S won’t have enough money to go on a safari this year.
2. If S knows that S won’t have enough money to go on a safari this year, then S is in a position to know that S will not win a major prize in a lottery this year.
3. Hence, S is in a position to know that S will not win a major prize in a lottery this year.

According to my account of epistemic responsibility, this argument will only be valid if the standards for responsible behavior remain constant. However, as I have argued above, this will not be the case as one moves from the safari belief to the lottery belief. Once the lottery proposition is brought to mind, the standards for intellectual behavior increase,

---

<sup>87</sup> Ibid.: p. 84

and without satisfaction knowledge is precluded. This preserves our intuitions that S can know that she won't have enough money, and that S isn't in a position to know that she won't win the lottery simply by meeting the requirements for knowledge with regards to the safari. With the lottery belief, more intellectual labor is required of her.

### V. Conclusion

Greco introduces his treatment the lottery problem by claiming that the “key idea here is not that knowledge requires responsibility *in* one's conduct...but that knowledge requires responsibility *for* true belief. Again, to say that someone knows is to say that his believing the truth can be credited to him. It is to say that the person got things right owing to his own abilities, efforts, and actions, rather than owing to dumb luck, or blind chance, or something else.”<sup>88</sup> It seems to me that in order to be responsible for one's true belief one has to be responsible in one's intellectual conduct. And that focusing on responsible intellectual conduct provides one with the resources for handling the lottery problem.

In this chapter I have argued that various attempts within virtue epistemology to solve the lottery problem are unsuccessful. I suggest that their failure in adequately addressing this problem stems from an impoverished understanding of intellectual virtues. This view is the result of virtue epistemologists' focusing solely on the problem of veritic luck. I agree with such epistemologists that veritic luck is worthy of our attention, and that any adequate account of knowledge must eliminate this sort of luck. However, I argue that the relevant difference between the lottery case and other cases of fallible knowledge does not concern the reliability of our intellectual faculties. Rather, the difference is that S's lottery belief is not responsibly formed whereas her garbage chute

---

<sup>88</sup>Greco (2003): p. 111.

belief is. This account of epistemic responsibility attempts to eliminate another form of luck besides veritic luck—what I have called subjective luck. I suggest that by understanding intellectual virtues as analogous to the classical conception of moral virtues, as something that requires responsible behavior, the problem of subjective luck can be adequately addressed. Given that this account of intellectual virtues does a better job addressing the lottery problem, this gives us reason to view the reliabilist versions of virtue epistemology as impoverished, and to make room for the responsibilist conception of intellectual virtues in our account of knowledge.

## Chapter 3: The Value of Virtuous Behavior: An Answer the Value Problem of Knowledge

### I. Introduction.

Although we value true beliefs, a true belief luckily acquired is not as valuable as knowledge. Indeed, we value knowledge more than any proper subset of its constituents. The amount of philosophical attention and labor devoted to the concept of knowledge confirms this. Considerable work has gone into understanding the concept of justification, but much with the ultimate goal of providing insight into the nature of knowledge.<sup>89</sup> However, explaining *why* knowledge is more valuable than an unknown true belief has proven extremely difficult. Jonathan Kvanvig recently argued that all attempts to explain the value of knowledge fail, and concludes that “knowledge does not have value exceeding that of its subparts.”<sup>90</sup> In this chapter I argue that employing the notions of intellectual virtues and of intellectually virtuous—or skillful—activity *can* provide an adequate explanation of value of knowledge. I first examine a promising first step to solving value problem of knowledge by appealing to intellectual virtues and credit. This attempt comes from the separate work of Ernest Sosa and Wayne Riggs. I argue that their solution is ultimately inadequate because of an impoverished conception of intellectual virtues. I then attempt to show how a different conception of intellectual virtue, one that requires responsible intellectual behavior—and which I call intellectual skill—is better suited to explain the value of knowledge.

---

<sup>89</sup> Mark Kaplan (1985) draws a different conclusion from this data. He suggests that the intuition that knowledge is distinctively valuable is mistaken, and that justification is really what we care about. However, this conclusion seems to spring from what I take to be an overly pessimistic stance on the likelihood of adequately addressing Gettier cases. Furthermore, Kaplan seems to assume that the relevant sort of value is practical value. However, as I will go on to discuss one may concede that knowledge has no more *practical* value than a justified true belief and maintain that there is still some other sort of value unique to knowledge—a sort of *epistemic* value.

<sup>90</sup>Jonathan L. Kvanvig (2003): p. 209.

## II. The Value Problem and The Swamping Problem.

The difficulty explaining the value of knowledge seems to stem from epistemologists' attraction to two related theses regarding epistemic value. First is the thesis that Ernest Sosa calls "epistemic truth monism:" the view that true belief is the fundamental epistemic value.<sup>91</sup> Much of its attraction comes from recognizing that the aim of beliefs is truth. Consequently, it seems natural to evaluate beliefs primarily on whether or not they achieve their aim.

Second, given "epistemic truth monism," many believe the value of all other epistemic properties—particularly, epistemic justification—derives from their instrumental connection to truth.<sup>92</sup> Both internalists and externalists express commitment to epistemic truth monism and the instrumental value of justification. According to Laurence Bonjour,

[the] basic role of justification is that of a means to truth, a more directly attainable mediating link between our subjective starting point and our objective goal... If epistemic justification were not conducive in this way, if finding epistemically justified beliefs did not substantially increase the likelihood of finding true ones, then epistemic justification would be irrelevant to our main cognitive goal and of dubious worth.<sup>93</sup>

For now let's suppose that justification converts a true belief into knowledge. If justification only has instrumental value for getting at the truth, then it's difficult to see why a justified true belief is more valuable than an unjustified true belief. In both cases the goal of getting a true belief will have been achieved. Many examples have been recently advanced to illustrate this point. For example, the value of a coffee machine

---

<sup>91</sup> Sosa (2007): pp. 70-1.

<sup>92</sup> Riggs (2002) suggests that the popularity of 'naturalized epistemology' is a major part of the problem. According to 'naturalized' theories, as Riggs understands them, all epistemic concepts are reducible to physical, naturalistic, concepts. (Does this apply to internalist theories like Bonjour's?)

<sup>93</sup> Bonjour (1985): p. 7-8.

seems completely dependent on the value of delicious and aromatic coffee; its value being proportionate to its reliability in giving us what we value. But now assume that we have two equally delicious and aromatic cups of coffee. A reliable coffee machine produced one cup, and an unreliable machine produced the other. It seems that the reliably produced cup gets no added value by being so produced. Surely it does not become more delicious or aromatic. When considering the value of a cup of coffee, whether a reliable machine produced it seems irrelevant.<sup>94</sup> Similarly, given one has a true belief, whether it was produced by a reliable process seems inconsequential to the belief's value.

This does not show that justification is not valuable. A justified belief *is* better than an unjustified belief. The lesson, rather, is that simply finding another valuable property of knowledge beyond truth is not sufficient to explain its value. The other value cannot be parasitic—in this instrumentalist sense—on the value of truth; otherwise it will be 'swamped' by truth's presence.<sup>95</sup> The machine-product analogy generated by the instrumentalist view of justification is incapable of adequately explaining the value of knowledge. If justification's value is only instrumental—like that of a machine—it will not add value to the product, i.e. the true belief.

### **III. Virtue Epistemology to the Rescue.**

Virtue epistemology can help at this point. It is common to think of knowledge as a non-accidentally true belief. And according to virtue theoretic accounts of knowledge what makes a belief non-accidentally true is some relation a believer bears to her true belief

---

<sup>94</sup> This example has been used in the works of Linda Zagzebski, Wayne Riggs, Ernest Sosa, Jonathan Kvanvig, to name a few.

<sup>95</sup> The term 'swamped' is taken from Kvanvig (2003) in which he calls the problem with epistemic truth monism the swamping problem.

making the belief virtuous. Many virtue epistemologists suggest that knowledge requires that the believer be sufficiently creditable for the true belief, and only epistemic agents who acquire a true belief through the use of their intellectual virtues deserve credit for reaching the truth. The following quote from Ernest Sosa seems to be the first incarnation of such an approach:

We have reached the view that knowledge is true belief out of intellectual virtue, belief that turns out right by reason of the virtue and not just by coincidence.<sup>96</sup>

Being credited for the true belief is intrinsically valuable, since it signifies a kind of cognitive achievement not obtained from non-knowledge states. Linda Zagzebski claims:

[k]nowledge is valuable not only because it involves having a valuable possession – the truth—but because it involves a valuable relation between the knower and the truth. This relation is a credit to the knower ... knowledge is not merely something that happens to us.<sup>97</sup>

John Greco agrees, stating that

when we attribute knowledge to someone we imply that it is to his credit that he got things right. It is not because the person is lucky that he believes the truth—it is because of his own cognitive abilities.<sup>98</sup>

Appealing to virtues and credit shifts the focal point of value. Instead of locating the added value in the reliably produced belief, virtue theoretic accounts locate the added value either in the believer or in her performance.

### *III.i. Riggs, Sosa, and the Value Problem*

Both Sosa and Riggs view intellectual virtues along reliabilist lines—as “psychological mechanism[s] that would deliver a high enough preponderance of true beliefs...., at least

---

<sup>96</sup> Ernest Sosa (1991): p. 277.

<sup>97</sup> Zagzebski (1996): pp 260-1.

<sup>98</sup> Greco (2003): p. 123.

in normal circumstances.”<sup>99</sup> Additionally, such mechanisms must be sufficiently stable to qualify as a virtue—it must be a disposition of the individual. If an individual is disposed to form true beliefs through the operation of some cognitive faculty, then it qualifies as an intellectual virtue. For example, Ernest Sosa recognizes eyesight as an intellectual virtue:

... if we include grasping the truth about one’s environment among the proper ends of a human being, then the faculty of sight would seem in a broad sense a virtue in human beings; and if grasping the truth is an intellectual matter then that virtue is also in a straightforward sense an intellectual virtue.<sup>100</sup>

We have various intellectual virtues—cognitive faculties, abilities, powers, and skills that reliably yield true beliefs—and when our true beliefs are sufficiently causally determined by the operation of such virtues, the resulting beliefs are non-accidentally produced. Riggs takes this as an answer to the value problem:

Bringing about some good end in a non-accidental way is more valuable than doing so accidentally, because one deserves more credit in the former case than in the latter.... A person who is causally efficacious in bringing about some positively valuable outcome is ‘due’ some amount of credit for having done so.<sup>101</sup>

If we take knowledge as a virtuously—and thereby, reliably—formed true belief, then the belief was formed *non-accidentally*. Consequently, the believer deserves credit for the true belief. According to Riggs, the added value of knowledge is attached to the believer—particularly, the achievement of the believer—rather than to the belief itself.

Sosa provides a similar explanation for the value of knowledge. He distinguishes three different kinds of evaluations we can make with regards to knowledge. We can

---

<sup>99</sup> Sosa (2003): p. 163. Riggs does not endorse this conception though he thinks it sufficient to deal with the value problem of knowledge. In Riggs (2002) he attempts to provide reliabilist theories with a response to the value problem. He does not explicitly mention virtue theoretic accounts of knowledge, but his repeated employment of notions like intellectual “skills,” “abilities,” and “powers” closely resembles the phraseology used in such accounts.

<sup>100</sup> Sosa (1991): p. 271.

<sup>101</sup> Riggs (2002): p. 92.

evaluate the goods produced; e.g., the belief. We can also evaluate how useful something is in securing goods we independently value. Sosa acknowledges that if we are limited to these two forms of evaluation then the value of knowledge remains unexplained.

However, “[the] evaluation of a particular *performance* is distinct from the evaluation of the artefact that then performs and of any performance-transcendent product of the performance.”<sup>102</sup> We might evaluate a performance positively even when we fail to secure any further goods. However, when such goods are secured because of one’s good performance then the success of the performance is credited to the agent/artefact. This credit has intrinsic value transcending the value of the product.

Both Sosa and Riggs offer analogous illustrations of the value of credit.

(Riggs): Consider a 10-year-old child shooting pool for the first time. She doesn’t know how to hold the cue properly, how to line up a shot properly, etc. Nevertheless, she manages to sink the ball she was aiming at into the pocket she intended. This happens, of course, only by the sheerest chance. Given a hundred chances to make a similar shot, she would make not even one.... What makes the child’s shot lucky or accidental is that the end or good produced was not brought about by any skill or ability the child has. If a skilled pool player attempted and made the same shot, we would assume that it was no accident—the ball went in because the pool player exercises the skills and abilities to make that happen, unlike the child.<sup>103</sup>

(Sosa): Someone with a barely competent tennis serve may blast an ace past his opponent at 130 mph. This is a most effective serve given its outcome: a ball streaking past the receiver untouched, having bounced within the service court. But from another point of view it may not have been so positively evaluable after all. If the player is a rank beginner, for example, one most unlikely to reproduce that performance or anything close to it, then one may reasonably withhold one’s encomium.... Creditable performances must be attributable to the agent’s skills and virtues, and thus attributable to the agent himself.<sup>104</sup>

Neither the immature child nor the novice tennis player deserve credit for their successful action because they do not have sufficient skill, making their success sheer

---

<sup>102</sup> Sosa (2003): p. 169, [emphasis added](#).

<sup>103</sup> Riggs (2002): p. 89.

<sup>104</sup> Sosa (2003): p. 167

luck. Even though the performance of the pool shark or Wimbledon champ is no more successful than their unskilled counterparts, they *do* receive credit for their accomplishments. And credit for a successful performance has added intrinsic value beyond the intrinsic value of the goods it produces. Therefore, knowledge has value beyond the intrinsic value of the true belief itself because it consists in being *credited* for the true belief.

#### **IV. Problems for Sosa and Riggs.**

To this point I have been connecting credit and luck in a way that might seem infelicitous. If we look at human actions there are seemingly clear cases in which a performer deserves credit despite her success being lucky. Consider the ballerina who gracefully and skillfully completes her dance although the stage on which she danced could have easily collapsed. Although in some sense lucky, she still deserves credit for her success.

However, other kinds of luck *are* sufficient to withhold credit. If one's success is not sufficiently caused by one's skillful behavior then not only is one's success lucky, but one is not credited for that success. This sense of luck is illustrated in the cases above. Neither the pool nor tennis novice's performance is sufficiently skillful to receive credit.

Whether appealing to credit provides a sufficient explanation of the value of knowledge will depend on what it means for an individual's performance to be sufficiently skillful. I argue Sosa's and Riggs's understanding of skillful (virtuous) performance is neither necessary nor sufficient for the kind of credit that provides additional value beyond the value of the product.

##### *IV.i. A Small Digression: Two Potential Worries*

One might worry that Riggs's and Sosa's solution provides only a partial explanation of the value of knowledge. Part of what makes being in a state of knowledge more valuable than being in a state of having an unknown true belief, is that the *belief itself* is more valuable than it would be otherwise. It seems more costly and lamentable to lose a known belief than to lose an unknown true belief. If giving a full explanation for the value of knowledge requires accommodating this intuition, then locating added value merely in an agent's performance is insufficient. One must also show how the product becomes more valuable because of the performance. However, as we've seen with the swamping problem, if "what makes a feature of a subject's psychology an 'intellectual virtue' is the reliable tendency of that feature to give rise to true beliefs on the part of that subject,"<sup>105</sup> then no added value will be conferred onto the belief.

Sosa may escape this worry by suggesting that forming and retaining a belief just *is* a performance, similar to a ballerina's dance or an athlete's run. In other words we shouldn't think of beliefs as products of a performance in the way a statue is the product of an artist's performance.<sup>106</sup> Evaluating Sosa's position will take us into the metaphysics of beliefs and beyond the scope of the chapter. However, some might find identifying beliefs with performances implausible. Toward the end of the paper I will address this problem and show how adopting a different conception of intellectual virtues can accommodate the intuition that the known belief gains added value without needing to identify it with a performance.

A second worry with the above solutions is that they fail to explain why an epistemically justified true belief is better than an unjustified true belief. I suspect that

---

<sup>105</sup> Ibid.; p. 163.

<sup>106</sup> Sosa (2007): p. 23.

those who have the intuition that knowledge is more valuable than an unknown true belief will also have the intuition that a justified true belief is more valuable than an unjustified true belief. I do. However, if epistemic justification is understood in purely reliabilist terms then it seems difficult to see how this intuition can be explained.

The reason why explaining the value of knowledge is so difficult given the J-T-B account of knowledge is because epistemic justification is viewed as merely an instrumental good to true beliefs. Sosa and Riggs don't abandon the instrumental conception of justification, and explain the value of knowledge by appealing to credit. But if a justified true belief is indeed more valuable than an unjustified true belief then the same examples showing that the value of justification is swamped by the value of the true belief—recall the coffee example above—will pose a problem for their account. This time the problem is not that they can't explain why knowledge is more valuable than a mere true belief, but that they can't explain why a justified true belief is more valuable than a mere true belief.

#### *IV.ii. On Necessity.*

According to Sosa and Riggs in order to receive credit for a true belief the truth of the belief should be explained by reliable mechanisms seated within the individual.

However when we look at human performances generally it seems we often assign credit to individuals even when they are unreliable in bringing about the desired end. Consider Riggs's 10-year-old pool player. Even though she lacks the skill to be a reliable shot, she has *some* skill and ability. Imagine that before taking her shot a pool shark informs her of the appropriate technique for making her intended shot. She then follows the expert's instructions perfectly and sinks the shot. Even though she is an unreliable player—she

couldn't reliably bring to mind and follow the relevant instructions on her own—surely she still deserves credit for her success.

Whether an individual is credited for her successful performance depends on i) the internal constitution of the performer and ii) her relationship to her surrounding environment.<sup>107</sup> It is possible for one's internal constitution to be properly aligned to one's environment, resulting in a successful performance, while still failing to receive credit. This depends on *why* one's internal constitution is properly aligned on the particular occasion. If one's character is unreliable we typically withhold credit because it suggests that the individual was not sufficiently responsible for the success. Perhaps there was an abnormal change in the environment or perhaps some external force changed one's internal constitution in a fortuitous way. But this may not be the case. Sometimes, the reason for successful performance and proper alignment with one's environment is sufficiently internal to the individual, despite one's general unreliability. For example, if the reason for one's proper alignment is one's intentional and responsible behavior, then, despite one's general unreliability in bringing about one's aim, one still deserves credit for one's success. We credit the 10-year old for making the shot because *she* was sufficiently responsible for properly aligning herself with her environment, even though it may have been an isolated event.

This point concerning credit applies to our intellectual performances as well.

Consider the following example concerning knowledge:

Jill has never been particularly good at math. She finds the subject matter extremely boring and tedious. As a result she never pays sufficient attention to the problems at hand, and repeatedly makes simple arithmetical errors. However, today is different.

---

<sup>107</sup> Cf. Sosa (2007): p. 81.

Jill is unusually focused and careful in calculating some difficult bit of arithmetic, and consequently gets the right answer.<sup>108</sup>

Although Jill acted out of character, she performed responsibly, and as a result deserves credit for her success.

Jennifer Lackey, who originally made this point, claims that this sort of example highlights a fundamental flaw with making the credit one receives for a successful performance proportionate to one's character:

Accordingly, when credit is analyzed in terms of revealing a person's character, such unexpectedly spectacular achievements turn out to be ones for which the subject in question fails to deserve credit, thereby leading to the consequence that mediocre thinkers cannot have knowledge of outstanding intellectual discoveries. But surely it is not just great people who deserve credit for or can have knowledge of great things. For a person's achievement can result from the working of her own faculties and skills and thus be properly creditable to her without necessarily revealing her character.<sup>109</sup>

Since Sosa and Riggs make credit dependent on virtuous activity, and because they understand virtuous activity in terms of the performance of stable and reliable processes/faculties within an intellectual agent, they fail to account for the possibility of creditably acting *out* of character.

Lackey goes on to argue that these cases pose problems for any virtue-theoretic account that claims the added value in knowing *p* is the result of the credit due to the believer for truly believing *p*. Moreover, she claims that knowledge does not even require credit. Since I will be proposing a solution to the value problem that does appeal to credit, I will return to Lackey's more general worries towards the end of this chapter.

#### *IV.iii. On Sufficiency.*

---

<sup>108</sup> This type of case is borrowed from Lackey (2003).

<sup>109</sup> *Ibid.*: pp. 16-17.

Let's suppose that success through generally reliable performances—even if not necessary for credit—is sufficient for credit. One might question the assumption that this sort of credit is really valuable. If we look at performances generally this assumption may seem doubtful.

James and John are identical twins. After extensive physical examinations immediately after birth, everything checked out fine with James, but John was diagnosed with an undersized heart. Given the size of his heart, John's health was expected to be worse than James's. However, through a stroke of luck (perhaps another deformity or malfunction off-setting the undersized heart), John experiences perfect health throughout his life. Although his heart is unreliable in sustaining healthy individuals, John's health was practically identical to James's.

According to the above view only James is credited for his good health. And since, presumably, being credited for a successful performance is more valuable than a non-credited performance, James's health should be more valuable than John's health despite being qualitatively identical. Is this evaluation right? It seems to me that it is not. What we value is whether or not John's health is inferior to James's (or to a normal human being's). Whether John's heart is generally unreliable in sustaining a healthy individual seems irrelevant to how we evaluate his health. Suppose we are given the choice between John's and James's physical makeup, knowing they result in qualitatively identical health for each individual. It seems to me we should be indifferent to whether we acquired John's or James's physical makeup. John's heart may have weaker beats and pump less blood than James's, making it less likely to sustain a healthy individual, but his health seems no less valuable.

Of course, one might deny that these states of affairs are performances, and adopt a more literal sense of our teleological terms. Accordingly, only agents will be performers. Circulating blood is certainly not a performance that *we* do. Even if our

hearts are generally reliable *we* are not credited for their success. The reason for this seems clear: we are not sufficiently responsible or in control of such events. But this move is not available to Sosa and Riggs. Under their account, success because of reliability is supposed to be sufficient for credit and for added value. Accordingly, they are faced with the following question: why is James credited for a true belief formed because of reliable cognitive processes/mechanism but not for good blood circulation produced by a reliable heart?<sup>110</sup> Both cases involve reliable processes and mechanisms, so reliability does not distinguish these cases. Perhaps there is something other than reliability that accounts for the difference. For example if belief formation and retention is something we are responsible for, then we might be able to account for the difference between the two cases, since James would be responsible for his true belief in a way he is not for his circulatory system. Unfortunately, Sosa and Riggs have not given us the resources to explain credit in this way.

Another way to highlight this shortcoming is by examining whether their account *fully* accommodates our intuitions of intellectual/epistemic value. If an agent receives credit for a true belief sufficiently caused by his reliable abilities, powers, virtues, etc., an agent should receive discredit for a false belief sufficiently caused by unreliable processes or mechanisms within the agent. Furthermore, if credit adds value when the product is a good, then it seems that discredit adds disvalue when the product is a bad. Just as a creditable true belief is better than a mere true belief, a creditable false belief is worse than a mere false belief.

---

<sup>110</sup> I'm using credit in a value laden sense. Perhaps Sosa and Riggs will claim that James does receive some credit for his blood circulation. But what I've tried to argue is that this is not the sort of credit that adds value. I acknowledge that there might be cases in which James clearly does deserve credit—in this value laden sense—for his good health. For example if James actively leads a healthy lifestyle by eating right and exercising. However these kinds of cases involve more than mere reliability.

However, the above analysis for false beliefs seems counter-intuitive. Consider someone whose cognitive faculties and processes, through no fault of one's own, are incredibly unreliable—perhaps an evil demon rewired his brain. Do we discredit such an individual for his false beliefs? Moreover, do we always evaluate such beliefs as epistemically worse than other false beliefs formed by processes that typically yield true beliefs? It seems to me that we don't. The false beliefs of the re-wired human are not worse than some others less reliably formed. In fact, we often mitigate our criticism of such beliefs precisely because the cognitive processes and faculties responsible for their formation are *incapable* of forming true beliefs. Isn't this the lesson of the problem of the New Evil Demon?

On the other hand we *do* discredit an individual for forming false beliefs from normally truth-conducive processes or faculties in situations where she 'should have known better' than to use such processes. Surely the beliefs of this irresponsible agent are worse than those of the evil demon victim or of the mentally insane. But, again, the reliabilist account of credit cannot appeal to responsibility in explaining our evaluations of false beliefs.

#### *IV.iv. In Conclusion.*

If the notion of credit adequately explains the value of knowledge it should account for the fact that we sometimes deserve credit for performances that are out of character, and that added value is not always attached to a non-accidental performance in Sosa's and Riggs's sense. I've suggested that our evaluations of added value depend not on reliability but on responsible and skillful behavior. We credit the unreliable math student because she acted responsibly in her calculations, using the appropriate means for ascertaining the truth. Also, we are reluctant to assign discredit (or discredit that implies

value) to the evil demon victim or the mentally insane/handicapped because they are not responsible for their false beliefs. Given this correlation between credit and responsibility, I claim we would do better to adopt a conception of intellectual virtues, and of epistemically virtuous performance more analogous to a conception of moral virtues and of morally virtuous performance in ethics, respectively.

### **V. Credit and Intellectually Skillful Behavior.**

I am not the first to pursue this line of thought. As I've discussed in the first chapter, Linda Zagzebski has employed the epistemic analogue of moral virtues to analyze knowledge, and she also specifically addresses the value problem. On her account, virtues—whether intellectual or moral—are characterized by a motivational component. Just as one cannot exhibit moral virtue without the proper motivation, one cannot exhibit intellectual virtue without the proper motivation. The only difference between the two classes of virtues is the *aim* of the motivation; the characteristic motive for intellectual virtues is a love of truth.

For Zagzebski, evaluations of actions depend on whether they are virtuously or viciously performed, and since virtues are characterized by having the right motivation, the evaluations will depend on the motivations of the agent. Furthermore, there is added value in being properly motivated because of “an accurate emotional fit” between the action/belief and its intentional object.<sup>111</sup> Acts of compassion are more valuable than disinterested acts that aim at alleviating suffering because of the proper fit between the suffering of others and one's disvaluing of such suffering. Similarly, knowledge (or a true belief virtuously formed) is better than a mere true belief because knowledge necessarily involves one being motivated by a love of truth.

---

<sup>111</sup> Zagzebski (2003): pp. 146-150.

Although, Zagzebski correctly points out that virtuously motivated acts are better than acts lacking such motivation, there are serious difficulties with her view. First, simply having the right motive is insufficient for morally virtuous action. I may be motivated to help others, but only if that motive is manifested in the *appropriate* action will my action be considered virtuous. Moral virtues have a motivational component, but this is just one component. In order to fully exemplify a virtue one's intended act must also be of the right sort. This is why *phronesis* (sometimes translated as 'practical wisdom') is so important for virtuous behavior. One must be able to see and do what is most appropriate in a given situation. If one fails in intending the right action, then, even if one's heart is in the right place, the action is not virtuous.

This is problematic for Zagzebski's account because a virtuously motivated true belief does not, by itself, qualify as knowledge. One can be motivated by a love for truth and be excessively dogmatic, credulous, cowardly, etc. Being virtuously motivated does not entail being intellectually responsible. And since part of the problem was to identify what makes knowledge more valuable than any of its subparts, some other feature beyond virtuous motivation—I argue, intellectual responsibility—will have to explain the value gap that exists between knowledge and non-knowledge states.

Second, it seems being motivated by a love of truth is not necessary for knowledge. What is important is my intellectual behavior, not my emotional state when performing. If I responsibly follow good belief forming procedures, my motivation seems inconsequential for knowledge. If having the right motivation is not necessary for knowledge, then having the right motivation cannot be what makes knowledge more valuable than any of its subparts.

Unlike in the moral realm, *intellectually* virtuous activity does not require proper motivation. However, given that the aim of intellectual inquiry is to acquire true belief and avoid false beliefs, in order to behave in an intellectually responsible manner one must engage in the intellectual practices and employ those cognitive faculties that one has good reason to believe will reliably yield true beliefs. For example the typical, cognitively mature, individual has good reason to believe that one's perceptual faculties in normal conditions will reliably yield true perceptual beliefs. Similarly, she will have good reason to believe that relying on a tarot card reader for beliefs about one's future is an unreliable guide for truth. How we approach our various processes and faculties will determine whether we are acting responsibly or not. If we show excessive trust in a tarot card reader, then we are acting irresponsibly. On the other hand if exhibit fair-mindedness in giving reasonable weight to the opinions of intellectual peers then we are behaving responsibly. With regard to our intellectual endeavors there are many different processes and mechanisms that generally lead to true beliefs, and sometimes a conflict may arise between two or more of them. Which processes we should pursue or show preference for, will depend heavily on our circumstances. But what's important for our purposes is that behavior flouting the guidelines for appropriate or responsible behavior will not be valued, even if successful. Intellectual behavior can be either responsible or irresponsible and we only value the former.

But even a responsibly successful belief is not enough to solve the value problem. We can be responsible and successful but still fail to acquire knowledge. Sometimes, because of one's background and upbringing going to a fortuneteller, or tarot card reader may still be within the parameters of responsible behavior. But if these means for

acquiring true beliefs are objectively unreliable, then even if one's behavior is successful knowledge will not be gained. Therefore, even if a successful responsible belief is valuable it does not equal the value of knowledge. In order to have knowledge our standards for behaving must *in fact* be reliable. Only when we responsibly engage in reliable intellectual practices, are we candidates for being knowers. I claim that intellectually virtuous behavior requires responsibly following reliable procedures for acquiring truth. Knowledge requires responsibility but it also requires skillful behavior.

This conception of intellectual virtues closely resembles an Aristotelian conception of skills.<sup>112</sup> Attempting to explicate an Aristotelian conception of skills James Wallace states:

A skill is a capacity to do something well in the sense of doing it proficiently. Proficiency involves the mastery of a technique, the mastery of something that is technically difficult.... A skill is the mastery of a technique, and it is something that generally is acquired through learning—through instruction and practice. [Moral virtues] are not masteries of techniques; technique has very little to do with being brave, generous, or honest, nor do these necessarily involve being proficient at any particular thing. Some virtues involve being able to do difficult things, but the difficulties involved are due to contrary inclinations, not to technical difficulties in the actions themselves.<sup>113</sup>

We see here that unlike moral virtues, manifesting skills does not require a particular emotional state. However, if one is proficient in the relevant technique then the skillful behavior is exemplified.<sup>114</sup>

---

<sup>112</sup> James Wallace in Wallace (1978) claims: "The Aristotelian 'intellectual' excellences are either skills or complexes of skills." If he is right about this then my conception of intellectual virtues is a move back to an Aristotelian conception of intellectual virtues. However, I think that there are still significant differences between my conception of intellectual virtues and an Aristotelian conception of skills.

<sup>113</sup> Wallace (1978): pp. 45-6.

<sup>114</sup> Zagzebski agrees with this distinction between skills and moral virtues but claims that intellectual virtues like "intellectual care, thoroughness, perseverance, fairness, and courage are not technically difficult. Their difficulties arise primarily from a lack of sufficient passion for truth or from a desire to appear right in one's own eyes or in the eyes of others or, perhaps, from just plain laziness." (1996, p. 108) However it seems to me that Zagzebski is mistaken on multiple points here. First, even if one does have a

I claim that understanding knowledge as intellectually virtuous, or skillful, behavior gives us the resources for solving the value problem of knowledge. All skillful behavior is directed at some end, and intellectually skillful behavior is directed at truth. But simply attaining the truth is not enough—we have to be credited for the attainment. As Sosa and Riggs point out, the agent's skillful behavior must sufficiently contribute to the attainment of the truth. Therefore, if one is successful in reaching the truth but not *because* of one's intellectually virtuous (or skillful) behavior, then one's true belief is lucky in a way that precludes knowledge. And it is instructive to see how this notion of luck differs from the notion Sosa and Riggs employ. On my account there are two ways in which a belief can be lucky (in the knowledge destroying sense): i) my intellectual behavior is not reliable in yielding true belief, ii) my intellectual behavior is not responsible (because I am following a method for which I lack good reason to believe is reliable). This account nicely explains why—despite their general unreliability—Jill (the mediocre math student) and the instructed pool novice both get credit for their success. For it is sufficiently brought about because of their intellectually virtuous (skillful) behavior.

In summary, in order for additional value to accrue to one's performance one's aim must be achieved through responsible behavior, and such behavior must align with the appropriate—reliable—means to one's aim. As Wallace mentions above, skillful behavior implies technical behavior. There is a right way to shoot an arrow and a wrong

---

passion (love) for truth one may not possess—even moderately—the above virtues. Second, pointing out that laziness, or a desire to be right in one's own eyes or in the eyes of others makes exercising (or acquiring) the various intellectual virtues difficult does not show that they are not skills. These vicious desires will make practically any skill (whether it be artistic, athletic, or any practical skill) difficult to exercise or acquire. Third, that one can exhibit intellectual care, thoroughness, perseverance, fairness and courage without a passion for truth seems manifest. I may display all of these 'intellectual virtues' because of a desire for public approval or because of my love for my wife.

way; a right way to bake a cake and a wrong way; and a right way to form a belief and a wrong way. If one's behavior is to be deemed skillful, then one's action must conform to a certain circumscribed set of guidelines; and like most crafts/skills what is considered proper behavior will depend on the situation. For example, in certain situations trust will be the prescribed behavior, in other situations caution. Knowing how much trust/caution to display in a given situation requires wisdom.

So I follow Sosa, Riggs, Greco, and Zagzebski, in explaining the value of knowledge in terms of credit for successfully attaining the truth. However, my account of credit is significantly different from their accounts. I have claimed that what is required for intellectually virtuous behavior—and therefore for credit—is reliable and responsible—in a word: skillful—behavior. Only when the truth is reached because of one's skillful behavior is the sort of credit that makes knowledge more valuable than its proper subparts obtained.

Issues of doxastic control immediately arise. It seems that in order to be responsible for our beliefs we have to have some sort of control over them. It also seems, however, that we do not have direct control over what we believe. Perhaps a sort of indirect control is all that's needed to be responsible for our beliefs. Unfortunately, dealing with these issues is beyond the scope of this chapter and will therefore be postponed for the next chapter. That said, I have argued that if we are hoping to explain the value of knowledge by appealing to credit, then we must understand credit in terms of intellectual responsibility.

## **VI. Tying Up Loose Ends.**

### *VI.i. Lackey and Credit*

We can now look back to Lackey's general objection against requiring credit for knowledge. One might wonder whether Lackey's worry regarding credit is also applicable to my proposed account of knowledge. My account makes intellectually virtuous behavior a component of knowledge, and intellectual virtues (or skills) are typically understood as integrated traits in an individual's character through habituation.<sup>115</sup> Therefore, it may seem that one is unable to commit an intellectually virtuous *act* without having the appropriate virtue integrated into one's character. However, as Lackey has shown knowledge can be obtained even when the performance is out of character, and when the individual lacks the relevant virtues.

Although, intellectual virtues—in the sense developed in this paper—are not perfectly analogous to moral virtues, they are acquired through habituation, and repeated practice. However, one can perform virtuously without having the virtue in question. Indeed, if habituation is required in order to acquire a virtue, then it *must* be possible to act virtuously prior to having the virtue. If one wants to be courageous (have the virtue of courage), then one has to repeatedly act courageously (and thereby habituate the trait into one's character). By defining knowledge in terms of acting virtuously, rather than being virtuous, one avoids the worry presented by Lackey. It is not the case that in order to receive credit one's act must arise out of one's character. The relevant traits need not be fully integrated, but they can be exemplified in performances that the individual is responsible for.

---

<sup>115</sup> Aristotle distinguishes intellectual virtues and moral virtues by the way that they are acquired. He states: "Intellectual virtue in the main owes its birth and its growth to teaching (for which reason it requires experience and time), while moral virtue comes about as a result of habit." (NE, II, 1, 1103a15-17) However, his notion of habituation is different from the sense in which I'm using here. By habituation I simply mean repeated behavior, where I take Aristotle to mean something more affectively laden where one's emotional state is endorsing one's intentional behavior.

The idea of being responsible for a skillful performance despite lacking the relevant skill may strike some as confused. But if we look at human performances generally I believe we'll find that it is not an uncommon phenomenon. People who are not sympathetic can on occasion show sympathy. People who are not courageous may perform great acts of courage. People who are not benevolent may selflessly give to others. One might suggest these acts of sympathy, courage, or benevolence reveal that the individuals had the underlying trait all along. But this seems unlikely since it is assumed that these acts may be extremely rare occurrences in an individual's life.

This also brings up a question about reliability. I claim that receiving credit requires skillful—responsible and reliable—behavior. I also claim that one can perform skillfully without having the underlying trait fully integrated in one's character. This might sound contradictory. If one doesn't have the relevant trait, then it seems that the individual is not reliable in exemplifying the trait. But if the individual is not reliable then it seems that her performance is not skillful.

It is important for my account that what needs to be reliable is not necessarily the individual but the individual's behavior. If she is aiming at the truth then in order for her to receive credit the *means* she employs must be truth-conducive. Recall the novice pool player. Although *she* is not a reliable player, by responsibly following the instructions of the pool shark she is following reliable means for her aim. And if she makes her shot in virtue of following these instructions then she is credited for her success. She may quickly forget the instructions after making her shot and never be able to repeat her performance. But this would not remove the original credit she received. I suggest that something similar is possible with our intellectual performances (as the Lackey-style

example shows). If we reach the truth by behaving in a way which we have good reason to believe is reliable in the relevant environment, and it is in fact reliable, then we receive credit for our success.

It is important to note that the reliability of a performance is indexed to the individual performer. Intellectual behavior may be reliable for one individual but not for another. For example, whether the intellectual behavior of trusting another's testimony is reliable will depend on whether one's surrounding community is trustworthy. Suppose John's intellectual community is prone to lie but James's isn't. Both may have good reason to think that trusting another's testimony is a reliable mechanism for getting at the truth, but only James's intellectual behavior of trusting another's testimony will be intellectually skillful. Again, James may not have the relevant virtue. His character may be overly untrusting. However, if he reaches the truth because he is able to bring himself to trust his reliable community, then he deserves credit, and acquires knowledge.

#### *VI.ii. The Value of True Belief*

Finally, it will be helpful to see how my account handles the other objections facing Sosa's and Riggs's account. First, I suggested that their account fails to accommodate the intuition that the belief itself is more valuable when it qualifies as an instance of knowledge. Even if their account can explain how the *state of knowing* has value beyond the value of the true belief, an adequate solution must also explain the added value to the *belief*. I claim that understanding knowledge as a true belief reached through skillful behavior can make sense of the intuition. We value responsible and skillful behavior, but it seems to me there is a special value that products of skillful behavior receive for being skillfully formed, exceeding their independent intrinsic value. This is most evident in cases where the performance and the product of the performance are indistinguishable.

Sosa mentions that the ballerina's movements are more valued when they are the produced through skill rather than through drunken stumbles: "We had paid to see *a performance, the product* of artistic excellence and control. We take pleasure in seeing the grace of the movements, true enough, but we take special pleasure in knowing it to be grace due to the ballerina and, more particularly, to her art. It is her actions that we normally admire and value."<sup>116</sup> Greco makes a similar comparison between Ken Griffey Jr making a great catch in virtue of his athletic prowess and his making the same catch in virtue of tripping and the ball fortuitously landing in his glove. Although, in the latter we may still value the catch for its contribution to the game, the first performance (and product) has added value because it manifested Griffey's skill.<sup>117</sup> Because we don't have direct control over our beliefs in the way the ballerina does over her dancing or in the way Griffey does over his catches, our beliefs are not performances in this sense, but rather the products of our performances. However, this does not change the fact that skillful behavior can give added value to the *products* such behavior. And it is also worth noting that this value exceeds that which obtained through mere reliability, since one can be reliable without behaving responsibly.

Sarah Broadie claims that Aristotle disagrees with me on this point. According to Broadie, Aristotle thinks

"we are satisfied with things which are normally produced by art or skill provided they are up to standard, even when they were produced by someone without skill. If we assess what such a doer has done by what he has made, we can say that what he has done is good. The lack of skill implies no defect in what he has done on this occasion, and it might reasonably be claimed that the skill is of value only because whoever possesses it is more likely to produce acceptable articles. Aristotle's point is that it is not like this with virtue and right actions (hence, he implies, virtue is too

---

<sup>116</sup> Sosa (2007): p. 76. Cf. Sosa (2003).

<sup>117</sup> Greco (2003): p. 122.

different from skill for one to be justified in drawing conclusions about virtue from premises about skill).”<sup>118</sup>

I defer to Broadie’s interpretation of Aristotle. But if she is right then I humbly disagree with Aristotle. It seems to me that there are clear cases in which we value the products more for being skillfully produced, and these cases extend to those in which the products are distinct from the performance, for example compare the accidental brush strokes of a child to a Picasso painting or the wild experimentation of rookie chef with a Wolfgang Puck dinner. Even if the paintings or the dinners are qualitatively identical, we value the painting and the dinner that was skillfully produced more highly.

#### *VI.ii. The Value of a Justified True Belief*

We can now also explain why a justified true belief is more valuable than an unjustified true belief. We value responsible behavior for it’s own sake and not simply because of it’s instrumental value to some further end. By taking an epitemically justified belief as an intellectually skillful belief, and incorporating the intellectual responsibility in what it means for a belief to be intellectual skillful, the value of a justified true belief is no longer swamped by the value of truth.

#### *VI.iii. The Value of Credit*

My last objection claimed that even if we accept the reliabilist’s conception of credit, we have reason to doubt its value. First, I argued that this notion of credit does not seem to carry value when we consider performances generally—recall the comparison between the performances of James’s and John’s heart. Second, I argued that this conception of credit is incapable of handling our evaluations of false beliefs produced through the operation of unreliable processes. It seems that if we don’t always disvalue false beliefs

---

<sup>118</sup> Broadie (1991): p. 83.

any more for having been produced by unreliable faculties (e.g. in the case of the evil demon victim or the mentally disabled/insane), then the reliabilist's conception of credit is inadequate.

The conception of credit I have proposed is better suited to accommodate these evaluations. The performances James's heart fails to carry added value because blood circulation is not something we can control or be held responsible for. Accordingly James does not get credit for his successful performance. Of course his *heart* may receive credit, but this sort of credit does not seem to give James's health added value. However, now suppose that the reason John and James have identical health is that John leads a responsible and healthy life while James leads a recklessly unhealthy one. Given identical behavior James would have superior health. It seems that although qualitatively identical, in this case *John's* health is more valuable. If given the choice we would choose John's life over James's. The difference between the two is not that one's life is more likely to result in good health—we can suppose that James's heart/lifestyle pair is equally likely to produce good health as John's heart/lifestyle pair. The difference is that John's life is responsibly lead—with respect to health—which gives him a sort of valuable credit for his health that James does not receive.

When we turn to intellectual performances, if we exhibit intellectually irresponsible behavior, and if such behavior sufficiently contributes to the formation of a false belief, then we receive discredit for the false belief. Under this sense of credit, the state of having a discredibly bad belief *is* worse than that of having a mere false belief. We don't criticize the false beliefs of the evil demon victim whose brain was rewired or of the mentally disabled any more than other false beliefs despite being formed by

unreliable mechanisms. The reason is we recognize that their failure in attaining the truth is not the result of irresponsible behavior on their part. However, if someone who is cognitively capable of acquiring true beliefs but fails through either inattention, laziness, or some other intellectual vice, then we censure his false belief more than had his error been the result of some cognitive malfunction beyond his control. His state is worse for having been the result of irresponsible behavior, and he thereby deserves discredit and blame for this failure.

### **VII. Conclusion.**

In this chapter I have tried to argue that the value of knowledge is difficult to account for if justification is merely understood along instrumentalist lines. In order for knowledge to be more valuable than any of its subparts, there has to be some value in knowledge that is not swamped by the value of a true belief and that exceeds the value of even a justified true belief. Recently, some virtue epistemologists have proposed that what makes knowledge more valuable than any of its subparts is the credit that is bestowed on the intellectual agent for having reached the truth. Having found this line of thought appealing and plausible, I have examined the various ways in which credit has been understood in much of the recent literature on the value problem. I have tried to show that understanding credit along the virtue-reliabilist lines of Riggs, and Sosa, or along the motivational lines of Zagzebski is insufficient for an adequate treatment of the value of knowledge. Instead I have offered my own account that takes credit to be achieved only by intellectually virtuous or skillful behavior on the part of the agent, and have tried to

show how many of the objections and problems besetting the other theories can either be addressed or simply do not arise for my theory.<sup>119</sup>

---

<sup>119</sup> Interestingly, Wayne Riggs has recently argued for a position similar to the one I defend here. He mentions that intention is necessary for the kind of credit needed to respond to the value problem of knowledge. However, in a couple of places he backs up a bit from this position and states that perhaps an appropriate desire will do the job. This retreat would align his view the one put forward by Zagzebski, and I have argued in this paper that this view is inadequate to address the value problem of knowledge.

## Chapter 4: Epistemic Responsibility and Doxastic Control

Ernest Sosa, in “The Raft and the Pyramid,”<sup>120</sup> suggested that epistemology (in the article he focuses on the foundationalist-coherentist debate) could benefit by making agents and their intellectual virtues the primary source of epistemic evaluation—particular beliefs receive a derivative evaluation “in virtue of their source in intellectual virtues or other such justified dispositions.”<sup>121</sup> Eventually he employs the notion of intellectual virtues in a theory of knowledge.<sup>122</sup> However, his reliabilist characterization of intellectual virtues—as some sort of “psychological mechanism that would deliver a high enough preponderance of true beliefs”<sup>123</sup>—parts ways from the more traditional conception of virtues. Sosa’s virtues are simply an individual’s cognitive faculties, such as perception and memory, which, when appropriately placed in one’s environment, are reliable in producing true beliefs.

I have followed Sosa in thinking that intellectual virtues can help make progress in many epistemological debates—particularly in giving a theory of knowledge—but have argued that his understanding of intellectual virtues fails to provide an adequate account of knowledge. Instead I have argued that knowledge requires intellectually skillful performance, understood in terms of reliability and responsibility. However, incorporating intellectual responsibility into my account of knowledge raises new questions. How do the differences between acts and beliefs affect the intelligibility of evaluating beliefs analogously to our evaluation of actions? More specifically, what sort

---

<sup>120</sup> Sosa (1991b).

<sup>121</sup> *Ibid.* 190.

<sup>122</sup> For his most recent treatment see Sosa (2007).

<sup>123</sup> Sosa (2003): p. 10.

of control must we exhibit over our beliefs in order to make evaluations of epistemic responsibility appropriate? This chapter attempts to address this question.

I noted in chapter 1 that drawing an analogy between beliefs and actions is not unnatural. Many of the same evaluative notions are employed to appraise agents, their actions, *and* their beliefs. Beliefs are judged as right and wrong, permissible and impermissible, obligatory and prohibited, responsible and irresponsible. Furthermore, we regularly engage in practices of praising and blaming individuals for their beliefs. We judge people as intellectually virtuous and vicious and hold them responsible in the same way as people who display moral virtues or vices.

Pettit and Smith (1996) have observed that our “conversational stance” toward ourselves and other believers implies responsibility for—and freedom over—our beliefs. Typically, when there is a discrepancy between two people’s beliefs, there is an attempt at resolution, and an assumption that such a discrepancy shows that someone is wrong, and reviewing the evidence will help to establish agreement between the differing parties. Once all the evidence is reviewed, if no resolution is reached, then we generally assume that this is the result of inattention, carelessness, dim-wittedness, or some other intellectual vice by one of the disagreeing parties. Pettit and Smith argue that this dialectical approach only makes sense, if (i) there are norms relevant to what one ought to believe, (ii) that the parties involved can recognize such norms and (iii) the parties involved can respond appropriately to those norms; that they have the ability to believe as they ought. In other words:

To see someone in this way is to see her as a responsible subject in the relevant domain. The interlocutor whose beliefs are engaged is depicted as someone who can

be made to answer to the norms of evidence governing what is the case.... The one person can be held responsible, as we say, for what she believes....<sup>124</sup>

There are at least two ways one can approach the above data. One could take the above observations to reveal something important about beliefs and of how we should evaluate them, or one could suggest that these are simply loose ways of talking and behaving; inaccurately depicting how we are actually related to our beliefs. I think we should avoid resorting to error theories whenever possible, and, therefore, should give *pro tanto* preference to a conception of intellectual virtues that can better accommodate these widespread intuitions and behaviors. With this in mind I now turn my attention to the issue of doxastic control and argue that, despite appearances and arguments to the contrary, we do exhibit sufficient control over our beliefs to make attributions of epistemic responsibility appropriate.

### I. The Voluntarism Argument

The idea that moral and epistemic evaluations are akin precedes the resurgence of virtues in epistemological discussion. A.J. Ayer speaks of epistemic rights, and Roderick Chisholm, Laurence Bonjour, Richard Feldman, Matthias Steup, Richard Fumerton, among others, employ the notions of epistemic requirements, obligations, and permissions.<sup>125</sup> However, some have objected that these deontological concepts are suitable for moral evaluations but malapropos for epistemology. The main reason for this is the apparent dissimilarity between actions and beliefs. Such evaluations are fitting for ethics because our actions are under our voluntary control; our beliefs, however, are not. Richard Feldman puts the argument as follows:

---

<sup>124</sup> Pettit and Smith (1996): p. 441.

<sup>125</sup> Cf. Ayer (1956); Chisholm (1977); Steup (1996); Fumerton (1985).

*The Deontological Voluntarism Argument (DVA)*

1. People do not have voluntary control over their beliefs.
2. If deontological judgments about beliefs are sometimes true, then people have voluntary control over their beliefs.
3. Deontological judgments about beliefs are not sometimes true.<sup>126</sup>

Constructing the argument against deontological theories may be a distraction since our main concern has been a virtue theoretic account of knowledge. However, a prima facie plausible argument can also be constructed against virtue theories. If virtue theories claim that we are responsible for whether our beliefs are viciously or virtuously formed in the way that we are responsible for whether our actions are vicious or virtuous, then our beliefs must be voluntary like our actions. Here is the analogously structured argument:

*The General Voluntarism Argument (GVA)*

- 1\*. People do not have voluntary control over their beliefs
- 2\*. If people are responsible for their beliefs, then people have voluntary control over their beliefs.
- 3\*. People are not responsible for their beliefs

William Alston (1988) meticulously examines the extent of control that we have over our beliefs, and concludes that at most we have ‘indirect voluntary influence’ over our beliefs. We typically can’t decide what to believe in the way we can decide to move our bodies (a basic action), or in the way we can decide to open a door (a non-basic but rather immediate action), or even in the way we can decide to remodel the kitchen (a non-basic and long term action). All of these actions exhibit some sort of voluntary control on the part of the agent for *a particular and determinate* action, and Alston claims that we typically lack such control over our beliefs. Alston acknowledges that we can influence our beliefs by engaging in behavior or developing traits that would make some beliefs more likely and others less likely to be formed. But he finds this sort of influence

---

<sup>126</sup> Feldman (2001): p. 79.

insufficient to ground moral-analogue evaluations for most, or all, of our beliefs—this requires more direct voluntary control.

One way to combat (GVA) is by providing a positive argument establishing the need for moral-analogue evaluations for beliefs. James Montmarquet and Lorraine Code argue that, in many cases, in order to be responsible for some action one must be responsible for those beliefs upon which the action was based.<sup>127</sup> I think there is much to recommend in Montmarquet's and Code's arguments. However, I will not discuss their merits. Instead, I will argue that there is reason to doubt both 1\* and 2\* of (GVA).

## II. Zagzebski and Doxastic Control

I take it as clear that for the most part we cannot directly voluntarily control what we believe.<sup>128</sup> Alas, I cannot now choose to believe that I am not sitting in front of my computer working on my dissertation. Furthermore, even in more complicated cases of inquiry where we consciously reflect on a body of evidence, “[we] believe, not because on reflection a certain thing seems worthy of belief, seems epistemically valuable, but because in reflecting we become vulnerable in certain ways to beliefs of certain sorts.”<sup>129</sup> Michael Stocker rightly claims that noting “our passivity helps remind us that beliefs are about the world and their truth determined by it, not by us: that the fit is of belief to the world, not the world to belief.”<sup>130</sup>

---

<sup>127</sup> Montmarquet (1993): pp. 1-18. Code (1987): pp. 68-9.

<sup>128</sup> Richard Feldman argues that we sometimes do have voluntary control over what we believe. If our beliefs track the world, and if we have voluntary control over some states of affairs in the world, then we have control over some of our beliefs. If it's up to me whether or not the lights in the room are off or on, I have control over whether I believe the lights in the room are off or on. See Feldman (2001): pp. 77-92. I will ignore Feldman's argument since it does not have the generality needed to ground the notion of epistemic responsibility I am seeking.

<sup>129</sup> Heil (1983): pp. 357.

<sup>130</sup> Stocker (1982): p. 408.

However, Linda Zagzebski argues that this fact should not deter us from employing the notion in epistemic responsibility in one's account of knowledge. In a section titled "The irrelevance of the [voluntarist] objection to virtue theory" she states:

First of all, the question of the voluntariness of beliefs has a much weaker force when set against a virtue theory than against other kinds of moral theory since the purpose of a virtue approach is to shift the locus of evaluation from the belief or act to the virtue or vice or other internal qualities of persons. The primary object of evaluation is inner traits of persons, and the evaluation of acts or beliefs is derivative. No one claims that our moral virtues and vices are under our complete control, but they are generally regarded as sufficiently voluntary to be the proper objects of moral evaluation, including moral praise and blame. The appropriate question to ask about using a virtue approach in epistemology, then, is whether our intellectual virtues and vices are as voluntary as our moral virtues and vices.<sup>131</sup>

Zagzebski's thought seems to be that what one's theory takes as the primary bearers of evaluation is significant for the effectiveness of (GVA). For example, in deontological theories the primary bearers are acts and beliefs, whereas in virtue theories the primary bearers are agents, and in particular their virtues. Since beliefs are involuntary in the way that Alston and others claim then deontological theories will have a hard time explaining why they are evaluated analogously to acts—which are voluntary. By contrast, virtue theories can accept the disanalogy because the evaluation of beliefs is derivative to the evaluation of agents and their virtues, and one's intellectual virtues and vices are no less voluntary than one's moral virtues and vices. Even if beliefs are mostly involuntary their evaluation is derived from something that is voluntary; namely one's intellectual character.

Although Zagzebski rightly notes that one of the distinctive features of virtue theories is having agents and their virtues the primary objects of evaluation, and that as

---

<sup>131</sup> Zagzebski (1996): pp. 59-60.

traditionally conceived moral virtues are acquired through voluntary action, more needs to be said about the relationship between one's intellectual virtues and the formation of one's beliefs. If Alston is right, and most of one's beliefs are not under our voluntary control, then it is not clear how intellectual virtues are formed. What sorts of actions are required to build intellectual courage, or honesty, or open-mindedness? In order to maintain the analogy with moral virtues, intellectual virtues need to be developed by intentional action. Either our beliefs are significantly under our control, and our character is formed by habitual voluntary belief formation, or there is some other sort of control that we display with respect to our beliefs that explains how our intellectual characters are formed. Alston and others seem to have discredited the first option and Zagzebski says little about the relation between one's intellectual virtues and one's beliefs. This is a deficiency in her account. Claiming that one's intellectual virtues are sufficiently under one's voluntary control does not remove the problem if the relationship between virtues and beliefs is left mysterious.

Zagzebski does, however, claim that "beliefs, like acts, arrange themselves on a continuum of degrees of voluntariness, ranging from quite a bit to none at all,"<sup>132</sup> and that epistemologists typically take as the paradigm case of an act those of direct choice—such as choosing to raise one's hand—while the paradigm case of a belief those seemingly completely involuntary—like perceptual beliefs. Zagzebski thinks this unfairly tips the scales against voluntarist theories, and that we are better off acknowledging that beliefs, like acts, come in varying degrees of voluntariness. Appreciating that many acts for which we are morally responsible, and which are voluntary, are not the result of one's direct choice (e.g. the impolite person not being able to withhold from falling asleep in

---

<sup>132</sup> Zagzebski (1996): p. 66.

the presence of company), should dissuade us from requiring that our beliefs be formed by a direct voluntary choice. We should also reject the idea that the paradigmatic case of justified beliefs are perceptual beliefs, and liken them more to sneezings and coughings. Rather, we should focus more on beliefs that result from rational deliberation, since these are more analogous to those acts for which we take one to be responsible for.

Again, I find this response unsatisfactory. Even if epistemologists have incorrectly categorized paradigmatic acts and paradigmatic beliefs, the categorization is not without motivation. Her claim that perceptual beliefs are analogous to involuntary behaviors like sneezes ignores this motivation. Sneezes and twitches are not things that we are held responsible for. In fact these behaviors are traditionally thought to fall outside the class of human actions. They do not receive moral evaluation since they are things that happen to us rather than things we do. However, it seems things are significantly different for perceptual beliefs. The person who believes every optical illusion (who takes the straw to be bent when partially immersed in water, who believes the road is wet on a hot day, etc.) should be reckoned intellectually vicious even under Zagzebski's lights. Or consider the skeptic who believes that things only *appear* to be a certain way but refuses to accept their appearances (if there exists such a person), surely he is overly cautious and untrusting; traits the intellectually virtuous would not possess. By likening perceptual beliefs to sneezes and twitches these intuitions become difficult to accommodate. Even if perceptual beliefs should not be considered a paradigm case of justified belief, they should not be relegated to things outside epistemic evaluation. Indeed any non-skeptical account of justification or knowledge should give most of these beliefs a positive epistemic status. But since perceptual beliefs are typically involuntary

something needs to be said for why they receive such evaluation if it is to mirror that in the ethical domain.

Both of Zagzebski's attempts to deal with (GVA) fail to adequately address its central concern. Given that many of our beliefs escape our direct voluntary control one needs to explain why they are open to moral-like evaluations. Claiming that a belief's evaluation is derivative on one's intellectual character, without explaining the relationship between one's virtues and one's beliefs leaves the problem unresolved. Similarly, claiming that many of the beliefs that epistemologists take as paradigm instances of beliefs with positive epistemic status, should rather be relegated to some class of beliefs outside the scope of epistemic evaluation—like sneezes and twitches in relation to genuine human actions—does not solve the problem but rather demonstrates why the problem is so threatening to the proposed association between ethics and epistemology. In what follows I will argue that there are two main forms of control that are mutually reinforcing and that are sufficient to warrant making moral-analogue evaluations of our beliefs.

### **III. Responsibility and Evaluative Control: A First Step Toward Explaining Doxastic Control.**

One main reason why the second premise in (GVA) is so attractive is that a parallel principle seems widely accepted in the ethics:

(P1) We are morally accountable only for voluntary actions and omissions.

If we are morally responsible only for those things that are under our voluntary control, then we are epistemically responsible only for those things that are under our voluntary control.

However, there are good reasons to worry about the truth of (P1). Robert Adams (1985) articulately presents these reasons, arguing that we are guilty of involuntary sins. His central claim is that we are morally accountable for our states of mind (i.e. our beliefs and desires) even if such states are not the result of our voluntary actions or omissions.

The subject of ethics is how we ought to live; and that is not reducible to what we ought to do or try to do, and what we ought to cause or produce. It includes just as fundamentally what we should be for and against in our hearts, what and how we ought to love and hate. It matters morally what we are for and what we are against, even if we do not have the power to do much for it or against it, and even if it was not by trying that we came to be for it or against it.<sup>133</sup>

Adams considers the beliefs of the graduate of Hitler's *Jugend*, and claims his beliefs are heinous and the fitting for moral reproach even if they were the result of indoctrination and enculturation beyond the graduate's control. In general, people who are self-righteous, easily angered, bigoted, prejudiced, etc. are morally worse (in this respect) than those who are not, regardless of whether these beliefs are the product of their voluntary actions or omissions.

If Adams is right to reject (P1), then a large motivation for 2\* is removed. If we are morally accountable for our beliefs (and desires) without needing voluntary control over them, then perhaps we are analogously epistemically accountable for our beliefs without needing voluntary control over them. This is conclusion Pamela Hieronymi (2006) draws.

Adams and Hieronymi reject the traditional understanding that one is morally responsible only for what is under one's control. But what is the other notion of responsibility employed in their accounts? According to Hieronymi:

---

<sup>133</sup> Adams (1985): p. 12.

To be morally responsible for a thing (whether an attitude, an action, or a state of affairs) is to be open to moral assessment or judgment on account of that thing, to be open to moral praise or blame for it, to be open to what have come to be called the “reactive attitudes” of resentment and indignation on account of it...<sup>134</sup>

The class of reactive attitudes is larger than Hieronymi states here and includes guilt remorse, reproach and blame. And to be morally responsible in the sense just sketched is to be open to any of these attitudes. But before this can be illuminating one has to understand when these ‘reactive attitudes’ are appropriate. In other words, if we do not display voluntary control over what we believe (or feel) why are we open to moral assessment on account of them? The answer to this question is found in examining the nature of these intentional attitudes. According to Hieronymi such attitudes

reveal something of one’s mind, of one’s take on the world and what is important or worthwhile or valuable in it. By settling certain questions for oneself, by having a take on what is true, what is important, and what is to be done, one thereby constitutes those bits of one’s mind relevant to the quality of one’s relations with others—and so establishes what we might call one’s moral personality, or in an older but apt phrase, the quality of one’s will. But that bit of one’s mind—one’s moral personality or one’s will (broadly construed)—just is the object of moral assessment and reaction. It is that which we assess, when we assess whether someone is morally praise- or blameworthy; it is that to which the reactive attitudes react.<sup>135</sup>

Our beliefs (and desires) are our answers to certain questions. We believe *p* when we have settled the question of whether *p*. We desire *q* when we have settled the question of whether *q* is valuable or worthwhile. As a result whenever we form a belief or desire we become answerable to the actual reasons available for the various answers. When such reasons have moral significance then our beliefs also acquire moral significance by comprising our moral characters, and by showing what we are for and against, what we value and disvalue; as Hieronymi puts it they display “the quality of one’s will”. This is

---

<sup>134</sup> Hieronymi (2008): p. 361.

<sup>135</sup> Hieronymi (2008): p. 362.

what makes them the proper subject of moral appraisal even if we do not have voluntary control over them.

Some might object that this conflates being open to moral appraisal and being morally responsible.<sup>136</sup> Such people will concede that the graduate of Hitler's *Jugend* is morally worse than Mother Teresa regardless whether he exhibits doxastic control but deny that the graduate is responsible for his moral standing. This is not implausible. For example, we are all morally imperfect and therefore subject to some measure of negative moral appraisal. But do we want to say that we are some how responsible and to blame for not being morally perfect? It might seem that we don't. Or take an example not involving moral perfection. Moral education presupposes the notion moral improvement; that we are not born as morally good as we can be. Children have vices that need to be eradicated from their characters, and only when properly educated will they become a morally better people. But, again, do we thereby hold them responsible, and blame them, for their moral standing? Again, it may seem that we don't blame them for not being as good as they can be from the get-go; we recognize that moral improvement takes time and practice, and children are not to blame for their current moral standing

But I think that distinguishing between blame/reproach on the one hand and punishment on the other can accommodate these appearances. It seems clear to me that we don't punish people for being less than perfect, or children for not being morally mature. Recognizing that their moral character escapes their voluntary control is a reason to be merciful. But not being deserving of punishment does not entail that one doesn't deserve reproach or blame. In fact, it seems that our practices of moral blaming reveals that we do hold people morally responsible for features of their character that escape their

---

<sup>136</sup> Indeed, I had this worry about this approach in an earlier draft.

voluntary control. When one is insensitive, ungrateful, malicious, vindictive, etc. we think that an apology is in order, and that one should ask for forgiveness. We are not satisfied if the offending party reacts to the offence with, “I’m not to blame, it was beyond my control,” or “That is just how I was raised.” We expect people to own their wrong doings. Moreover, it seems clear that one’s moral imperfections (even those beyond one’s voluntary control) give people reason to think poorly of her. But it seems odd to suggest that though you do not blame someone, you think poorly of her.<sup>137</sup> Of course, it may sometimes be inappropriate to voice one’s reproach when directed at another’s moral faults that escape one’s voluntary control. But there are many reasons for this; one being that voicing one’s reproach is sometimes a form of punishment.

Admittedly, it is difficult to determine when one’s beliefs and other propositional attitudes comprise one’s moral character and when they don’t. Clearly, if I hold a belief because of hypnosis or because of being drugged, then I am not subject to moral blame for them. Such beliefs do not arise from me in the appropriate sort of way needed to reveal *my* moral character and *my* take on the world. But most of our beliefs are not formed in this way and are rather the result of more natural processes and therefore more appropriately represent who we are and what we stand for.

Once we accept that we are rightly open to moral appraisal for our beliefs even though they are not under our voluntary control, it is not difficult to extend this line of reasoning to epistemic evaluations. For not only do our beliefs constitute our moral characters, they also constitute our epistemic characters. They indicate how much we value the true and eschew the false. Accordingly, just as our moral characters are the

---

<sup>137</sup> This observation was made in Adams (1985): p. 21.

proper subjects of moral appraisal, our epistemic characters are the proper subjects of epistemic appraisal—whether or not we exhibit voluntary control over our beliefs.

Does denying voluntary control over our beliefs entail that we have no control over them? I don't think so. First, Heironymi claims that insofar as our beliefs are our “answer to some question(s), we exercise a distinctive form of control over them.... Because these attitudes embody our take on the world, on what is or is not true or important or worthwhile in it, we control them by thinking about the world, about what is or is not true or important or worthwhile in it. Because our minds change as our take on the world changes—because our minds change as we change our minds—we can be said to be “in control” of our commitment-constituted attitudes.”<sup>138</sup> She admits that this form of ‘evaluative control’ is an odd sort of control and that some might plausibly think that this is no form of control at all. However, in what follows I argue that this does not exhaust the sort of control that we exhibit over our beliefs. I claim that this sort of evaluative control opens the door to a more robust form of control that more closely resembles the control we exhibit over our actions.

Before moving on, however, I think it will be instructive to contrast the view just outlined with a similar view that also denies that we are responsible for our beliefs only insofar as we exhibit voluntary control over them.

### *III.i Feldman*

Richard Feldman proposes a solution to (GVA) that also denies 2, claiming that obligations can still exist even if we have no voluntary control over our beliefs.<sup>139</sup> Unlike

---

<sup>138</sup> Heironymi (2008): pp. 370-1.

<sup>139</sup> Of course my account does the same thing. However, I account for claim that responsibility requires some voluntary control by claiming that such control exists with what one accepts. Feldman here claims that obligations can exist even without voluntary control of any sort.

Adams and Hieronymi, Feldman accepts (P1). He, therefore, thinks epistemic obligations are significantly different from moral obligations in that only the latter have application on things we have voluntary control over. Instead, Feldman appeals to role obligations to draw the analogy with epistemic obligations. He claims that role obligations (or role ‘oughts’), unlike moral obligations, are still binding even when one does not have the ability to fulfill them. For example teachers ought to explain things clearly and actors ought to be convincing. His thought is that these ‘oughts’ explain what comprises good performances in these roles. Teachers who fail to teach clearly fail to satisfy a condition for being a *good* teacher. Even if one is simply unable to communicate clearly, if one is in a teaching role, then one fails to fulfill one’s obligation for being a good teacher.

Being intellectual agents are also roles people participate in. And there are certain things that we ought to do in order to be good in this role. For example Feldman proposes that one thing “they ought to do is to follow their evidence (rather than their wishes or fears).”<sup>140</sup> To the extent that we fail to follow this rule we fail to fulfill our obligations for being good intellectual agents, even if we have no control over our beliefs. What it takes to be a good teacher, actor, or intellectual agent does not depend on our abilities, and so even if we have no control over our actions or our beliefs these obligations still exist.

Feldman correctly points out that the requirements for excelling in various roles do not depend on our abilities. However, there is still reason to worry about Feldman’s treatment of (GVA). One important difference between teachers and actors, on the one hand, and intellectual agents, on the other, is that the former roles are typically engaged in voluntarily whereas we have no choice about whether to be intellectual agents.

Feldman recognizes this point stating that “the role of a believer is not one that we have

---

<sup>140</sup> Feldman (2001): p. 88.

any real choice about taking on. It differs in this way from the other roles mentioned.”<sup>141</sup>

This difference is important since it makes it difficult to explain the phenomenon of praising and blaming people for their intellectual behavior. We don't just think that people are responsible for certain activities only if they want to be excellent intellectual agents, but that they are responsible for *actually being* excellent intellectual agents. If people are too hasty, or not thoughtful, or are prone to follow their emotions rather than their evidence, we not only say they have failed in being intellectually excellent, but that they should be blamed for their shortcomings.

Feldman's role norms are what Phil Goggans has identified as 'valuative norms' in that they give the conditions for the value of something relative to the end of performing some role sufficiently well.<sup>142</sup> However, these are different from traditional 'deontic norms' in that we don't necessarily hold subjects accountable for satisfying them. This feature of valuative norms will be more or less worrisome depending on how much similarity we expect our epistemic evaluations to have with our moral evaluations. I have argued that having the notion of epistemic responsibility is important for an adequate account of knowledge, and therefore recommend a stronger connection between our epistemic and moral evaluations than Feldman's appeal to role norms allows. In fact Feldman's likening epistemic norms to role norms can be happily embraced by many of the externalist theories that Feldman hopes to part ways with.

I have surveyed two attempts to respond to (GVA) by rejecting 5. I have defended the first attempt that justifies the rejection by maintaining a strong analogy between epistemic and moral evaluation, and arguing that one can morally responsible for one's

---

<sup>141</sup> Ibid.: p. 88.

<sup>142</sup> Goggans (1991): pp. 103-5.

beliefs despite lacking voluntary control over them. The second attempt justifies the rejection by parting ways with the moral analogy and claiming that there other kinds of norms—role norms, or ‘valuative norms’—that can be properly applied even when we lack voluntary control. Unfortunately, this move abandons the main motivation for appealing to notions of epistemic responsibility that gave the voluntarism argument its significance.

#### **IV. Reflective Control: A Second Step Toward Explaining Doxastic Control**

Many authors have recognized that although our beliefs are not subject to our will in the way our actions are, we still exercise a form of reflective control over them. As reflective beings our mental lives are not consumed by the external world; we have the ability to turn our attention inward and evaluate our mental states. What is relevant for this discussion is that we have the ability to examine our beliefs and hold them up to our epistemic standards. We can often determine whether what we believe is inconsistent with our other beliefs, or whether it is sufficiently supported by the evidence. If, upon examination, we discover that our beliefs do not meet our standards for them, then we evaluate them negatively. And it is here where we have some control over our mental lives, since some of our mental states are sensitive to how we judge them; they are what Scanlon (1998) calls ‘judgment-sensitive attitudes’.

These are attitudes that an ideally rational person would come to have whenever that person judged there to be sufficient reason for them and that would, in an ideally rational person, ‘extinguish’ when that person judged them not to be supported by reasons of the appropriate kind.<sup>143</sup>

Our beliefs are one kind of these ‘judgment-sensitive attitudes’. But of course most of our beliefs are formed without such reflection. Before writing this chapter I was

---

<sup>143</sup> Scanlon (1998): p. 20.

aware of all the first-order reasons indicating that I was sitting in front of a computer, and as a result formed the corresponding belief. And although I did not reflect on whether these reasons were sufficient to justify the belief, surely the belief was still justified and rational—something I ought to believe. Therefore, it cannot be that our beliefs must have undergone reflective evaluation before they can be open to epistemic appraisal. It is sufficient that I have the ability to reflect on my reasons for my belief and that such reflection would motivate a revision if my evaluation is negative. As Pettit and Smith note:

Responsible believers and desirers are orthonomous subjects, in the sense that they recognize certain yardsticks of right belief and right desire and can respond to the demands of the right in their own case. They may vary among themselves in how far they actually conform their beliefs and desires to those yardsticks; they may be more or less thoroughly ruled by the right. But outside the domain of disabling obstacles, they are all equally orthonomous in at least this sense; they are all able to answer the call that the right makes upon them.<sup>144</sup>

To hold a belief or desire freely is to hold it in the presence of an ability, should the belief or desire be wrong, to get it right.<sup>145</sup>

This does not mean that there are not cases where we are unable to believe what we know has overriding evidential support, or where we are unable to disbelieve what we know is irrational or unjustified. One should not deny the occurrence of incontinent beliefs.

However, one who holds a belief despite giving it a negative epistemic evaluation, will thereby believe irrationally.

It is important to note that one's standard for epistemic rightness need not be very articulate or sophisticated. And one's recognition of whether one's belief satisfies one's standard need not involve a precise itemized list of all the reasons and a clear calculation

---

<sup>144</sup> Pettit and Smith (1996): p. 442.

<sup>145</sup> *Ibid.*: p. 445.

of their relative weights. Recognizing that there are certain norms that constrain belief formation and determining whether one's belief has sufficiently satisfied such norms may be nothing more than a general seeming that one's belief is appropriate (inappropriate) for the given circumstances and given the evidence in its favor. One need not be an epistemologist nor even cognitively mature to meet these requirements and thereby have control over one's belief.

Clearly, this sort of control and freedom over our mental lives is significantly different from the standard conception of freedom. Under the standard conception one's belief is free only if, for whatever one believes, one had the ability to believe otherwise. Under the current conception one has freedom of belief to the extent that one can conform one's doxastic life to an appropriate standard of epistemic justifiedness. But this more restricted notion of control may amount to nothing more than the evaluative control discussed above. There are surely times when we voluntarily decide to be reflective and examine our beliefs. This is common practice in the philosophy classroom. But there are also times when we seem compelled to turn reflective. Sometimes upon forming a belief we are struck with the uncomfortable feeling that our doxastic structure may not be perfectly consistent, and we instinctively examine which of our beliefs are justified and which are not. Moreover, the standards that we apply to our beliefs, and our determination of whether our beliefs satisfy this standard seem no more up to us than the first order beliefs which are being examined.

I don't take this as a problem for those hoping to explain how we can be responsible for our beliefs. I think that sort of evaluative control discussed in section II is

sufficient to establish that. However, this sort of control is relevant in that it provides the resources for how we can exhibit a more robust kind of control over our beliefs.

### **V. Monitoring our Doxastic Processes and Faculties: A Final Step toward Explaining Doxastic Control**

Having reflective control allows us to evaluate whether our beliefs are justified or unjustified, but it also allows us to evaluate whether our belief forming processes and faculties are reliable at reaching the truth. Indeed, often we are unable to evaluate our beliefs without evaluating those processes and faculties from which the belief was formed. Our ability to take a critical stance towards our belief forming practices and faculties grounds another way in which we exhibit control over our doxastic lives. For if we determine that our belief forming practices and faculties are unreliable we then have the ability to modify and regulate our trust of these processes and faculties. Sometimes this will come by evaluating the faculties themselves and determining whether they are generally reliable, or determining in which situations a particular faculty is reliable or unreliable. We learn to not trust our eyesight for telling us whether something is bent or straight, if the object we are looking at is partially immersed in water. We learn not to trust a testifier if she fails to make eye-contact, seems nervous, and stutters. This sort of regulation will be available for most of our belief forming faculties, and once we recognize the situations in which our faculties are unreliable in reaching the truth we will then be able to attenuate our trust of such faculties in those situations.

However, we do not just evaluate our belief forming faculties but also our approach to given inquiry. For a given investigation we can determine whether we properly gathered evidence, gave adequate care and attention to the possessed evidence,

calculated and inferred too hastily, were overly trusting or dogmatic, etc. In discussing the extent of control that we exhibit over our beliefs John Locke states:

All that is voluntary in our knowledge is the employing or withholding of any of our faculties from this or that sort of objects, and a more or less accurate survey of them: but they being employed, our will hath no power to determine the knowledge of the mind one way or the other; that is done only by the objects themselves, as far as they are clearly discovered.<sup>146</sup>

In a much more recent assessment John Heil states:

“[it] is still possible to speak of epistemic responsibility and agency ... if one focuses, not on the ways (if indeed there are any) in which agents select beliefs, but on the ways in which they select belief-generating procedures. It is here that talk of doxastic agency appears to have its most obvious application.”<sup>147</sup>

I am mostly in agreement with Locke and Heil here. Even if we do not have direct voluntary control over my beliefs in the way that we typically have over our actions, our reflective abilities give us control over how we approach our belief forming faculties. And how we employ this sort of control will have a significant impact on what we end up believing. Therefore not only can we be held responsible for these intellectual activities, we can also be held responsible for the beliefs that arise from such activities.

That we can legitimately apply moral-analogue evaluations to beliefs even though we exhibit control only at a stage before, but heavily influential to, the formation of the belief, is buttressed by the fact that something similar may be true of many actions that are subject to moral evaluations. Over the last 50 years there has been abundant research suggesting that many of our morally significant actions are more heavily influenced by situational pressures than by our personal traits or characters.<sup>148</sup> For example, one study

---

<sup>146</sup> Locke (1975): p. 650-1.

<sup>147</sup> Heil (1983): p. 363.

<sup>148</sup> For a good discussion of the most relevant ‘situationalist’ research see Doris (2002).

found that people coming out of a phone booth are more likely to help a stranger pick up a folder full of papers if they found a dime in the coin slot (14 of 16 subjects) than if they didn't (1 of 25 subjects). Moreover, this feature was a better predictor of future behavior than any other general, cross-situational feature. In another study a high proportion of people were found easily manipulated to inflict what they took to be intense pain on another individual by an experimenter telling them that they 'must' continue administering higher levels of electrical shock.

Some moral psychologists and philosophers find the research above threatening to the viability of virtue ethics as a legitimate normative theory in ethics. Gil Harman has even suggested that these findings cast doubt on the idea that we have characters of any kind.<sup>149</sup> I find the arguments for these conclusions less than compelling. However, the data does suggest that we may have less control over some of our morally relevant behaviors than previously thought. It suggests that our voluntary control only extends over which situations we place ourselves rather than over our responses to those situations. If this is right, then John Doris correctly suggests, "[the] way to get things right more often...is by attending to the determinative features of situations. We should try, so far as we are able, to avoid 'near occasions for sin'—morally dangerous circumstances. At the same time, we should seek near occasions for happier behaviors—situations conducive to ethically desirable conduct. This means that the determinants of moral success or failure may emerge earlier in an activity than we might think."<sup>150</sup>

But this fact should not compromise our responsibility for our actions. Even if we can best influence our morally relevant behaviors via our control over what situations we

---

<sup>149</sup> Harman (2000).

<sup>150</sup> Doris (1998): p. 517.

place ourselves in, we are still responsible for not only our being in those situations but also for the resulting behaviors we display. And now the similarity with beliefs becomes evident. Even though we lack significant control over how our cognitive characters react to a given body of evidence, we do exhibit significant control over whether we place ourselves in epistemically dangerous or healthy circumstances. And just like with actions we become responsible for our beliefs in virtue of being responsible for those behaviors that heavily influence belief formation. If we are responsible in avoiding epistemically dangerous circumstances, then our resulting beliefs will receive positive epistemic evaluation, and if we are irresponsible then the resulting beliefs will receive negative epistemic evaluation. As Michael Stocker observes

In these and in so many other cases, all we can do is put ourselves in as good a position as possible to intervene successfully at an opportune moment, and then make our ‘contribution’ to the world. Success may come (and with it may also come activity and responsibility). But it comes without certainty. As with beliefs, so with many physical acts and conditions: we must have an adequate degree of humility.<sup>151</sup>

But it is important to note that limiting our moral (and epistemic) responsibility to the situations we put ourselves in does not account for the full scope of what we are responsible for and is not supported by the sort of studies mentioned above. It is important that in none of the studies conducted, was the behavior invariant across subjects. Although there is a strong correlation between the kind of behavior displayed and the situational pressure exhibited on the subject, some people were able to overcome the situational pressure and act against the moral grain. This suggests that even when placed in a given situation, one’s moral character does have a significant role in explaining one’s behavior. If the context is fixed what else is supposed to explain the

---

<sup>151</sup> Stocker (1982): p. 410.

variant behavior? I believe that the lesson from these studies is not that our characters do not influence our behavior but rather that our behavior is more susceptible to situational pressure than we would like to admit. But this is perfectly consistent with the idea that we are responsible for the continued development of our moral characters as well as the behavior that emanates from them. We can, and ought to, strive to have robust characters that are not easily swayed by morally irrelevant features of our environment. Failing to develop such a character makes one open to moral blame and reproach.

Things should be no different in the epistemic realm. Once in an epistemic context we are moved by what we take to be the reasons supporting some proposition. We find ourselves compelled to believe what is best supported by the evidence. But this does not mean that we are not able to positively shape our epistemic characters, or that we are not responsible for such beliefs. For we can still evaluate our practices and faculties, and modify our intellectual characters based on what we determine is a reliable or unreliable approach to the truth. If we fail to adequately modify our approach to our various doxastic practices and faculties—e.g. failing to attenuate one’s trust of those faculties that have proved themselves to be unreliable, or failing to give proper attention to the evidence, allow ourselves to be overcome with emotion and bias—then we will be subject to epistemic blame and reproach. This is so even if when in the particular situation one does not have the ability to believe otherwise.

### **VI. Conclusion: A Supplement to Zagzebski**

By connecting our intellectual characters with our reflective control, we now can see how our intellectual characters help ground our epistemic evaluations despite our beliefs being outside the control of our will—something that Zagzebski was unable to provide. We don’t voluntarily choose our beliefs, in the way we choose our actions; as a

result our intellectual characters are not developed in the same way as our moral characters. Instead of being developed by repeated voluntary behavior, our intellectual characters are formed by taking a reflective perspective on our beliefs and on one's belief forming processes. We are able to evaluate how we engage our faculties and if we determine this to be an unreliable way to the truth then we have the ability to modify and regulate our approach.

We can now also accommodate the intuition that our perceptual beliefs are significantly different from twitches and sneezes. Perceptual beliefs, like all of our beliefs, are the result of our intellectual characters. For most of us we have been given very little reason to doubt that our perceptual faculties are generally unreliable and so have had no reason to attenuate our general trust on them. However, there are specific contexts in which we have learned that trusting our perceptual faculties is not a reliable guide to truth, and most of us have regulated our intellectual practices accordingly. Those of us who fail to regulate in this way will be subject to error and will be blameworthy for it.

I conclude that the voluntarism argument against moral-analogue evaluations for beliefs fails. Although for the most part we do not have voluntary control over our beliefs, such control is not needed for epistemic (or moral) responsibility. Moreover, we do exhibit significant control over our belief forming practices and over our intellectual investigations. This sort of control is made possible because of our ability for higher-order reflection. And even though we may not be able to voluntarily determine what particular belief we adopt, this may be no less control than we exhibit in many of our morally significant actions.

## **Chapter 5: Testimonial Knowledge in Children: I Can Trust You Now ... But Not Later.**

*When I was a child, I used to speak like a child, think like a child, reason like a child; when I became a man, I did away with childish things. 1 Corinthians 13:11*

### **I. Introduction**

If we think of all the things we know, it will not take long to realize that much (most?) of what we know depends, partially, or entirely, on testimony. Our political, social, medical, scientific, and geographical information largely relies on testimony. And this is not all. We learn about our more immediate surroundings from schedules, menus, traffic signs, business cards and flyers. These beliefs are also testimony-based.

That testimony is an indispensable medium for knowledge in our adult lives is indisputable. But our reliance on testimony was even more dramatic during the first stages of life. As Thomas Reid points out:

The wise author of nature hath planted in the human mind a propensity to rely upon human testimony before we can give a reason for doing so. This, indeed, puts our judgments almost entirely in the power of those who are about us in the first period of life; but this is necessary both to our preservation and to our improvement. If children were so framed as to pay no regard to testimony or authority, they must, in the literal sense, perish for lack of knowledge.<sup>152</sup>

A child's reliance on testimony provides more than a practical advantage; it also provides an epistemic advantage. Children learn and come to know things about the world at a very young age through the testimony of their caregivers. All this seems intuitive and well and good.

The challenge comes in trying to explain *how* children acquire such knowledge. The reason for the challenge is that children come to know things about the world from

---

<sup>152</sup> Reid (1983): p. 281.

testimony despite having gullible characters, and most accounts of knowledge require that a belief be reliably formed. Since children indiscriminately receive testimony, their testimony-based beliefs seem unreliable, and, consequently, should fail to qualify as knowledge. Greco formulates the problem as a triad of inconsistent propositions:

1. Young children can learn from the testimony of their caregivers; i.e. they can come to know through such testimony.
2. Testimonial knowledge requires a reliable consumer of testimony; i.e. the hearer can reliably discriminate between reliable and unreliable sources of testimony.
3. Young children are not reliable consumers of testimony.<sup>153</sup>

It seems that if we want to retain the intuition that children do acquire testimonial knowledge, then we either have to reject either 2 or 3. In this chapter I discuss some attempted solutions by Sandy Goldberg and John Greco that reject 3. I argue that their solutions fail. I go on to suggest that what generates the problem is a hidden assumption supporting 2, that the standards for testimonial knowledge should be invariant between children and cognitively mature adults. I propose that in order to adequately explain how children acquire testimonial knowledge we should reject this hidden assumption. I then argue that understanding knowledge in terms of intellectual skills gives us a plausible framework to do so.

## **II. Reliability and Testimonial Knowledge**

What is required for testimony-based knowledge? Does the testifier need to know (or even believe) the proposition testified? It is often assumed that testimony, like memory, does not generate knowledge—it only transmits knowledge. In order for a hearer to know the received proposition the speaker must know the delivered proposition. Jennifer Lackey has recently challenged this “bucket-brigade” conception of testimony, and

---

<sup>153</sup> Greco (2008): p. 335.

argues that testimony can generate new knowledge.<sup>154</sup> I will not address this very important issue regarding the transmission of epistemic properties via testimony.

However, it is more certain that testimonial knowledge requires the speaker be reliable at communicating the truth. Someone disposed to lying or asserting what is unjustified or false will not transmit testimonial knowledge. Whether or not a speaker needs to be a reliable *believer*, it seems clear that she needs to be a reliable *testifier*.

What about the hearer? Suppose that the speaker knows *p*, and is also a conscientious and reliable testifier. Is this sufficient for those who come to believe *p* on the basis of her testimony to acquire knowledge? The answer seems to be no. Consider the following story presented in Lackey (2008):

**Compulsively Trusting:** Bill is a compulsively trusting person with respect to the testimony of his neighbor, Jill, in whom he has an obsessive romantic interest. Not only does he always trust Jill when he has very good reason to believe her, but he is incapable of distrusting her when he has very good reason to not believe her.... Indeed, Bill is such that there is no amount of evidence that would convince him to not trust Jill. Yesterday, [Jill told Bill] that she had seen an orca while boating earlier that day. Bill, of course, readily accepted Jill's testimony. It turns out that Jill did in fact see an orca whale..., that she is very reliable with respect to her epistemic practices, both in general and in this particular instance, and that Bill has [very little reason] to doubt the proffered testimony. Given his compulsively trusting nature with respect to Jill, however, even if he had had massive amounts of evidence available to him indicating, for instance, that Jill did not see an orca whale, that she is an unreliable epistemic agent, that she is an unreliable testifier,... Bill would have just as readily accepted Jill's testimony.<sup>155</sup>

---

<sup>154</sup> Lackey (2008). She also argues that memory can also generate knowledge.

<sup>155</sup> Lackey (2008): p. 66. I think it is important that Bill has some, but very little reason to doubt Jill's testimony. Lackey's original case claims that Bill has *no* reason to doubt Jill. The reason for this change is that I think Bill can still come to know things despite being evidentially insensitive. For example, if Jill tells Bill that she had toast for breakfast, or that she sleeps on a queen sized bed, and if Bill has absolutely no reason to doubt this testimony, then I take it that Bill does come to know these facts about Jill. The cases I focus on are such that Bill has some reason, though normally not overriding reason to doubt Jill. In the case above I take it as somewhat surprising that Jill saw a whale, though under normal circumstances, and with normal recipients of testimony, this would not be enough to justify doubt the proffered testimony. This slight modification to Lackey's case should not have serious effects to the main argument in this chapter. Thanks to Alex Jackson for pushing me to address this issue.

According to Lackey, Bill's inability to reliably monitor for potential defeaters makes him an unreliable recipient of Jill's testimony, and consequently, "Bill's belief that there was an orca whale in the relevant body of water is *evidentially insensitive* in a way that is clearly incompatible with justification, warrant, and knowledge."<sup>156</sup> I am inclined to agree with Lackey's intuition that Bill fails to acquire knowledge.

Goldberg presents a similar case in which the lack of testimonial knowledge is even more apparent:

**Gullible in a Room Full of Liars:** Sid is gullible in the extreme: he accepts anything anyone says merely in virtue of the fact that someone said so. Sid is in a room full of inveterate liars. He immediately and uncritically believes everything each of them says. At one point he happens to bump into Nancy, the only reliable person around. Nancy reliably tells Sid that p, and (as a matter of course) he believes her.<sup>157</sup>

Sid is epistemically worse off than Bill for at least a couple of reasons. First, Sid's inability to reliably discriminate true from false testimony extends beyond just one individual. Bill accepts whatever Jill tells him, but Sid accepts whatever *anyone* tells him. Second, Sid's local environment is not suitable for testimonial knowledge. Because liars surround Sid, his bumping into Nancy—a reliable testifier—is a stroke of good luck. As a result, his testimony-based belief is epistemically perniciously unreliable. Sid's testimonial environment is equivalent of Goldman's 'Fake Barn County'. As a result of these two features, Sid's clearly fails to acquire knowledge.

Neither Sid nor Bill have knowledge because it seems that testimony-based knowledge typically requires not only that the testifier be a reliable at communicating

---

<sup>156</sup> Ibid.: p. 67.

<sup>157</sup> Goldberg (2008): p. 8.

true propositions, but also that the recipient be able to reliably discriminate between true and false testimony. Both Sid and Bill fail to satisfy the latter condition.

### III. Gullible Children

Although there are an increasing number of empirical studies suggesting that children are not as uncritical of other's assertions as many have supposed, there are numerous studies demonstrating that young children are unable to reliably discriminate true from false testimony.<sup>158</sup> At first these findings were used to put pressure on theories of justification requiring that one possess positive reasons for taking one's testifier (or perhaps testimony in general, or perhaps most of the testifiers one has encountered) to be reliable or truthful. Since children are unable to appreciate such reasons, then these theories have the unwelcome consequence that children's testimony-based beliefs are unjustified. Jennifer Lackey (2008), however, provides persuasive arguments that: (1) if the problem of cognitively immature children/infants is a genuine problem for theories requiring the possession positive reasons, then it is also a problem for theories that do not have that requirement; and (2) that the problem of cognitively immature children/infants is not a genuine problem for theories requiring the possession of positive reasons. Again, I will not survey the details of her arguments in this paper because there is a more debilitating problem for theories of testimonial justification or knowledge that Lackey does not address.

Whether or not testimonial justification requires positive reasons, **Compulsively Trusting** and **Gullible in a Room Full of Liars** seem to demonstrate that testimonial justification—if it is going to be relevant for testimonial knowledge—requires that the

---

<sup>158</sup> Goldberg (2008) cites a large number of such studies.

hearer of testimony be reliable in discriminating true from false testimony. Now consider two analogous cases:

**Babe with Mother:** Babe is a compulsively trusting child with respect to the testimony of his mother. Not only does he always trust mother when he has very good reason to believe her, but he is typically incapable of distrusting her when he has very good reason to not believe her. Yesterday, Babe's mother told him that she had seen a whale boating earlier that day. Babe, of course, readily accepted his mother's testimony. It turns out that his mother did see a whale, that she is very reliable with respect to her epistemic practices, both in general and in this particular instance, and that Babe has very little reason to doubt the proffered testimony. Given his compulsively trusting nature with respect to his mother, however, even if he had had evidence available to him indicating, for instance, that his mother did not see an orca whale, that she is an unreliable epistemic agent, that she is an unreliable testifier, etc., Babe would have just as readily accepted his mother's testimony.

**Babe in Preschool:** Babe is gullible in the extreme: like other preschool children, he accepts almost anything anyone says merely in virtue of the fact that someone said so. Babe is told by his preschool teacher that frogs eat bugs, and this testimony is both true and reliable. However, Babe is surrounded by his preschool classmates, who are notoriously unreliable reporters. One of them, who happens to be confused on the issue, could easily have reported that frogs never eat bugs, and Babe would have believed this.<sup>159</sup>

These two cases seem to resemble **Compulsively Trusting** and **Gullible in a Room Full of Liars** in all epistemically relevant respects. In the first case Babe is compulsively trusting of his mom, in just the way that Bill is compulsively trusting of Jill. In the second case Babe is not only generally gullible, but is in an environment where he could have very easily received false testimony. What makes these similarities particularly problematic for theories of testimonial justification and knowledge is that it is generally acknowledged that in the last two cases Babe's testimony-based belief qualifies as knowledge. The description of Babe's cognitive abilities is not far off from what most young children actually possess, so unless we are willing to claim that young children generally do not acquire testimonial knowledge—particularly from the testimony of their

---

<sup>159</sup> Greco (2008): pp. 336-7.

parents and caregivers—then there must be something that distinguishes these last two cases from the first two and explains why we are willing to countenance Babe as knowing but not Bill or Sid.

In the next two sections I will look at two attempts—one by Sandy Goldberg and one by John Greco—to provide such an explanation. Both Goldberg and Greco attempt to resolve the problem of indiscriminating children by suggesting that certain social features in a child’s environment qualify them as sufficiently reliable consumers of testimony for their beliefs based on the testimony of their caregivers to qualify as instances of knowledge.

#### **IV. Goldberg and Epistemic Monitoring**

Goldberg accepts an externalist condition on knowledge, and claims that an individual gains testimonial knowledge only if the following condition holds:

(RL\*) In most or all of the nearby worlds in which S forms the testimonial belief that p, p.

He also claims that whether young children satisfy this condition regarding their testimony-based beliefs will be highly dependent on their environment. Because young children are typically uncritical and overly credulous with other people’s testimony, (RL\*) will be satisfied only if the child is in an environment in which she will encounter only—or mostly—reliable testimony.

The problem, again, is that there are cases where we intuitively regard a child as knowing her testimony-based belief—particularly when the testifier is the child’s parent, teacher, or caregiver—despite the child being in an environment where unreliable

testifiers are present or nearby. **Babe in Preschool** seems to be such a case, and not out of the ordinary for many young children.

Goldberg claims that if we want to maintain an externalist condition—like (RL\*)—on knowledge and that children acquire testimony-based knowledge despite being indiscriminate recipients of testimony, then we should “reconceive the process eventuating in the child’s testimonial belief.”<sup>160</sup> He proposes the following analogy to help make his point:

Samantha is a teenager who has just received her drivers’ permit, which allows her to drive during the day, in good weather, and when in the company of an experienced driver. If we consider Samantha herself, independent of these restrictions on her driving, she is not a particularly safe driver (yet)... At the same time, when she drives under the restricted conditions set forth by her permit—and these are the only conditions under which she drives!—she drives very safely.... Under these conditions, it would be correct to say that her driving is safe to the extent that her experienced cohort is vigilant in policing her. This is so for two reasons: first, the occasions on which she confronts the sort of conditions that would elicit her driving flaws are minimized...; and second, her experienced cohort serves as an external reminder, guiding Samantha through difficulties and correcting Samantha’s driving errors as soon as they arise....<sup>161</sup>

Although Samantha is not a safe driver by herself, she is once we take into account her environment and social context. Similarly, Goldberg claims that, similarly, although a young child is not a reliable receiver of testimony by herself, if we take into account her social environment when evaluating the process involved in her consumption of testimony—in particular the role that her adult guardians play in monitoring her consumption—she is reliable.

Goldberg argues that adult guardians provide “(i) access-restriction, (ii) pro-active monitoring for credibility, and (iii) reactive monitoring for credibility,” for their

---

<sup>160</sup> Goldberg (2008): p. 17.

<sup>161</sup> Ibid.: p. 17.

children.<sup>162</sup> They shield children from those who might be unreliable testifiers, continually filter incoming testimony, and actively screen a child's beliefs by correcting false beliefs that may have crept in to a child's web of beliefs by way of false or unreliable testimony. Goldberg, therefore concludes, that "the roles adult guardians play ... enhance the reliability of a good many of the beliefs that are elicited by the child's encounters with testimony."<sup>163</sup>

When we look at the original triad of inconsistent propositions Goldberg's move is tantamount to a rejection of

3. Young children are not reliable consumers of testimony.

What Goldberg accepts is

- 3'. Young children are not reliable consumers of testimony by themselves.

But this is not inconsistent with 1 and 2 since we can still accept

4. Young children are reliable consumers of testimony when considered within their social environment.

#### *IV.i. Objection to Goldberg*

Whether appealing to a child's social environment succeeds in addressing the current problem depends on whether all the cases in which a child acquires testimonial knowledge are also cases in which the child's guardians play the epistemic roles in monitoring the child's consumption of testimony that Goldberg propounds. In other words, all cases in which 1 is true should be the cases in which 4 is true. But, unfortunately for Goldberg this equivalence does not seem to hold. There seem to be clear cases where a child's testimony-based belief qualifies as knowledge despite her

---

<sup>162</sup> Ibid.: p. 18.

<sup>163</sup> Ibid.: p. 19.

caregivers not providing the epistemic monitoring that Goldberg claims is necessary for testimonial knowledge.<sup>164</sup>

Consider again **Babe in Preschool**. It seems true that typically preschool teachers do provide access-restriction, pro-active monitoring for credibility, and reactive monitoring for credibility for their students. But this is not always the case. Sometimes teachers don't feel well and simply go through the motions in teaching the day's material, without giving much attention to the children they supervise. Should we deny Babe knowledge of those beliefs he forms on the basis of his teacher's assertions simply because his teacher was not monitoring the other students as carefully as she should, or could, be? Moreover, as class sizes grow teachers become less able to provide the kind of epistemic monitoring Goldberg describes. Suppose that within a class of thirty students, a teacher is able to monitor ten to fifteen students at a given time. Surely we do not want to preclude the other half of class from knowing those things that their teacher asserts simply in virtue of the class size preventing the teacher from effectively monitor their consumption of testimony.

Goldberg's analogy with Samantha fails because in order for Samantha to be deemed a good driver it has to be the case that she only drives in the restricting conditions that the permit places on her. If she fails to heed those restrictions and drives without supervision, and in less than optimal conditions, then she is no longer a safe driver. Even if she *usually* obeys the restrictions, this is not enough to change our evaluation of her driving when she fails to obey them. If our evaluation of Samantha's driving is supposed to be analogous to our evaluation of a young child's consumption of testimony, then we should always withhold knowledge when a child believes another's

---

<sup>164</sup> This objection is also raised by Greco (2008).

assertions without the active monitoring of her adult guardian. But as I've tried to show above, we don't.

Goldberg recognizes this consequence of his view but claims that it is in *accord* with our intuitions, not in *conflict* with them. He claims that it helps provide an easy explanation for why children fail to acquire knowledge from reliable testifiers in poorly-monitored contexts with other unreliable testifiers near by. Consider the following case:

**Babe in Used Car Lot:** Babe is taken by his parents to a used car lot. As his parents are engaged by a salesperson, he wanders off by himself and overhears someone saying of a particular car that it was owned by a grandmother. Credulous as ever, Babe accepts this testimony, and so comes to believe of that car that it was owned by a grandmother. It just so happens that the testimony was reliable (the speaker was the grandson of the car owner.).

Goldberg's theory rightly claims that Babe, in this case, does not know that the car was owned by a grandmother. However, given that it does not seem to yield the right answers in other cases (e.g. **Babe in Preschool**, in which the class size is sufficiently large), it seems it still will not do as an adequate theory of testimonial knowledge.

One challenge this presents, however, is explaining why Babe acquires knowledge in his preschool class, but not at the used car lot, despite lacking active epistemic monitoring from his adult guardians in both cases? What is the epistemically relevant difference between these two cases? One difference is that at preschool, but not in the used car lot, Babe bases his belief on the testimony of a socially approved source of information, namely his teacher. John Greco identifies this difference and claims it is the key to explaining how children acquire testimonial knowledge despite being indiscriminate consumers of testimony. I now turn to his proposed solution to our problem.

## V. Greco and Socially Practices of Information Sharing

Greco's approaches the problem from a virtue theoretic conception of knowledge:

(AR) S knows p only if (i) p is true; (ii) S's believing p is the result of dispositions that S manifests when S is thinking conscientiously; (iii) such dispositions make S reliable in the present conditions, with respect to p.<sup>165</sup>

The dispositions that Greco refers to in (ii) are an individual's cognitive virtues or abilities, which are understood as follows:

(CA) S has a cognitive ability A(C, F) with respect to propositions p if and only if there is a field of propositions F and a set of conditions C such that (i) p is in F, and (ii) across the range of close possible worlds where S is in C, S has a high rate of success with respect to believing correctly about propositions in F.<sup>166</sup>

Determining whether an individual, S, knows some proposition, p, requires first determining the mechanism or process S uses in coming to believe p, and then whether that mechanism or process is reliable in his current conditions, C, for forming true beliefs in the field of propositions, F, of which p is a member.

With this framework in mind Greco takes the problem of explaining how children acquire testimonial knowledge from their caregivers to be nothing more than the problem of identifying the relevant mechanism or process that a child uses in forming their testimony-based belief. In **Babe with Mother**, Babe forms his belief on the basis of *testimony from his mom*, and in **Babe in Preschool**, he forms his belief on the basis of *testimony from his teacher*, but in **Babe in Used Car Lot**, he forms his belief on the basis of *testimony from a stranger in a used car lot*, or something relevantly similar. Are these processes reliable, in Babe's circumstances, for forming true beliefs of the kind Babe forms? Consider the following table:

---

<sup>165</sup> Greco (2000): p. 218.

<sup>166</sup> Ibid.: p. 216.

Process Type	Conditions	Field of Proposition	Reliable?
Trusting testimony from mom	Lacking capacity to discriminate reliable from unreliable testifiers	Proposition about what mom saw in the ocean	Yes
Trusting testimony of teacher	Lacking capacity to discriminate reliable from unreliable testifiers and unreliable testifiers nearby	Proposition about what frogs eat	Yes
Trusting a stranger in a used car lot	Lacking capacity to discriminate reliable from unreliable testifiers and unreliable testifiers nearby	Proposition about who was the previous owner of a car	No

We can see that the process that Babe uses in **Babe with Mother** and **Babe in Preschool** are sufficiently reliable to yield true beliefs in the conditions that Babe is in and for the field of propositions of which his belief is a member. Accordingly, we countenance his belief in these cases as knowledge. This is not so in **Babe in Used Car Lot**. Because there are other strangers in the used car lot that are unreliable testifiers, Babe could easily go wrong by trusting a stranger's testimony about the previous owner of a car. Accordingly we don't countenance his belief in this case as knowledge.

One, now, might wonder whether there is a principled reason for identifying the processes as such? Recall that the testifier in **Babe in Used Car Lot** is reliable—he is the grandson of the previous owner. Why isn't the relevant process type in this case, *testimony from the grandson of the previous owner*? Greco provides the following response:

[First], our concept of knowledge answers to our interests and purposes as social, information-sharing being. In order to do so, it must track “fields and circumstances”

that are relevant to those interests and purposes, i.e. fields and circumstances that cut up our cognitive activities at useful joints....

[Second], our “interests and purposes” as information-sharing beings include our interest in identifying reliable sources of information. More specifically, they include our interest in identifying secondary sources of information, or those sources that facilitate reliable information flow. Wedding these two lines of thought, we get the following result: Relevant levels of generality are at least partly a function of our interest in identifying reliable secondary sources of information.

To complete the argument we need one more plausible idea: that reliable secondary sources are (at least sometimes) specified by virtue of their social status. That is, one way we identify people as reliable sources of information is by virtue of their social roles as parent, nanny, teacher, expert..<sup>167</sup>

According to Greco, the relevant process or mechanism for evaluating a belief is determined in part by our interests as information sharing beings. Such interests include identifying reliable sources of information, and sometimes we identify such sources by virtue of their social roles, e.g. parents, teachers and experts. But, of course, it isn't enough for the testifier to satisfy the description of one of these socially accepted roles. The testifier in **Babe in Used Car Lot** is a *grandson*, and for all we know he may be someone's *parent*, or *teacher*. What is important for identifying the relevant process in a given situation is the social practices that are being employed and whether these practices are recognized as reliable means for information sharing. Since Babe and the stranger he overhears are not participating in any recognized social practice (such as parent-to-child or teacher-to-pupil communication), we don't identify the relevant type in **Babe in Used Car Lot** according to any of the accepted social roles that the stranger actualizes.

With this we have a possible explanation for how children acquire testimonial knowledge within a general theory. When children consume the testimony of their parents or caregivers they are participating in a socially accepted practice of reliable

---

<sup>167</sup> Greco (2008): p. 346.

information sharing. This helps us identify the relevant type by which we evaluate their testimony-based beliefs according to the field of propositions of which they are members and the conditions in which they are formed. Once having taken into account these other parameters (i.e. the type of proposition communicated, and the relevant conditions), if the relevant process is reliable in reaching the truth, then the child acquires knowledge.

*V.i. Objections to Greco*

Although I find Greco's attempted solution attractive in many ways, I will argue that it is unsatisfactory. Recall the first two cases that were supposed to buttress the intuition that testimonial knowledge requires reliable consumption: **Compulsively Trusting** and **Gullible in a Room of Liars**. In these cases we deny Bill and Sid knowledge because they are insufficiently reliable consumers of testimony. They are insensitive to defeaters and incapable of properly monitoring the delivered testimony. At first glance it appears that Greco can easily explain why Bill and Sid lack knowledge. Since neither of them are engaging in approved social practices of reliable information sharing the relevant processes by which we evaluate their resulting beliefs are unreliable in yielding the truth. Consider the following table:

	Process Type	Conditions	Field of Proposition	Reliable?
<b>Compulsively Trusting</b>	Trusting the testimony of someone you are romantically obsessed with	Lacking capacity to discriminate reliable from unreliable testifiers.	Proposition about what one's neighbor saw in the ocean.	No
<b>Gullible in a Room Full of</b>	Trusting a stranger's	Lacking capacity to	Proposition about some	No

<b>Liars</b>	testimony.	discriminate reliable from unreliable testifiers and with unreliable testifiers nearby.	current event. <sup>168</sup>	
--------------	------------	---	-------------------------------	--

Given the relevant process that Sid and Bill use in forming their belief, they could very easily have gone wrong. This makes their behavior unreliable (or epistemically vicious), thereby precluding knowledge.

But let's focus on **Compulsively Trusting** for the moment. Why isn't the relevant process *trusting one's neighbor*? It is not clear to me that *trusting one's neighbor* is not a socially accepted practice of reliable information sharing. Perhaps it is not in today's over-individualistic American culture. Many of us live for years next to someone without even knowing his/her name. But perhaps in the 1940's or 1950's one had closer relationships with one's neighbors. Would we have evaluated Bill differently if we were looking at the case through 1940's (or 1950's) glasses? I don't think so. Suppose we grant that *accepting the testimony of one's neighbors* has never been a socially accepted practice of reliable information sharing. What about *accepting the testimony of a friend*, or *accepting the testimony of a best friend*? These seem as much a socially accepted practice as anything. But, again, if we suppose Jill is Bill's best friend, in whom he has an obsessive romantic interest, it still seems that we should deny that he acquires knowledge on the basis of Jill's testimony. Perhaps caregivers have a special epistemic role that friends don't (assuming our friends are not our caregivers!). It still seems that we can modify the case to put pressure on Greco's account.

---

<sup>168</sup> In the case provided no specific proposition was given. I've here made the innocuous supposition that it was a proposition about some current event.

**Compulsively Trusting Mom:** Bill is a compulsively trusting adult with respect to the testimony of his mother, whom he idolizes. Not only does he always trust his mother when he has very good reason to believe her, but he is typically incapable of distrusting her when he has very good reason to not believe her. His epistemic practices are virtually identical to those he exhibited as a young child. Yesterday, Bill's mother told him that she had seen a whale boating earlier that day. Bill, of course, readily accepted his mother's testimony. It turns out that his mother did see a whale, that she is very reliable with respect to her epistemic practices, both in general and in this particular instance, and that Bill has very little reason to doubt the proffered testimony. Given his compulsively trusting nature with respect to his mother, however, even if he had had evidence available to him indicating, for instance, that his mother did not see an orca whale, that she is an unreliable epistemic agent, that she is an unreliable testifier, etc., Bill would have just as readily accepted his mother's testimony.

If we are inclined to think Bill does not know in the original **Compulsively Trusting** case, I suggest we should be equally inclined to think Bill does not know in

**Compulsively Trusting Mom.** Consider another case involving a pupil-teacher communication:

**Dogmatically Trusting Professor:** Bill has just enrolled into State University and is immediately deeply impressed with his political science professor Dr. Knowsit. Bill's impression of Dr. Knowsit is so strong that he puts all his trust regarding issues of political science in Dr. Knowsit. Not only does Bill always trust his professor when he has very good reason to believe her, but he is typically incapable of distrusting her when he has very good reason to not believe her. His epistemic practices are virtually identical to those he exhibited as a young child toward his pre-school teacher. Yesterday, Dr. Knowsit explains to Bill the major influencing factors leading up to the start of WWII. However, there are other professors in the department who disagree with Dr. Knowsit regarding the causes of WWII. It turns out that Dr. Knowsit is correct in her analysis and is a reliable testifier. However, given his compulsively trusting nature with respect to Dr. Knowsit, even if Bill had had evidence available indicating, for instance, Dr. Knowsit's analysis was incorrect, Bill would have just as readily accepted her testimony.

Again, it seems to me that that in **Dogmatically Trusting Professor** Bill does not know what his professor reliably testifies.

According to Greco's analysis we seem forced to evaluate Bill's beliefs as knowledge. In both cases Bill is engaging in a socially accepted practice of information

sharing. Therefore, the relevant process for evaluating Bill's belief is *trusting the testimony of one's mom* and *trusting the testimony of one's professor*. And since his mom and professor are reliable testifiers, Bill would not very easily go wrong by using such a processes in the conditions he finds himself.

The problem, as I see it, is that some epistemic practices that are necessary and appropriate for a child are not for a fully developed, cognitively mature, adult. By characterizing the relevant process according the accepted social practice that Bill is engaged in, Greco's account ignores the fact that Bill fails to meet his—more demanding—standard for skillful performance, and should therefore be denied knowledge.

#### **VI. Proposed Solution: Knowledge in Terms of Intellectual Skill**

Both Goldberg and Greco approach the problem of testimonial knowledge in children by showing how their testimonial practices are sufficiently reliable for knowledge. Goldberg argues that a child's social environment (i.e. their caregivers) helps make their testimonial consumption reliable. Greco argues that, by engaging in a socially accepted practice of information sharing, the relevant belief forming process by which the child's belief is evaluated, is reliable. However, I have argued that Goldberg's account fails by being overly restrictive whereas Greco's account fails by being overly permissive. First, it seems that children can still acquire knowledge by basing their beliefs on the testimony of their caregivers even if their caregivers are unable to effectively monitor their testimonial uptake. Second, it seems that an individual—particularly, cognitively mature adults—can engage in a socially accepted, and reliable, practice of information sharing without acquiring knowledge.

Both have the background assumption that the standards for evaluation for children should be identical to their adult counterparts. This assumption is seen in the way the problem is generated. We look at cases in which adults fail to have knowledge and wonder why children are evaluated as knowing in seemingly epistemically identical cases. I suggest that the real problem is not as simple as Greco's initial formulation.

1. Young children can learn from the testimony of their caregivers; i.e. that they can come to know through such testimony.
2. Testimonial knowledge requires that cognitively developed adults are reliable consumers of testimony; i.e. that they can reliably discriminate between reliable and unreliable sources of testimony.
3. Young children are not reliable consumers of testimony.
5. The standards for testimonial knowledge should be uniform for adults and children.

But is this extra assumption, 5, warranted? I want to propose that an adequate explanation for how children acquire testimonial knowledge will reject this, heretofore, unchallenged assumption.

Over the last four chapters I have defended the view that knowledge is best understood in terms of intellectually skillful performance. In order to behave intellectually skillful one's behavior must be responsible and reliable. And I have argued that what makes us responsible for our beliefs is not whether we have willed to believe what our evidence supports but whether we have approached our belief forming processes or faculties appropriately, according to what we have good reason to believe will lead to the truth. In order to believe responsibly we must be sufficiently attentive, careful, and thorough in evaluating the evidence possessed. Moreover, we must display similar traits in gathering evidence and in performing our intellectual inquiries. Only

when our intellectual behavior is reliable *and* responsible will we be performing intellectually skillfully.<sup>169</sup>

But must the standard for skillful performance be rigidly set across all people and all times? Consider the following quote from James Wallace:

A skill is the capacity ... to perform a certain sort of action well in the sense of performing the action proficiently. Proficiency in the action is a matter of successfully overcoming the technical difficulty inherent in the action, thereby being able with some regularity to achieve the desired result efficiently and economically. A particular action will be characteristic of a certain skill if the action satisfies the standards of proficiency for actions of that kind. The relevant standards of proficiency, of course, will vary with the kind of action. Also, there are degrees of skill, and that standard by which a particular performance properly is judged is often relative to the degree of experience of the agent. There are often different standards for a beginner and for an expert, and what is a good performance for a beginner is apt to be a poor performance for someone more advanced.<sup>170</sup>

Wallace's observation about how standards for skillful performances shift depending on whether the agent is a beginner or expert seems accurate when we look at our evaluations of more ordinary performances as skillful. We credit the little leaguer for a skillful performance even though her performance would not be impressive when compared to how we evaluate a major leaguer. As she develops as a player the standards for skillful performance become more demanding. She is expected to approach hitting differently than when she was less mature—similarly for fielding and base running. If she keeps her old habits and maintains the same approach to the game it is likely that even when successful she will not be performing skillfully, and therefore will not deserve credit for her success.

---

<sup>169</sup> The proposed solution to this problem can be adopted by virtue reliabilist and virtue responsibilists alike. However, it seems to me important that framework of intellectual skills is adopted since it provides the motivation for having different standards for a credited performance.

<sup>170</sup> Wallace (1978): p. 48.

Although the expert's approach to a situation is different from the beginner's it would be a mistake to view it as completely different. This is important for distinguishing the beginner from the individual who lacks skill altogether. The beginner, but not the unskilled, receives credit for a successful performance because her approach significantly resembles the approach of the expert. Julia Annas notes that many of the ancient moral theories took moral virtues as a kind of skill, and recognized the similarities between the behavior of the learner of skill and the behavior of the expert.

In ancient theories ... [we find] an attempt to show how ordinary reasoning can first, in the person who aims to be virtuous, start getting it right, and then develop to the point of full virtue. *The fully virtuous person may reason differently from the beginner; but the beginner's reasoning is transformed, rather than abandoned or downgraded.* Again, skills give us a useful general picture: the expert musician will approach the playing of a piece of music differently from the learner; but his is because she no longer needs to go through the thoughts that the beginner has, not because the beginner is having the wrong kind of thoughts.<sup>171</sup>

The Stoic moral beginner learns to be aware of the kinds of reason she has for acting.... Learning to be good starts with learning right behaviour. *The person who has become virtuous, then, will still do the same kind of action as he did before, but will stand in a different relation to it.* The beginner did the due action because there was good reason for him to do it, and he did it because of that reason. But this was compatible with ... failing to realize that there were other moral requirements in that situation; failing to grasp similar moral requirements in other situations. *The virtuous person will do the same action, but ... will discern and balance all the relevant values in that situation; will reliably discern similar values in all relevantly similar situations. This is just the distinction, familiar by now, between the beginner and the fully virtuous person with phronesis, developed practical intelligence.*<sup>172</sup>

There are two things to note about how the ancients saw the relationship between the beginner and the expert. First, the beginner's *action* and *reason for acting* may be identical to that of the expert. The expert musician's performance may still be distinguished from the beginner's despite both playing the same note for the same reason.

---

<sup>171</sup> Annas (1993): p. 85 (emphasis added).

<sup>172</sup> Annas (1993): p. 97 (emphasis added).

Second, what does distinguish their performances is the *way* that each is related to their respective actions. The expert has a kind of *understanding* or *grasp* that the beginner lacks. This sort of understanding need not be something that the expert can articulate to others or even to oneself, nor need it be a precise algorithm that easily calculates the appropriate behavior in every situation. This deeper, or fuller, understanding/grasp is manifested in the expert's ability to see all of the relevant values in a given situation, and in her ability to discern similar values in different, but relevantly similar, situations. The expert and the beginner may do the same action for the same reason, but the expert will be better able to recognize the counter-reasons for acting, even if on the whole they are outweighed. The expert, therefore, has a sort of sensitivity absent in the beginner, and consequently she is not as limited in her performance. Had the balance of reasons changed slightly, the expert, but not the beginner, would be able to adapt and still perform appropriately.

In contrast to the expert, the beginner performs the right action when she employs the right rules in a given situation, but she does not have internalized the rules into more general principles for acting. She is reliant on rules in the way the expert is not and therefore is less flexible when it comes to performing the right action in similar but significantly different situations. So how does the beginner become an expert? Through experience. Aristotle viewed virtues as analogous to skills in that we acquire virtues

just as we acquire crafts, by having first activated them. For we learn to craft by producing the same product that we must produce when we have learned it; we become builders, for instance by building, and we become harpists by playing the harp. Similarly, then, we become just by doing just actions, temperate by doing temperate actions, brave by doing brave actions.<sup>173</sup>

---

<sup>173</sup> Aristotle (1999): *NE* 1103a32–1103b3.

As the beginner confronts new situations she will hopefully recognize the reasons and purpose for the rules she employs. She will see how and when it is appropriate to adjust the rules given the demands of the situation.

It is important not to over-intellectualize the notion of understanding employed above. Some ancient schools claimed that an expert must be able to explicitly grasp the unifying principles underlying the relevant skill and be able to articulate and teach this understanding to others. This seems to be the view of Socrates and Plato.<sup>174</sup> However, not all schools were in agreement. Some took a more ‘empiricist’ approach to skillful performances in which the expert, through repeated experience, acquires an ability to ‘see’ what the right thing to do is in a given situation.<sup>175</sup> I don’t have the space to adequately engage this very interesting debate, but I do think that there are intuitive and empirical reasons for adopting the less intellectualist approach to skills, and it is the view I adopt here.

These general considerations regarding the nature of skillful performances, I believe, give us the resources to better understand the epistemic problem of explaining how children acquire testimonial knowledge. Knowledge is a non-accidentally true belief; that is, a true belief for which the believer deserves credit. Now in order to receive credit for reaching the truth one’s intellectual performance must be sufficiently skillful, and one’s success must be because of the skillful performance. But, as Wallace observes, the standards for beginners are typically lower than for experts or those more advanced. It is a mistake to hold children to the standards of adults. We don’t do this when assigning

---

<sup>174</sup> C.f. Annas (1995).

<sup>175</sup> According to Stichter (2007) this is Aristotle’s view. Hutchinson (1988) argues that the Isocrates and other rhetoricians were opposed to Socrates and Plato in holding an empiricist approach to skills.

credit for a successful performance in other, more ordinary domains, and we shouldn't do it when assigning credit to our intellectual performances either.

Of course, this does not mean that no standards apply to the child. If we look back to **Babe with Mother** and **Compulsively trusting Mom** we see that both Babe and Bill engaging in *trusting mom's testimony* behavior. Given that both Bill's and Babe's mom are reliable testifiers this sort of behavior is not, by itself, objectionable. What is potentially objectionable is their inability to be sensitive to signs of unreliable or false testimony. I say "potentially" because this inability seems to compromise only Bill's credit for his true belief. Why? Because Bill is judged according to standards appropriate for cognitively mature individuals. This requires that he have a better grasp of his behavior and how it is appropriate in his circumstance. This is not required of Babe, on the other hand. Because Babe is still cognitively immature and still developing, he has not had the relevant experiences needed to develop the sensitivities characteristic of those more mature. As Annas points out

The learner, paradigmatically the young learner, begins by picking up what to do in particular cases; he copies his elders or follows their advice. But if he is intelligent he does not remain stuck at the stage of depending on models for each new case or memorizing a list of cases and dealing with each new one by comparing it with the past ones. Rather, he develops a sense of the point of doing these specific things, and when he grasps this he has a sense of the basis of these previous judgments, which will enable him to go on to fresh cases without mechanically referring back. Like the person who has acquired an expertise in a skill, the learner has acquired understanding of what he is doing, an understanding that can be represented as a unified grasp of the principles that underlie his actions and decisions.<sup>176</sup>

Babe, like most children, simply follows the directive of his caregivers to believe p but not q, r but not s. And he is right to do so. This sort of rule or model following is necessary for the proper development of intellectual skill. As he gains experience he will

---

<sup>176</sup> Annas (1993): p. 67 (emphasis added).

hopefully see the point of trusting his caregivers, and learn that this rule of thumb is defeasible in certain environments. As a result Babe will hopefully learn to monitor his environment for such situations so that in the event that they are present he can adjust his intellectual practices. Bill's intellectual behavior though seemingly identical to Babe's is not sufficiently responsible and thereby, not sufficiently skillful given his stage of cognitive development, and as Babe matures more will also be required of him. Eventually the intellectual behavior that Babe exhibits in **Babe with Mother** will be insufficient to credit him for a true belief.

*VI.i. The Generality Problem and Characterizing Babe's Behavior*

My approach to explaining testimonial knowledge in children will be adequate only if the intellectual behavior that children engage in is properly identified as reliable. Recall that I described Babe's relevant behavior as *trusting mom's testimony* and concluded that, since his mom is a reliable testifier, this behavior is reliable in reaching the truth. However, some may think that I have simply helped myself to this without giving any reasons for choosing this type of behavior for evaluating Babe's belief, rather than some other type of behavior that also characterizes Babe's intellectual performance. Greco was able to provide a response to this generality worry by claiming that we should understand our concepts by looking at what purpose they serve and how we use the concepts. He goes on to argue, following Edward Craig, that part of the function of our concept 'knowledge' is to identify reliable sources of information. Consequently, when one engages in a socially accepted practice for reliable information sharing, then the relevant process type by which one's belief is evaluated is one that identifies the reliable source. But since I have

rejected Greco's account one may think I am without resources to make a principled identification in the way Greco does.

I follow Zagzebski's approach to the generality problem by looking at the habits that one's behavior exemplifies or promotes.<sup>177</sup> But care is needed in determining which habit is being exemplified. Suppose I have the habit of believing whatever I read in print—whether it is a news article from the New York Times or an article from the National Inquirer. However, suppose that I also place special weight on the reports from a reliable newspaper; the New York Times, let's say. In the face of conflicting reports between the Times and other publications I believe the Times. Moreover, if given the option I would prefer to look at the Times for my information, and in many cases I actively seek out to find a copy of the Times rather than other publications. Finally, suppose that on Tuesday I go to the newsstand, consciously buy the Times and come to believe that the president is in Florida. Question: What habit does my belief forming behavior exemplify? Although I do have the tendency to believe whatever is in print, and though my belief is based on what I read from the newspaper, it seems to me that my behavior is best characterized, not by *believing what I see in print*, but by *believing what the Times reports*. My preferential standing with respect to what the Times reports suggests that the relevant habit for describing my behavior is more narrow than it might first appear.

So when Babe believes his mom's (or his teacher's) testimony is his intellectual behavior representative of the habit of believing what his mother tells him, or the habit of believing what anyone tells him? Providing an answer to this question will depend in part on how indiscriminate children are in their testimony consumption. Thus far we have

---

<sup>177</sup> Zagzebski (1996): p. 310-11.

assumed that children are generally unreliable in monitoring good from bad testimony. If Babe's mother were an unreliable testifier, or if someone else (e.g. a classmate or a deceiving relative) were to provide unreliable testimony, Babe would be unable to detect it, and would readily believe the testified proposition. However, this does not mean that Babe is indiscriminate in all respects. In fact, there are some recent empirical studies that suggest that even at a young age children do show preferences in whom they trust.

Kathleen Corriveau and Paul Harris have provided a series of studies suggesting that children are more likely to seek and endorse information from their mother over that of a stranger, and from a familiar rather than an unfamiliar teacher. One study (Corriveau, Harris, et al., forthcoming), conducted on children of 50 to 61 months of age, showed that when presented with novel objects or with hybrid creatures falling equally into two different categories children preferred (i) to ask their mother, rather than a stranger, for the name or function of the object, and (ii) to endorse the names and functions provided by their mother over those provided by the stranger. In another study (Corriveau and Harris, 2008), conducted on children of 3, 4 and 5 years of age, showed that when presented with novel objects children preferred to ask for, and endorse, the name and function provided by a familiar teacher rather than an unfamiliar teacher. Moreover, the younger children (3 and 4 year olds) were more resistant in modifying this behavior in the face of strong evidence that the familiar teacher was significantly less reliable in accurately labeling objects than the unfamiliar teacher.

What these studies suggest for our purposes is that when children base their beliefs on the their caregivers (e.g. their parents and teachers) testimony, they are typically not manifesting a habit of believing what anyone says but rather of believing

what their caregivers testify. This seems especially true of younger children (younger than 4 years of age), who seem less able to attenuate this preferential trust on the basis of strong evidence suggesting that their caregivers are unreliable testifiers. Of course this is not shown simply by the fact that preferential trust is given to one's parents and teachers. Babe in the used car lot may show a preferential trust to the grandson, but we still don't think this is enough to categorize his belief as *trusting the grandson's testimony*. The reason is that he does not know that the testifier is the grandson and therefore does not approach the testifier *as* the grandson of the car. I suspect that things are different when children acquire testimonial knowledge from their caregivers. Not only do they give special weight on the testimony of their parents and teachers, they approach them as such. Indeed it is because they recognize them as parents and teachers that they exhibit preferential trust. The reason why the studies by Corriveau and Harris are relevant is that systematic preferential trust that children exhibit toward their parents and teachers suggests that they are approaching them as their parents and teachers. And it is this that is relevant to their habit formation and for categorizing their belief forming behavior.<sup>178</sup>

Once we have identified the relevant habit, and thereby the relevant type of behavior exhibited by the child, we can evaluate whether it is a reliable guide to truth. And provided that the caregiver is a reliable testifier, then the child's reliance on the caregiver for her testimony-based belief is also a reliable means for truth.

We now can also explain why Babe acquires knowledge in **Babe in Preschool** but not in **Babe in Used Car Lot**. The sort of behavior that Babe exhibits in coming to believe that the previous owner of the car was a grandmother does not exemplify the habit of believing whatever her grandson says, but rather the habit of believing whatever

---

<sup>178</sup> Thanks to Ernie Sosa for pressing me to clarify this point.

anyone says, or, perhaps, whatever a stranger says. Babe is unaware that the individual he overhears is the grandson of the previous owner, and as a result would not have any preference for this person's report over that of another stranger. By contrast, when Babe is in the classroom and comes to believe what his teacher tells him he exemplifies the habit of believing whatever his teacher tells him. This can be seen in the sort of preferential trusting—described in Corriveau's and Harris's studies—that children exemplify toward their teachers.

## 6. Conclusion

Empirical studies provide strong evidence that children are unable to reliably detect reliable from unreliable testifiers. Young children are generally poor at attenuating their trust on the basis of bad track records and are typically insensitive to reasons suggesting that a testifier or a piece of testimony is unreliable. This provides a challenge for providing a general account of testimonial knowledge since when looking at cases of involving adult testimonial consumption we require that the hearer be able to reliably discriminate reliable from unreliable testimony. But we still want to maintain that children acquire testimonial knowledge, especially when they form beliefs based on the reports of their caregivers.

Goldberg and Greco attempt to resolve this tension by trying to explain how children are more reliable than it might first appear. They argue that their testimony-consumption practices are sufficiently like their adult counterparts to acquire knowledge. However, I have argued that such attempts go wrong for different reasons. Goldberg argues that children are as discriminate as adults since their caregivers closely monitor their testimonial consumption. However, his explanation fails to account for cases where

children acquire testimonial knowledge despite not being monitored. Greco opts for a more lenient account of knowledge. Knowledge, Greco argues, requires a reliable process of belief formation, and though children are generally indiscriminate consumers, when they believe the reports of their parents and caregivers they are engaging in a socially accepted practice of reliable information sharing. After taking into account the reliability of the testifier, the conditions at the time of consumption, and the type of proposition testified, if the process is reliable at reaching the truth, then the child acquires knowledge. Although this explanation seems accommodate the cases in which, intuitively, children acquire testimonial knowledge, it makes knowledge too easy for the more cognitively mature. In many cases, if adults exhibit the sort of indiscriminate behavior that characterizes children, then—even when they are engaged in a socially approved practice of reliable information sharing—they fail to acquire testimonial knowledge.

In summary, while Goldberg raises the standards for testimonial knowledge for children to the level of their adult counterparts, Greco lowers the standards for adults to the level of children. Both moves have debilitating consequences. I have argued that instead of trying to find a standard that both classes of individuals satisfy, we should recognize that there are two different standards for testimonial knowledge one for children and one for adults. I attempted to make this move plausible by understanding knowledge in terms of intellectual skills, such knowledge is a true belief for which one deserves credit, and whether an individual deserves credit for a true belief requires that the individual's doxastic performance be sufficiently skillful. But like other skills, the standards for a skillful performance will depend if one is a beginner or an expert. An intellectually skillful performance for a child may be intellectually irresponsible for an

adult, and the sort of trust that I displayed as a child may now be inadequate to acquire knowledge.

## Curriculum Vita

Joshue Orozco

## Education

- 2004-2009 Rutgers University, New Brunswick – Philosophy, Ph.D.
- 2001-2003 Biola University – Philosophy of Religion and Ethics, M.A
- 1997-2001 University of California, Los Angeles – Philosophy (Double Major in Biology), B.A

## Occupation

- 2001-2003 Teaching Assistant, Biola University
- 2004-2005 Graduate Fellow, Rutgers University
- 2006-2007 Teaching Assistant, Rutgers University
- 2007-2009 Graduate Assistant, Rutgers University