

**University of Alberta**

**When We Should Know Better: Investigative Virtues and Moral  
Blameworthiness for Ignorance**

by

Stephanie Joanne Buri Choy



A thesis submitted to the Faculty of Graduate Studies and Research  
in partial fulfillment of the requirements for the degree of

**Master of Arts**

Department of Philosophy

Edmonton, Alberta  
Fall 2008



Library and  
Archives Canada

Bibliothèque et  
Archives Canada

Published Heritage  
Branch

Direction du  
Patrimoine de l'édition

395 Wellington Street  
Ottawa ON K1A 0N4  
Canada

395, rue Wellington  
Ottawa ON K1A 0N4  
Canada

*Your file* *Votre référence*

*ISBN: 978-0-494-46953-8*

*Our file* *Notre référence*

*ISBN: 978-0-494-46953-8*

**NOTICE:**

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

**AVIS:**

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

---

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.



**Canada**

## Abstract

This thesis explores the issue, important for the fields of ethics and epistemology, of when persons are morally blameworthy for the results of their failure to know certain propositions. The issue is examined from the perspective of virtue ethics and follows Aristotle in the basic conception of “virtue” used. Two kinds of virtues are discussed: Aristotelian moral virtues, and a cluster of pragmatic virtues here termed investigative virtues. The argument is made that persons are blameworthy for the morally negative results of their failure to know if and only if they have failed to undergo their epistemic investigations in accordance with the moral and investigative virtues. The concepts of relevant alternatives, right action and virtue education are utilized in making this argument.

## Acknowledgment

I would like to thank my examining committee, Adam Morton, Jennifer Welchman, Bruce Hunter and Stephen Norris for your input and help in finishing this thesis. I would especially like to thank Adam Morton and Jennifer Welchman for your fantastic input throughout the writing process. Your ideas and comments helped me greatly along the way.

I would also like to thank everyone who has supported me throughout the past years: Mom and Dad, for instilling in me a love of learning and for always encouraging me to do what I love; Katherine, George and Maryann, my siblings, for their constant support; my good friends Carlie, Katie, Kris, Amanda, Jen and Gerald; and Kai, my husband, who is always there for me, and whose many thoughts and discussions are essential to this work. The impact of you all in my life has informed every word I write, and helped me to make this thesis what it is.

**For Mom and Dad, and for Kai**

## Table of Contents

Introduction	1
Chapter One: Relevant Alternatives	6
Chapter Two: Virtue Ethics	24
Chapter Three: Virtue Ethics and Right Action	31
Chapter Four: Virtue and Relevant Alternatives	52
Chapter Five: Moral and Investigative Virtues and Relevant Alternatives	71
Conclusion	89
Bibliography	92

## **Introduction**

Despite the arguments of the skeptic, most people are quite certain that we know a great number of things. We believe that we know facts about our current personal experiences, such as that we see, for example, the page before us. We also believe that we know more speculative claims, such as that there are people in New York City at this moment, or that the Sun is approximately 8 light-minutes from Earth. It is also true (though it may seem silly or even tautologous to state it) that, while there are many things that we are certain we know, there are also many things that we are certain we don't know and many things that we aren't certain whether we know or not. For example, I am certain that I don't know whether or not there is life on other planets, and I am not certain whether or not I know that some complex laws of science of which I have some understanding, such as Heisenberg's Uncertainty Principle, are true. One of the main tasks of epistemology has been to determine what sorts of things we really can and do know. There are also ethical questions related to epistemology, that explore the moral obligations that persons have with respect to formulating and expressing what we believe to be knowledge. It is one such ethical issue that I will explore in this thesis.

Consider the following situation faced by an office worker, Alex: his co-worker, Jane, has been lethargic at work, seems not to care about her work, has been getting reports in late, and so on. Alex also notices that Jane is on the phone a lot, and has been coming in to work with dark circles under her eyes as though she hasn't been sleeping. From this evidence, Alex concludes that Jane has been staying

out late. He further believes that he knows that she is a lazy party animal who is hurting the whole office. After three straight weeks of Jane behaving in this manner and the fifth late report, Alex confronts Jane about her behavior. He tries not to be unkind, but makes it clear that he is aware of her behavior, that she is hurting the office, and that if she doesn't start doing her job he will notify a higher authority in the company. Alex's rationale for this confrontation is that he is helping to fix a bad situation, and Jane will have a chance to change her behavior before losing her job. Alex is surprised when Jane gets visibly upset and runs off after he confronts her. Afterwards, he is told by an angry officemate that Jane's husband has been in an accident and has been in a hospital over the last three weeks and that Jane has been spending most of her time off work with him. The officemate tells Alex that he "should have known better" before confronting Jane, since everyone else in the office seems to have known and he should have discussed the situation with someone else before simply giving Jane a piece of his mind.

What has happened here? First of all, Alex made a knowledge claim to himself, that he knew that Jane was a lazy party animal.<sup>1</sup> This knowledge claim directly led to an action, that of confronting Jane. However, this action had an unforeseen negative consequence: it made Jane, who was already in a difficult and painful situation, feel even worse. So Alex's erroneous conclusion led him to perform a morally wrong action. Alex's officemate is holding him morally blameworthy for the reason that, even though he did not expect the morally wrong

---

<sup>1</sup> Alex may not consciously think "I *know* that Jane is a lazy party animal." However, if asked whether or not he *knows* his conclusion to be true, he would likely say yes.



result, he *should have known* that his actions would have such an effect on Jane.

This example will serve as the template for my discussion. To put it in general terms, the sort of case I will consider goes as follows: a person, P, believes that she knows that some claim, Q, is true. Believing that Q leads P to take some action, T, that she would not have taken if she did not believe that Q. However, Q is not, in fact, true. Doing action T leads to some morally bad consequence, C, that was unforeseen by P. Furthermore, there is some fact or information, F, such that if P knew that F, then P would no longer believe that she knew that Q. Instead, she would believe some other claim, A<sub>1</sub>-A<sub>n</sub>, to be true, which would not lead P to take action T. Note that A<sub>1</sub>-A<sub>n</sub> may only be the claim that Q is false, or it might be some alternate claim. In the case above, the alternative A would be that Jane is doing the best she can despite horrible circumstances and needs to be cut some slack.<sup>2</sup>

This fairly complex template boils down to one concept: “should have known better” (SKB). The way I am using SKB has a normative dimension; to say that someone SKB is to imply that she is morally responsible for the consequences of her ignorance. There are other ways in which the expressions “should have

---

<sup>2</sup> There will be slight variations on this template that are equally relevant to my discussion. It may be the case that, rather than claiming to know a false proposition, the agent, P, may not make any knowledge claim at all, for the reason that she has failed to put an adequate amount of thought into the situation. This lack of knowledge may likewise lead her to take action T, once again causing the bad consequence, C. In this case, we could also say that P *should have known* some proposition A<sub>n</sub>, such as that Jane's husband is in the hospital.

known” or “should have known better” are used that have no moral dimension. For example, I might decide to take a certain bus route at 5:00 that is slowed by traffic. While on the route, I might say to myself in exasperation, “I should have known better than to take this route during rush hour.” Other ways in which the expression may be used are when a person who spent an entire afternoon looking for her car keys might say “I should have known to look in the drawer first,” or when a person is offered comfort from a friend and says “I should have known that you would know exactly what to say to make me feel better.” My discussion does not apply to these other uses of “should have known better,” but only to those uses in which moral responsibility is implied.

However, it is possible for someone to hold a false belief under circumstances where it would be false to say that they SKB and where we cannot hold them morally responsible for the consequences of that ignorance. For example, consider a variation of the Jane case above in which Jane has told no one about her situation and in which Alex has discussed Jane with other office mates and received their support before confronting her. The consequences of making Jane feel worse and embarrassing her will still follow from Alex's action, but in this case it is not clear that he SKB and is morally blameworthy.

What I want to explore is how we may be able to determine in which sorts of cases an agent SKB in the moral sense. Fitting this question into my template gives us the following question: in what sorts of cases are we morally blameworthy for concluding that we know  $Q$  as opposed to  $A_1-A_n$  (or for failing to reach a conclusion

at all)? These alternatives  $A_1$ - $A_n$  represent any situation in which the alternative is true and the claim,  $Q$ , is false, including the alternative “ $Q$  is false.” I will explore this question by utilizing the concepts of relevant alternatives and virtue theory. In order to be able to confidently make a knowledge claim and avoid the possibility that she SKB, an agent must exercise the necessary virtues when making this knowledge claim. The necessary virtues will include both Aristotelian moral virtues, as well as other virtues that are similar in structure, but are not Aristotelian moral virtues. I will call these the “investigative virtues.” I will argue that a person is morally responsible for her ignorance if she fails to act in accordance with moral and investigative virtues.

## Chapter 1: Relevant Alternatives

The question of knowing that some proposition is true instead of an alternative proposition naturally leads us to relevant alternatives literature. Discussion of relevant alternatives is typically in the context of disproving the claims of the epistemic skeptic. Hence, I will first briefly summarize the skeptic's position. The skeptic claims that we do not, in fact, *know* many things that we think we do. The reason for our lack of knowledge is that we cannot prove wrong certain propositions that undermine our knowledge claims. A common example employed by the skeptic is as follows: suppose you are at a zoo and are viewing an animal that appears to be, and is labeled as, a zebra. You make the knowledge claim "That is a zebra," meaning, "I *know* that that is a zebra." However, asks the skeptic, can you be certain that what you see is actually a zebra and not really just a cleverly painted mule? Most people (with the exception, presumably, of a zoologist who is an expert in mules or zebras) cannot be certain that they are not looking at a painted mule rather than a zebra. This demonstrates, according to the skeptic, that you *do not know* that the animal you see is a zebra.

In fact, many alternatives, such as that what you are looking at is a painted mule, can be found for almost all knowledge claims. Even mundane claims, such as that I know that there is a computer screen in front of me, are subject to the skeptical alternative that I am actually being deceived by an evil genius into believing that I see a computer screen, when in reality there is no screen there at all. Many

philosophers have rejected the skeptics' arguments on a number of different grounds. One approach is to answer the skeptic using relevant alternatives theory. I will briefly discuss how proponents of relevant alternatives theory seek to answer the skeptic, but, for the purposes of this thesis, I will be focusing on accounts of what a "relevant alternative" is and how relevant alternatives relate to the question of when an agent "should have known better."

### *Dretske's Formulation*

One of the earliest proponents of relevant alternatives (RA) theory was Fred Dretske, who set out his theories in his article "Epistemic Operators" (1970). Here Dretske argues that, though there are a large number of possible alternatives to most knowledge claims, not all of these alternatives are relevant to our knowledge claims. Dretske begins with a discussion of language operators, arguing that not all such operators are "penetrating". An operator is fully penetrating if, for every  $q$  such that  $p$  entails  $q$ , the statement we get by applying the operator to  $p$  entails the statement we get by applying the operator to  $q$  (1007). A statement is *not* fully penetrating if this is *not* so. Operators such as "it is strange that" are not fully penetrating, or are "closed under entailment": the fact that Frank holds winning tickets in all the races entails that Frank holds a winning ticket in one specific race, say the third race. However, although it is true that it is strange that Frank holds a winning ticket in all the races, it is not true that it is strange that Frank holds a winning ticket in the third race (1008-9).

Dretske goes on to apply his arguments to *knowledge*, arguing that epistemic operators are not fully penetrating. His claim, fully spelled out, is as follows: it is *not* true that, for any and all  $q$ , if an agent  $A$  *knows* that  $p$  entails  $q$ , then if  $A$  knows that  $p$  then  $A$  knows that  $q$  (1010). So,  $A$  may know that  $p$  entails  $q$  and may know that  $p$ , but may nonetheless not have knowledge that  $q$ . This view goes against what Dretske identifies as the common practice in philosophy, which is to treat epistemic operators as penetrating (1011). According to Dretske, the reason why  $A$  may not have knowledge that  $q$  in cases such as this is that  $q$  is not one of the consequences of  $p$  to which the operator “knows” penetrates;  $q$  is not relevant to  $A$ ’s justification for believing that  $p$  (1012). So, not all the entailments of a statement,  $p$ , are operated on by the epistemic operator that operates on  $p$  (1014).

Dretske goes on to apply these arguments to skepticism. The skeptic argues that any proposition,  $p$ , entails that a number of alternatives,  $q_1$ - $q_n$ , are false. For example, the proposition “this wall is red” entails a number of other alternate propositions, including the proposition “the wall is not actually white illuminated to look red.” According to the skeptic, since we do not normally know that the wall is not actually white illuminated to look red, we do not truly know that the wall is red. Dretske disagrees with this conclusion, arguing that the proposition that the wall is not really white illuminated to look red is *not* one of the entailments of “the wall is red” to which the epistemic operators penetrate (1015). Hence, the wall’s being white illuminated to look red is not a relevant alternative to the proposition that the wall is red. To look at it another way, since the wall actually being white

illuminated to look red is not a relevant alternative to the wall being red, we do not need to know that the wall is not white illuminated to look red in order to know that the wall is red.

This is a hasty look at Dretske's arguments, but it serves to demonstrate the importance of having some way of determining *which* alternatives are relevant in a given situation. Different authors have attempted to give an account of what it means to say that some alternative is relevant. I want to go over some of these views and see whether or not they are satisfactory. In "Epistemic Operators," Dretske is mostly concerned with setting out his reply to skepticism, but he does provide a brief account of what it means to say that an alternative is relevant. Dretske states that a relevant alternative is "an alternative that *might have been realized* in the existing circumstances if the actual state of affairs had not materialized" (1021, my italics).

### *Stine's Criticism*

Gail Stine points out that this is a fairly ambiguous definition, as it is unclear what it means to say that the alternative "might have been realized." She argues that there are two possible interpretations: i) an alternative is relevant if there is some reason to think that it *could be* true and ii) an alternative is relevant if there is some reason to think that it *is* true ("Skepticism..." 252). Stine contends that Dretske is properly interpreted as claiming the first option to be correct, but that this is not satisfactory, as it would allow too much to be considered as a relevant alternative

(253). “Could,” says Stine, can be interpreted to indicate either logical or physical possibility. Clearly, if it were interpreted to mean logical possibility then all the skeptical alternatives would be relevant, which is not what Dretske wants. However, the situation would not be much better if “could” were interpreted to indicate physical possibility. In this case, skeptical alternatives such as a zebra actually being a cleverly painted mule, which are physically possible, would also be relevant (253).

Stine argues that, while it is admittedly unclear what it means to say that there is reason to believe that an alternative *is* true, this second interpretation captures the purpose of RA theory much better, and rules out many skeptical alternatives (253). In the painted mule case, we would be allowed to rule out the possibility that the zebra we think we see is actually a cleverly painted mule unless we had a particular reason to think that we are looking at a cleverly painted mule (for example, we recently read a newspaper article about some zoos painting mules to look like zebras). While I agree with Stine that interpreting a “relevant alternative” as an alternative that *could* be true would allow too much to be considered relevant, I believe that her interpretation would allow too *little* to be considered relevant. We might ask what kinds of reasons Stine would accept or how good a reason for believing that an alternative is true would have to be for that alternative to count as relevant. Would the agent need direct evidence, or would a suspicion based on intuition or a complicated inference be enough?

Even if we interpret “reason” in the loosest way possible to mean any reason



at all, Stine's view still rules out genuine relevant alternatives. It is possible that a person may face a case in which some alternative is clearly relevant, even though she has no particular reason to believe that the alternative *is* true (although she will presumably have a reason to think that it *could be*). Furthermore, if we interpret "reason" too loosely to mean any kind of reason, we risk being forced to accept alternatives as relevant that we do not wish to view as relevant. Suppose a superstitious doctor is attempting to diagnose a patient. All the symptoms point to condition A, a common disease that can easily be treated if addressed promptly. However, suppose that the doctor is a minor astrologer whose study of the stars has led her to conclude that her patient has condition B, a disease which doesn't quite fit the symptoms and that has not been diagnosed for decades. Nonetheless, the doctor is convinced that there *is* a genuine possibility that the patient has condition B and so she begins a lengthy series of tests to make sure. In other words, she considers herself to have a (non-conclusive) reason to believe that it *is* true that her patient has condition B. I find it very unlikely that Stine would want to include such reasons in stating that a relevant alternative is an alternative that the agent has reason to believe *is* true. This leaves her in need of further criteria, however, for determining what sorts of reasons do, in fact, count as good enough reasons.

#### *Dretske's Five Criteria*

In a later article, "The Pragmatic Dimension of Knowledge," Dretske gives a more detailed account of what it means for an alternative to be relevant. In this

article, he introduces the concepts of a “contrasting set” and a “relevancy set.” The contrasting set is “the class of situations that are necessarily eliminated by what is known to be the case,” while the relevancy set is the “set of possible alternatives that a person must be in an evidential position to exclude (when he knows that P)” (371). In other words, the contrasting set is the set of all alternatives, while the relevancy set is the set of relevant alternatives. A skeptic will argue that the contrasting set and relevancy set are equivalent, a position that Dretske obviously looks to reject (372). Dretske then gives a list of the five considerations that he believes determine the makeup of the relevancy set. I will set out these conditions as best I can, and then consider criticism of them afterwards.

The first consideration has to do with the context of knowledge. It is what Dretske calls “contrastive focusing”: the placement of the emphasis in a sentence determines what claim, exactly, is being made (373). For example, I could say that I know that *your friend* is tall or that I know that your friend is *tall*. In the first case, I am making the claim that I know that the tall person I see standing there is your friend. In the second case, I am making the claim that the person standing there who is your friend is tall. In the first case, the relevancy set will consist of alternatives to the subject being your friend; for example, she might be your sister, your co-worker, or even a complete stranger. In the second case, the relevancy set will consist of alternatives to her being tall; for example, she might be standing on a box, she might look tall because the person beside her is extremely short, or she might be wearing very high heels. In the first case, in which I am claiming that I know that *this* tall

person is your friend, the alternative that she is standing on a box is not relevant, as she could stand on a box whether she is your friend or not.

Dretske's second and third considerations have to do with what we can legitimately assume when making knowledge claims. The second thing to consider, then, is the subject term that we choose to utilize (373). For example, if I claim that your sister (or, to avoid borderline cases of "sister", female sibling) was amused by the story I told, I am making certain assumptions about the subject, your female sibling: namely, that she is a human being (not a robot or any other sort of creature) and hence capable of feeling the amusement that I have ascribed to her. I am not claiming to *know* that your sibling is a human being; I am simply presupposing that she is. Hence, it is not a relevant alternative that she is a robot and as such incapable of actually being amused. (Although it might be a relevant alternative that she is angry and pretending to be amused.)

The third, related, thing to consider is our expression of *how* we know. For example, I might claim that I know that your tire is flat by saying "I *see* that your tire is flat" (374). When we make such claims, we are assuming that the way in which we know is generally reliable. Likewise I might also say that I *heard* two people fighting or *used my calculator* to get an answer. In all of these cases, alternatives that involve my source of information being inaccurate are not relevant. For example, in the case of me seeing that you have a flat tire, it is not a relevant alternative that my eyesight has experienced a strange and sudden form of failure that causes me to see tires as flat.

The fourth consideration that Dretske identifies is far less straightforward.<sup>3</sup>

Dretske argues that the importance of the information to speakers and listeners partially determines which alternatives belong to the relevancy set (375). The greater the importance of the knowledge, the more alternatives will be relevant. This is so because as the importance of the knowledge claim to the speakers and listeners increases, it becomes more crucial to omit any error. For instance, it is fairly trivial for me to claim that I know that summers are hotter than winters. However, it is not at all trivial for me to claim to know that summers are hotter than they used to be when I was a child. This second claim in itself may seem almost as trivial as the first, but the potential implications of the claim (i.e. that we are experiencing climate change) make it more important to myself as the speaker and to my listeners. Many more alternatives, then, will be relevant to the second claim. For example, the alternatives that I am biased or being deliberately misled by others are not relevant to the first claim, but may have to be ruled out with regard to the second claim.

Dretske's fifth consideration is the remoteness of the possibility of a claim (376). An alternative not ruled out by the above considerations might be ruled out because it is just too remote or unlikely. Dretske does not spell out further what it

---

<sup>3</sup> Dretske gives an account of this consideration and then applies it to an example which, in my opinion, does not fit with our intuitions. Cohen gives an account of where Dretske has gone wrong here (Cohen 30-31), but I will try to present what I think is key in Dretske's discussion without dealing with this example.

means for an alternative to be too “remote,” but he seems to want to say simply that an alternative which has a very low probability is not a relevant alternative. After setting out these five considerations, Dretske contends that, when considering which alternatives are relevant, we must consider what really *is* relevant in the actual situation, and not what we *regard* as relevant (377). This implies that we may *believe* that we have ruled out all RA's in cases where we do not actually have knowledge (378). This point is very important to the question of an agent's moral blameworthiness, since the lists of which alternatives truly are relevant and which alternatives we *think* are relevant might be quite different.

Ultimately, the five criteria given by Dretske are vague. All are good criteria in the sense that, in cases where they can be clearly applied, they do a good job of ruling out many irrelevant alternatives. However, the limits of application of each of them are unclear. Regarding the first criterion, sentence emphasis certainly has a role to play in determining context, as Dretske's examples show. However, there may be cases in which an alternative to a non-emphasized aspect of a knowledge claim statement might be relevant. Suppose a scientist were testing a sample, say XYZ, in some lab apparatus. She experiences problems with her results and, after running several failed experiments, makes the statement that “*The shutter* (on my apparatus) is ruining my sample.” Following what Dretske says, RA's to this claim will be that it is some other part of the apparatus that is ruining the sample. The nature of the emphasis means that we can assume that the sample is being ruined. However, suppose that the scientist has inadvertently been using a different

substance, UVW, instead of XYZ out of lack of due diligence. In this case, it would be a relevant alternative that the sample is not actually ruined at all! The scientist believes the sample to be ruined because she is achieving results that are inconsistent with XYZ, but these results would actually make sense with sample UVW. In this case, then, an alternative to a non-emphasized aspect of the sentence is relevant.

The second and third criteria likewise seem to be defeasible. There will certainly be cases in which it is neither legitimate to assume that the subject is truly who or what we think he, she or it is, nor to assume that our source is reliable. First of all, there may be circumstances under which it is not legitimate to assume certain things about the subject of our knowledge claim statement. Dretske would certainly agree that there are some human qualities that we cannot assume the subject possesses, such as, perhaps, a certain level of aesthetic appreciation. At some point there must be a line between qualities we can assume a person to possess and those we cannot, but the grey areas around this line are undefined. As for the assumption of the reliability of sources, there will be many cases in which we cannot trust a source. For example, it might not be legitimate to assume that any newspaper we read from is reputable. We should take into account the possibility of reporter error or bias. We have a new question, then: how do we know when we *can* make assumptions regarding our subject or our source?

The fourth criterion is also unclear with regard to the limits of its application. It is intuitively plausible that the standards for knowledge, and hence the number of

relevant alternatives, increase as the importance of the knowledge does. But the question that remains is what the nature of this correlation might be. How might we judge the relative importance of knowledge claims? Even given a plausible account of levels of importance, what account might we give of what sorts of alternatives become relevant as importance goes up? This is left very unclear. An alternative such as bias may not be relevant when I am telling you what the weather is like, but will be relevant when I am claiming that the climate is changing. Nonetheless, there will be a grey area encompassing cases where it is unclear whether the information is important enough for my biases to be relevant.

Finally, the fifth criterion is also an obviously vague one. It is not clear how remote an alternative would have to be before we could dismiss it as irrelevant. Dretske wants to argue that skeptical alternatives such as that of the apparent zebra actually being a cleverly painted mule are too remote to be relevant. At the same time, some alternatives are clearly relevant. For example, when I claim that *your female sibling* is tall in a situation in which I have never before met her, it is a relevant alternative that she is actually someone else, such as your friend who lives next door. In between these two cases, however, is a lot of grey area. Is it an RA to my claim that my friend is in her office at this moment that she is getting coffee? Or that she left early, even though she rarely does? Or that she had a family emergency? Or that there was a fire in the building? Some of these may seem to be RA's, but others may not be as clear.

### *Lewis' Criteria*

The defeasibility of the above five criteria does not imply that the criteria are incorrect. My point is, rather, that these criteria are vague, and hence more needs to be said about them to make the application clearer. We stand in need of some further criteria or discussion to indicate when we can, for example, make assumptions about a source or view an alternative as too remote. Another set of criteria is one proposed by David Lewis in his article "Elusive Knowledge." Lewis provides three rules governing which alternatives *cannot* be ruled out, three rules governing which alternatives *can* be ruled out and a final rule that he calls the "rule of attention" (554-559). Like Dretske's criteria, I think that Lewis' seven rules are generally defeasible and vague. In the first three rules, Lewis tells us that we cannot rule out: a) the alternative that reflects the reality of the situation (what really is the case), b) any alternative that the agent believes or should (in an epistemic, not moral, sense) believe is true and c) any alternative that saliently resembles an already relevant alternative (554-556). The second rule here is problematic in that it is far from clear in any situation what an agent *ought* to believe to be true. As an agent would likely only believe one alternative to be true at any given time, the first two rules do not leave much as relevant alternatives. However, they set up Lewis' third rule, which seems to be the most important of this theory: that any alternative that saliently resembles an already relevant alternative is also relevant (556). I will come back to this.



The next three rules Lewis provides give guidelines for determining which rules can be ignored. The fourth and fifth rules tell us that we are allowed to (very defeasibly) presuppose: a) the reliability of perception, memory and testimony and b) that a sample is representative and that the best explanation of a situation is the correct one (558-559). The sixth rule tells us that possibilities generally ignored in our epistemic community can be ignored (559). Lewis makes it very clear that he recognizes the defeasibility of all these rules. Further, however, the limits of application of rules five and six are vague. Opinions might differ as to what the “best explanation” of a phenomenon is, with some putting more emphasis on the simplicity of an explanation and others emphasizing the explanatory power. There may also be cases when two explanations seem about equal. As well, it may be less than obvious whether or not a possibility is being generally ignored in an epistemic community. For example, at some point in the past it may have been unclear whether or not the possibility of the Earth being flat was being generally ignored amongst peasants in Europe.

Lewis’ final rule states that once an agent pays attention to an alternative, then that alternative can no longer be ignored (559). This implies that at the moment I consider a skeptical alternative to knowledge, that alternative becomes relevant and can defeat my knowledge claim. This claim is important to Lewis’ argument that “epistemology destroys knowledge”: the very act of doing philosophy, which involves being aware of and considering skeptical alternatives, undermines our ability to make knowledge claims (559). This is an interesting position, which

questions whether or not any *argument* against skepticism could work. However, this aspect of Lewis' argument is not relevant to the current discussion, so I will leave it there.

Rule three gets at the heart of what Lewis believes it is that makes an alternative relevant, with the other rules serving to help pare things down. However, rule three is also by far the vaguest. Given rules one and two, rule three indicates that we cannot ignore alternatives that saliently resemble reality or that saliently resemble what the agent believes or should believe to be the case. The problem is that it is not clear what sorts of resemblance are "salient." Lewis believes that the concept of "salience" is clear enough to deal with some common problems with knowledge claims, such as Gettier problems and the lottery paradox. One of several examples that Lewis gives is the classic Gettier case in which I have a justified true belief that one of my two friends, Nogot or Havit, owns a Ford. I hold this belief because I think (based on very strong evidence) that Nogot owns a Ford, when in reality Nogot does not own a Ford but Havit does. Lewis argues that the reason I do not have knowledge in this case is that I have not eliminated the relevant possibility that Nogot drives a Ford he does not own while Havit does not drive nor own a car. This possibility is relevant, says Lewis, because it resembles reality in many salient ways. It resembles reality exactly with respect to Nogot, and very closely with regard to Havit's "carless" habits and the fact that people who exhibit carless habits are usually carless (557).

Obviously, in order for his arguments to work in refuting skepticism, Lewis

must be able to claim that an evil genius world does not saliently resemble the real world. For instance, I may claim to know that there is a water glass in front of me because I see one. Let us suppose that the actual situation is that there is, indeed, a glass in front of me. One possible alternative is that I am actually being deceived by an evil genius, who is implanting the image of a glass in my mind. This evil genius world resembles the actual world in many ways, in fact, perhaps in all ways except one, if we assume that the only action the genius is taking is to induce an image of a glass in my mind. Some of these resemblances, then, are: that I exist, that I am conscious, that I am capable of seeing things, that I see the rest of my apartment, that my hometown exists, *that I am having glass-type sensations*, and so on.

Given what Lewis says when discussing Gettier problems, I think that these resemblances are salient. The evil genius alternative resembles reality perfectly with respect to the rest of the world, just as in the Ford example the alternative of neither friend owning a car resembles reality perfectly with regard to Nogot. In the Ford example, there is a salient resemblance to reality in the aspect that differs from it: Havit exhibits carless habits. Similarly, in my glass example, there is a resemblance in the aspect of the alternative that differs from reality: I am experiencing glass-type sensations. I see no reason to consider the carless habits as salient but the glass-type sensations as non-salient. It might be objected that the reason that the glass-type sensations are non-salient is that their source is different in the alternative than in reality. However, Havit's carless habits also have a different source in reality than in the alternative: in the alternative, his carless habits would stem from not owning a

car, while in reality his habits may stem from, for example, a hatred of driving, a desire to save fuel, and so on. In the end, it seems to me that many alternatives that seem to be obviously irrelevant are not ruled out by Lewis' criteria because the concept of salience does not do the work that Lewis believes it does. No matter how strange or outlandish the evil genius alternatives might be, I think Lewis is stuck because of the resemblances that still exist between these alternatives and actuality.

### *Other Attempts*

Are there other accounts of "relevance" that are more successful? One suggestion, from Cohen, that seems to improve upon Lewis' third criterion, is that an alternative is relevant if and only if it is salient and sufficiently probable (Feldman 109). Once again, however, the application of this rule is vague. Whether or not an alternative is salient may be unclear, as may the question of how probable an alternative would have to be in order to be "sufficiently probable." Keith Lehrer argues that what determines what is relevant depends on which epistemological theory you are using, but that the context for determining what is relevant must be the world as we know it and not some cluster of nearby possible worlds. For example, Lehrer argues, for a reliabilist an alternative is relevant if it is germane to the actual reliability of some belief-forming process in the world as we know it (176). For example, the skeptical alternative of someone controlling your brain is not relevant because brain controllers are not part of our world and hence cannot affect the reliability of belief-formation in the world as we know it (176). Lehrer

seems to be begging the question here; the skeptic's challenge is that we cannot know that brain controllers do not exist, and Lehrer's reliabilist would answer this objection with an argument that depends on the premise that brain controllers do not exist.

There may be many possible variations on the attempts above to give an account of what makes an alternative relevant. All these attempts are alike in that they attempt to formulate some rule or set of rules that can be applied to a case where an agent claims to have knowledge in order to determine whether or not that agent has properly ruled out the relevant alternatives. In the examples that we have, the sets of rules invariably carry rules that are defeasible or in need of a debilitatingly long list of provisos, while the single rules are problematically vague.

## **Chapter 2: Virtue Ethics**

It is at this point that I suggest a shift to a new kind of approach to the question of what it means to say that someone “should have known better,” namely, a virtue approach. This approach will give us a different kind of answer to this question than those provided by relevant alternatives; an answer that I believe is more helpful.

One major reason to make this shift is parallel to one that many virtue ethicists have cited as a reason to make the switch from traditional, rules-based, ethical theories to virtue ethics. The claim is that, ultimately, rules-based accounts fail to succeed with respect to one aspect that is supposed to partially constitute their advantage over non-rules-based theories: providing concrete answers to ethical issues that can be applied to various situations. So, what is “virtue ethics”? The term itself does not refer to a specific theory, but rather is an umbrella term referring to several different kinds of theories. The use of the virtue approach stems from a reaction against traditional ethical theories (called “duty ethics”), which are consequentialist or deontological in nature and which define terms like “right” and “wrong” without reference to aretaic concepts, and that may go so far as to define aretaic concepts themselves in terms of acts, duties, and so on. Utilitarianism, a consequentialist theory, argues that the term “morally right” refers to acts that bring about the greatest amount of happiness to the greatest number of people. A utilitarian who wanted to discuss virtues might define a “virtue” as a trait that leads an agent to act in a way that brings about the greatest overall amount of happiness for the greatest

number of people. Virtue approaches, on the other hand, argue that agents, as opposed to acts or duties, are central with respect to judgments in ethics and that aretaic concepts are, at the very least, conceptually independent of acts, duties and so on (Statman 7).

There are two different general ways in which philosophers have used virtue in their theories. Following Driver, we can call these two approaches “virtue theory” and “virtue ethics” (Driver “The Virtues and Human Nature” 111). Virtue theory holds that aretaic notions are very important to ethics, and that they should be defined independently from the concepts of “acts” and “duties”. However, the virtue theorist does not claim that aretaic concepts are prior to moral terms like “morally right” and “morally good”. For a virtue theorist, a virtue is more than a disposition to act in accordance with a moral rule or duty, but virtues *can* be defined in terms of “right” and “good.” (I.e., we could say that a virtue is a disposition to behave in a way that is conducive to moral good, and then define “good” separately.) The virtue ethicist, on the other hand, not only argues that aretaic concepts are independent of concepts of acts and duties, but makes the further claim that concepts such “morally right” should be defined in terms of virtues. Roughly, then, to a virtue ethicist, something is morally good if it stems from virtues. In this thesis I am working from the perspective of virtue ethics, not virtue theory.

The move toward using virtue approaches to philosophical problems that has been going on in the last half-century began in the field of ethics. This move is actually a revival, as virtue approaches to ethics date from the ancient Greeks, and

Aristotle in particular. One of the major motivations in taking a virtue approach to ethics is the view that more traditional, principles-based, theories (or “duty ethics”) misrepresent what is truly important with respect to moral right and wrong (Brady and Pritchard “Introduction,” 1). Traditional theories focus on either the external features of an act, such as its consequences, or on some guiding principle, such as Kant's categorical imperative. These traditional theories represent moral competency as the ability to solve moral problems or dilemmas. However, virtue theorists argue, morality is about much more than these aspects; the agent performing the action and making the moral decisions, and specifically the agent's character, plays a crucial role that only virtue ethics can properly account for. Furthermore, proponents of virtue ethics argue, it has the added advantage of providing better reasons for being moral, as it entails that the life of virtue is the good life.

### *Aristotle*

Aristotle was one of the earliest philosophers to talk about virtue as the basis for an ethical theory. As the conception of “virtue” that I will be working from in this thesis will be an Aristotelian one, I will give a brief account of what Aristotle's views on virtue were, while leaving out his arguments in favour of his views. Aristotle discusses his virtue theory in *The Nicomachean Ethics*, where he seeks to determine what “the good” consists in. Aristotle distinguishes between things that are intrinsically good, or good in themselves, and things that are instrumentally



good, or good because they lead to goods or to some final or non-derivative good. Furthermore, the instrumental virtues can be broken down into those which are constitutive of the good and those that are not (1096b 10-15). Virtues are constitutively instrumentally good; not only do virtues dispose an agent to act in a way that will lead to the good life (*eudaimonia*), but the exercise of the virtues themselves is also constitutive of *eudaimonia* (1097b 5).

A necessary component of virtue is virtuous action; in other words, a person cannot be said to have a certain virtue unless he or she performs activities stemming from that virtue (1103b 30). However, doing virtuous acts is not enough to have virtues. A virtuous person must have a “firm and unchangeable character” from which virtuous acts proceed and must choose to perform virtuous action for its own sake (1105b 5). Humans are not born virtuous or vicious, according to Aristotle (1103a 25). Rather, these dispositions are acquired through habituation that begins in our childhood. We begin developing virtuous dispositions by “practicing” the performance of virtuous acts, and we require other persons with knowledge of virtue to habituate us to perform virtuous action. Since moral virtue is inextricably tied to emotion, these others “train” us to feel joy and grief about the right sorts of things (1104b 10). We are taught to feel joy regarding instances of justice, giving to others, and courageous acts, and to feel grief regarding injustice, avarice and cowardice. Eventually, we will develop the correct habits regarding how we regulate our desires, emotions and appetites and will form our dispositions accordingly. These dispositions, which are acquired through habituation and which

manage and channel our desires, emotions and appetites, are moral virtues.

The virtuous person also requires practical wisdom, or *phronesis*, in order to be virtuous. *Phronesis* involves being able to determine what *eudaimonia* consists in and hence being able to determine how to best employ one's virtues to achieve a good life and how to best employ the virtues in situations in which two or more actions that seem to be required by the virtues may be incompatible (1140b 15). Although, according to Aristotle, virtues themselves can never conflict, there may be cases in which our different virtues point us towards conflicting actions; in these cases, practical wisdom is crucial to making the correct decision on which course of action to pursue.

According to Aristotle, virtue is typically a mean between two extremes of a state of character. The extremes are vices; one extreme is an excess of a particular state of character, and the other extreme is a deficiency of that state (1104a 25). For example, consider pride. A certain degree of pride is a virtue. However, too much pride leads to arrogance and vanity, while too little pride leads to being unduly humble (1123b 5-15). Aristotle gives examples of many virtues, such as liberality, good temper, and truthfulness, which are all means between two extremes. He also describes the virtue of justice at some length. Justice differs from the virtues just mentioned in that it does not represent a mean between extremes; an agent cannot be *too* just.

To summarize, Aristotle characterizes virtues as lasting or permanent dispositions or states of character, developed through instruction or habituation,

which involve having the correct attitude towards situations the agent encounters, and that lead the agent to take correct action. In its current resurgence, there are many variations of virtue ethics being discussed and supported. Some of these theories are Aristotelian, while some are not.

### *Moral Virtues and Investigative Virtues*

As I indicated in my introduction, I want to argue that moral virtues are necessary for the agent in identifying and considering relevant alternatives in order to avoid moral blameworthiness for ignorance. However, there are other virtues besides moral virtues that are necessary to the agent when dealing with RA cases. These virtues are similar in structure to Aristotle's moral virtues; they are deep, permanent states of character that involve the agent being disposed to act and think in certain ways. I will be using the term "investigative virtues" to refer to these non-moral virtues for the reason that they have an important role to play in carrying out epistemic investigations. Many of these investigative virtues appear on current philosophers' lists of intellectual virtues, while others are not so easily classified. I will discuss the investigative virtues further down, but first the question of the role of classification of virtues in my discussion needs to be addressed.

The issue of classifying virtues into specific categories has been discussed by some contemporary philosophers. Linda Zagzebski, in her book *Virtues of the Mind: An Inquiry into the Ethical Foundations of Knowledge*, makes the argument that intellectual virtues are actually a subset of moral virtues and not a different

category of virtue at all. Her arguments are worth study, but the issue of exact classification of virtues is actually very unimportant to this thesis. My argument that certain sorts of virtues are necessary to properly identify RA's and avoid moral blameworthiness for ignorance does not involve making a claim that certain whole categories of virtues are necessary for identification of RA's. Rather, I want to discuss which sorts of virtues, taken from any category a person might wish to put them in, are necessary for the task. The category is unimportant; what matters is how the virtues operate in RA situations.

Nonetheless, I have made a verbal distinction in this thesis between moral virtues and investigative virtues. Hence, I stand in need of some justification for making that degree of distinction. The reason I have is simple, and has to do with one difference between moral virtues and the others I wish to discuss. Moral virtues, as discussed by Aristotle, invariably involve the regulation or direction of emotions, desires or appetites. Investigative virtues, on the other hand, do not. While moral virtues have to do with being a good sort of person with respect to moral actions, investigative virtues are pragmatic virtues that have to do with being someone who is good at formulating useful beliefs.

### **Chapter 3: Virtue Ethics and Right Action**

Now we come to the point where we can begin to examine how we might apply virtue ethics to the relevant alternatives issue. As I have indicated, I will argue that, in addition to moral virtues, other virtues that I have called “investigative” virtues will be needed for identification of relevant alternatives. However, before discussing what these investigative virtues might be, I want to take some time to discuss how my claim that certain virtues are required to identify relevant alternatives derives from the claims of virtue ethics. This discussion will address the issue of virtue ethics and right action, which will also provide me with a framework for demonstrating how virtues can be seen as being required in an epistemic context.

As I discussed above, there are two different ethical approaches that employ virtues: virtue theory and virtue ethics. Virtue ethics claims that aretaic concepts, such as the character of the moral agent, are conceptually prior to all other concepts in moral discussion. This focus on the agent instead of the act has drawn much criticism. One criticism is the claim that virtue ethics is not sufficiently action-guiding. An important part of ethical theory, critics argue, is determining not only how to assign moral praise or blame in a situation, but also determining what sorts of things we should and shouldn't do in the future. A good moral theory should be able to give us a way to determine which action we ought to take when we face a moral dilemma. For example, if we follow utilitarianism we can determine what to

do by performing a utilitarian calculation to figure out which action will create the most happiness for the greatest number of people. The critic of virtue ethics claims that this sort of action-guiding is necessary to a good moral theory. In fact, virtue ethicists themselves claim that a central aspect of morality or ethics is becoming a morally good agent, which includes morally good behavior. But, is virtue ethics truly sufficiently action-guiding to complete this necessary task of any moral theory?

Virtue ethicists have sought to address this criticism by providing virtue-ethical definitions of “morally right action.” For a virtue ethicist, a definition can still be given for a “morally right act,” it simply must be in terms of moral character. One such definition is provided by Rosalind Hursthouse, and is very close to the definitions used by Linda Zagzebski in *Virtues of the Mind*. Hursthouse states that an act is morally right “if and only if it is what a virtuous agent would do in the circumstances” (Hursthouse “Virtue Theory and Abortion” 225). Notice that the question of whether or not the agent who performed the act is herself virtuous does not come into play here. This implies that an isolated act can be considered morally right regardless of the character of the agent performing that act. This lines up with the possibility that a morally vicious, or at least non-virtuous, person could “fluke off” a morally correct act.<sup>4</sup> However, when evaluating this sort of case we would not consider the agent herself to be morally good, as she would not be virtuous.

---

<sup>4</sup> Additionally, it implies that a morally imperfect agent could attempt to improve her character by investigating and trying to acquire virtuous dispositions.

### *An Objection*

There is a criticism of Hursthouse's account, however, that led her to formulate a slightly different definition. It has been pointed out that Hursthouse's account faces a challenge if we consider the possibility that a virtuous person may still be fallible. In this case, the act that a virtuous person would do might be the wrong act in a case where the virtuous person is, for example, misinformed through no fault of her own (Das 331). For a truly virtuous person this case would be a rare possibility, as a part of the virtuous character is being able to avoid error, but it is still a possibility. This risk leads to Hursthouse's slightly modified account: "An action is right iff [if and only if] it is what a virtuous agent would *characteristically* (i.e. acting in character) do in the circumstances" (Hursthouse "Normative Virtue Ethics" 28, my italics). Under this definition, the rare situation in which a virtuous person is misinformed through no fault of her own does not challenge the definition, as characteristically she would not be so misinformed. This modified definition, however, is not free from criticism. Questions remain as to whether there might not be cases in which an agent could characteristically get a certain *kind* of moral problem wrong and, additionally, how we can know what a virtuous agent would do in any case.

Swanton formulates the first question by stating that if there is some defect of knowledge throughout the human race then any agent, no matter how virtuous, will characteristically fail to do the morally right thing (35). She gives the example of climate change. Swanton supposes that a species wishes to combat climate

change but is ignorant as to the best ways to do so because of a deficiency in scientific knowledge. In this case, a moral agent may characteristically take certain actions  $T_1$ - $T_n$  that, according to the best knowledge of herself and her species, will help prevent ecological disaster and the loss of lives. However, it may actually be the case that these actions do not prevent disaster but in fact help bring it about more quickly (35). In this case the virtuous agent cannot be blamed for lacking knowledge that her entire species does not possess but nonetheless she will fail to perform in a morally good manner. In fact, her virtuousness will lead her to act in the morally wrong manner that she believes is best.

What sort of reply might a virtue ethicist make to such an objection? She can deny the virtuousness of the agent in some way, change her definition of a right act or deny that the agent's actions in the climate change case are morally wrong. The first option will lead to some difficulties. How might we parse out the claim that the agent in the climate change example is not truly virtuous? An ideally virtuous agent is someone who, presumably, possesses all possible virtues. So, if the agent in the climate change example is not virtuous, then she must lack some virtue or virtues. What virtue, however, could it be that would allow her to gain a scientific understanding that her species is not yet capable of? Perhaps a certain kind of imagination would allow a virtuous agent to see solutions others miss. Nonetheless, even a great imagination combined with all the intellectual virtues possible could not allow an agent to, in a short period of time, overcome a deficiency of understanding of her species as a whole.



Alternately, one might answer that what is lacking is not a virtue itself, but the capacity to recognize the agent's limited knowledge, a capacity that is a feature of practical wisdom. However, we can respond to this possibility: if the species as a whole is incapable of knowing some information or of even recognizing that deficiency, then it is unreasonable to argue that any agent should be able to do what her species cannot. Recall that Swanton's example is set up so that the information the agent needs to make the decision that will actually prevent ecological disaster is *impossible* for her to attain. The second option, redefining "right action," may be more promising, but it is very open-ended. One concern with redefining morally good action would be to avoid a definition in terms of the consequences of the actions, as such an answer would be contrary to the basic tenets of virtue ethics.

A simpler answer than redefining morally good action, and one that is properly in line with virtue ethics, is to deny that  $T_1$ - $T_n$  are morally bad actions, and to assert that taking actions  $T_1$ - $T_n$  is the morally right thing to do, despite the consequences. This is really the very claim that the virtue ethicist wants to make: what is morally right is a person's *character* and, therefore, what she is characteristically prompted to do by her virtuous character is a morally good thing to do. The consequences of an action may not be good, but that doesn't change the fact that it was still morally right to do what virtue prompted you to. Following virtue ethics, morality has to do with what sort of person you want to be, where the right sort of person to be is the virtuous person. The virtuous person has certain attitudes, assignments of value, motivations and so on having to do with her virtues.

Being a virtuous person sometimes necessarily involves taking actions: there are cases in which being a virtuous person is incompatible with not taking a certain action. For example, take the virtue of generosity. A generous person is the kind of person who will give assistance to others who are in need, even at a cost to herself. If a person chooses not to take the action of giving to those in need even when she has a clear opportunity to do so (and no other reason to believe such an act would be harmful), then that person is not generous. In this way, being generous is incompatible with not giving to someone in need when you have a clear opportunity to do so. There may be a case in which this act of giving has negative consequences that were unforeseeable by the virtuous agent. However, this does not change the fact that a person possessing the virtue of generosity *would still give* in that case; doing so is a necessary part of being generous. So, we judge both the agent and her act of giving to be morally good. The later consequences may be negative, but the act itself was the morally right one under the circumstances.

Applied to the climate change example, the kind of person that the virtuous agent is will entail that she takes the actions that she has every reason to believe will help prevent catastrophe. Even if those actions lead to catastrophe, we still give the agent herself, as well as her actions, to the extent that they express virtue, a positive moral evaluation. I find this view to be very appealing, as it also has the consequence that we need not fear being morally blameworthy for negative consequences that we could not have possibly foreseen. It may be objected that this view is deficient in that it does not hold persons properly responsible for the

consequences of their actions. However, notice that after the negative consequence resulted from the agent's initial action, and was recognized by the agent as a negative consequence, this would constitute reason to re-evaluate such action before repeating it in future situations. At this point, virtues such as humility, perseverance and courage are very important; the agent must be humble enough to realize her error and be brave and tenacious enough to seek alternate routes. It is her responsibility, then, to be aware of and learn from the consequences of her actions. So, a person can be held fully responsible for her actions and be expected to apply the knowledge that she gains from those actions to future situations, whether or not those initial actions were morally right or morally wrong.

### *Das' Objection*

The second question involving the virtue theoretic account of morally right action still has to be addressed: how might we be able to know what sorts of things the virtuous person would do? Ramon Das sees a weakness in virtue ethics at this point, as he believes that any attempt to explicate the sorts of things that a virtuous person would do inevitably relies on having a concept of "morally right action" and hence renders definitions such as Hursthouse's circular (Das 332). He argues that a natural virtue ethical answer to the question of what sorts of things a virtuous person would do is that the virtuous person does things that lead to *eudaimonia*. However, Das then points out that the concept of *eudaimonia* might rely on an unexplained conception of "right action" in the form of Hursthouse's v-rules (Das 332). As

Hursthouse conceives of them, “v-rules” are rules of thumb that explicate the sorts of behavior that tend to be morally good (Hursthouse “Normative Virtue Ethics” 25-27). V-rules are often rules such as “Do not lie,” or “Help others in need,” and describe the sorts of behaviors that virtuous persons typically perform. A young person learning to be virtuous begins by following the v-rules until she begins to be able to recognize for herself what a virtuous agent would do. Das argues that this leads to circularity: a right action is one that a virtuous person would characteristically do, but a virtuous person will characteristically take right actions as described by v-rules.

To illustrate this supposed circularity, take an example of a situation in which a student has to decide whether or not to lie to get an extension on a paper. This student follows virtue ethics, and so asks herself (or perhaps a parent or friend whom she views as more virtuous than herself) what the virtuous agent would characteristically do in such a situation. She reaches the conclusion that the virtuous person would not lie and would instead face the consequence of incurring a late penalty on the assignment. On what basis is this conclusion reached by the student? According to Das, the process would be something like what follows. The student asks what the virtuous person would do, and finds that the virtuous person is the kind of person who would do whatever leads to the best possible life (*eudaimonia*). She further finds that the virtuous person will live the best possible life by following certain moral rules, including the rule “Do not lie,” that indicate what sorts of actions are morally right. Ultimately, then, the student decides that the morally right

thing to do is to tell the truth, for the reason that the action of lying is morally wrong.

There is an answer to the circularity objection that is in part formulated from comments Hursthouse makes in her 2001 book *On Virtue Ethics*. Das' argument is that virtue ethical accounts of right action are circular because "right action" is defined in terms of the concept of the "virtuous agent" which in turn is, ultimately, defined in terms of specific right actions. However, I think Das' arguments result from a misunderstanding of the role that v-rules actually play in virtue ethics. For Das' objections to have force, he must be able to show that the concept of a virtuous agent *logically* relies on a concept of "right action." The problem for Das is that virtue ethicists do not want to use v-rules as part of the definitions of "*eudaimonia*" and "virtuous agent" or even as a necessary part of these concepts. Rather, v-rules serve as guidelines and are helpful in education and communication about moral action. These v-rules, further, use not only "thin" rules such as "Do not lie," or "Do not steal," but also employ "thick" concepts such as "Don't be cruel," or "Be fair to each other," which do not suggest specific actions (38).

Nonetheless, the opponent of virtue ethics might demand, if the virtue ethicist does not want to give an account of what kinds of actions the virtuous agent will take in terms of v-rules, then what sort of account can he or she give? If the virtue ethicist rejects the argument that the virtuous person does what leads to *eudaimonia* and that *eudaimonia* is reached by following v-rules, then what argument does she offer in its place? Here Hursthouse has a couple of suggestions.

If we want to know what the virtuous person would do, she says, then we could simply ask a virtuous agent (35). Of course, we may never have a virtuous person on hand to talk to, but we still have resources for determining what a virtuous person would probably do. We know that a virtuous person can be described by a list of dispositions, such as honesty, generosity, courage and so on, and we can figure out what sorts of actions are in line with these characteristics (36). It is true that there will be cases in which we can't be certain what actions are in line with these traits, but virtue ethics has no more trouble here than traditional moral theories like utilitarianism and hence is no less action-guiding than these theories. In a tricky moral case, it will be no clearer which action will lead to the greatest amount of happiness for the greatest number of people than which action is most in line with virtue.

### *Right Action and Moral Education*

Accounts of how the concept of right action can be defined within virtue ethics will provide us with a framework for discussing the ways in which virtues are required for practical wisdom and, hence, right action. On the face of it, it may seem that the statement that virtues are required for right action follows from the central claims of virtue ethics. If the right action is, by definition, what the virtuous person would do, and if the virtuous person is a person who possesses and exercises the virtues, then it seems to follow that virtues are a prerequisite for right moral action. However, as was indicated above, the virtue-ethical definition of "right

action” leaves open the possibility that a person not in possession of the virtues may still do the action that a virtuous person would do. Given the reality that most people are not perfectly virtuous, virtue ethics must either state that very few actions performed by human beings are actually right or accommodate the possibility that a non-virtuous person can still perform right actions. Since the claim that very few human actions are the right actions is an unappealing one, virtue ethicists typically take the latter point of view. In fact, this circumstance is an important component of moral education, which in turn is a significant part of virtue ethics.

Given the virtue theorist’s views on moral education, then, it may seem that virtues are not required for right action, at least not in individual cases. However, the process of moral education itself requires that there be persons who do possess sufficient practical wisdom to recognize the importance of specific virtuous dispositions, whether the agent in question fully possesses those virtues or not. For example, I might not need to possess the virtue of generosity in order to decide to give money to charity and for this action of giving money to be a right action for me to perform. However, at some point in my moral education I will have known persons who, although probably not possessing the virtue of generosity in the sense of being ideally generous persons, were nonetheless more virtuous than myself in that regard and possessed the practical wisdom to recognize in what situations a generous disposition ought to be exercised. I might have learned about generosity from these persons, who would be recognizably more virtuous than me due to their similarity to properly virtuous agents. When I make this decision to give to charity,

the influence of those persons is instrumental, whether or not I am consciously thinking of them or trying to be like them. To put it simply: at some point I must have learned about generosity from someone (my parents, teachers, friends) and am now taking an action such as giving to charity as a result. In short, the existence of persons with sufficient practical wisdom to see the value of having such a disposition is necessary for the moral education of others and the performance of right acts by others. As practical wisdom is a virtue itself and requires possession of at least some virtues, it follows that virtues are necessary for right action.

Given this account, any right action that is done by a virtuous person or by a person in the process of becoming more virtuous will require, or stem from, virtues and practical wisdom regarding *eudaimonia*. Nonetheless, there might still be cases where someone who is not virtuous takes an action that a virtuous person would characteristically take in similar circumstances but not because of any influence by a virtuous person. For example, someone might give to charity in order to make a good impression on someone whom she hopes might give her a job, or because she had some extra money and decided arbitrarily to give it away (say, she flipped a coin) or because she was coerced into doing so. In all of these cases, giving to charity might be what the virtuous person would characteristically do and hence would be the “right action.” However, virtue did not play a part in any of these cases. Do examples such as this undermine the thesis that virtue is required for right action? Not in any significant way.

What the example above of a person giving to charity for reasons unrelated



to any virtue demonstrates is that virtues are not required for every individual right action. A person can perform isolated incidents of right action without virtues. However, surely the goal of morality is not to “fluke off” morally right actions in single cases but to consistently perform morally right actions and avoid morally wrong ones. The more important question, then, is whether virtues are required for *consistent* morally right action. While the thesis of virtues being required for all individual right actions was able to be refuted by a single example, the case is more complex with respect to consistent right action. Presumably, though, a compelling example of a person consistently performing right action and avoiding wrong actions without any influence of virtue would damage the thesis that virtues are required for right action. Can we think of such an example? Earlier, I gave three different examples of non-virtuous reasons for acting in an individual case that led to right action. Would consistently acting from these reasons lead to consistent right action?

First, consider the person who acts in order to impress someone whom she hopes will give her a job. As it is unlikely that she will take many actions from this specific motive (i.e. getting the job), let us say that she typically or consistently acts in order to impress others so that she will benefit from doing so. It seems in this case that the nature of her actions will depend solely on whom she is trying to make a good impression on. If the persons she tries to impress are uniformly impressed by consistent virtuousness, then she will indeed likely tend to consistently do right actions *if* she is consistently able to correctly gauge what the person she is trying to

impress will be impressed by.

We could object to this example on grounds the obvious low probability that every person the agent wants to impress is impressed by consistent virtuous action, and the low probability that she will consistently be able to figure out what will impress. However, I think that the example fails even if we grant that these things are possible. If it is the case that every person the agent is trying to impress is only impressed by virtuous action and the agent can correctly gauge what will impress, then is it true that virtue has no role to play here? The persons who are only impressed by virtuous action must, at the very least, be influenced in some way by virtuous agents, if they are not somewhat virtuous themselves. Further, if the agent herself is able to consistently recognize what she must do to impress these people, i.e. the virtuous action, then she must also be familiar with virtuous persons. Even if her motives remain unvirtuous (i.e. to selfishly further her own ends), virtues have had an influence on her choice of action, if only indirectly.

Next, let us look at the case of the agent who makes decisions regarding what action to take in a completely arbitrary manner, say by flipping a coin or rolling a die. In this case it seems quite clear that she will not consistently perform right actions. For one thing, the right action might not even be one of the actions that she is choosing between or among. Even if the right action is among these choices, odds are that she will only pick it a percentage of the time, and hence will not *consistently* perform right actions and avoid wrong ones. As for the case of coercion, it doesn't seem reasonable for there to exist any agent who only acts under

coercion. However, assuming this to be possible it is still highly unlikely that those persons doing the coercing would consistently coerce the agent into doing the right action. However, if the coercers did consistently coerce the agent to do the right action then, similarly to the first case, those doing the coercing must be under the influence of virtues.

### *An Illustration*

I want, now, to give a brief illustration of how virtues lead to the consistent performance of right action. Specifically, I will give an account of how virtues are required for consistent right action in the case of a teacher dealing with her students. I hope this discussion will make clear how, under a virtue ethical conception of “right action,” virtues can be shown to indeed be necessary for consistent right action. My methodology here will be to first give an account of the sorts of situations in which the teacher will face moral decisions, then to go through a list of character traits typically considered to be moral virtues and demonstrate how these traits will lead the teacher to be the sort of person who performs moral actions by looking closely at two of them.

Teachers face many moral or ethical decisions in their dealings with students. I will consider here teachers of adult students, say in a university setting, since this setting is one I am more familiar with. The ethical choices faced by a teacher of adults will be somewhat different in the case of a high school teacher, or an elementary school teacher. The teacher of adults must ensure that her

expectations are made clear and are fair, she must treat all of her students equitably regardless of gender, race and so on, she must listen to her students and treat their ideas with respect, she must mark them fairly, she must judge when it is right to allow students extensions on their assignments, she must be able to determine when the students might be cheating or plagiarizing while allowing these students the opportunity to defend themselves, and she must be able to deal with any other ethical issues that might arise.

Throughout the remainder of this thesis, I will make reference to several moral virtues. I will not even attempt to provide a full and complete list of the moral virtues here, but will use several virtues that appear on many lists, and hence by consensus are indeed moral virtues.<sup>5</sup> I think that these selected virtues run the gamut of kinds of virtues and will be sufficient for my purposes. The virtues I will consider are: courage, justice, generosity, thoughtfulness, compassion, dependability, industriousness, honesty, self-control and self-confidence.<sup>6</sup>

Courage is considered by most philosophers to be one of the core moral virtues. Briefly, the courageous individual is the one willing to put herself at risk or in danger of some sort in order to take some action that she judges to be the right one. At the same time, a person willing to throw herself into severe danger for a

---

<sup>5</sup> Not all of these moral virtues will appear in Aristotle's discussion. However, they all look like Aristotelian moral virtues, in that they are deep, unchanging, states of character that have to do with the regulation of desires, appetites, and emotions.

<sup>6</sup> My discussion of the specific moral and epistemic virtues here and elsewhere will draw from selected writings of: Aristotle, James Wallace, Linda Zagzebski, Rosalind Hursthouse, James Rachels, Michael Slote, Lorraine Code, Edmund Pincoffs, Jonathan Dancy, Lawrence Bonjour, James Monmarquet, Juli Eflin and Richard Paul.

more trivial cause is reckless and not courageous at all. So, courage represents a mean between cowardice and recklessness. There is, not surprisingly, debate about difficult cases in which it is not clear whether the agent concerned is courageous or not. As well, what counts as a normal, everyday action for one person may be a courageous act for another. For most people, stepping into a crowded bus is not an act of courage (to make this a moral action, let us say that the action of stepping onto the bus is for the purpose of delivering a warning to the passengers), but for someone who suffers from claustrophobia it might be an extremely courageous act indeed.

Courage is necessary for the teacher who wishes to act ethically in all her dealings with her students. As a teacher, there is a lot of incentive to be “the good guy” or the “cool” teacher who is well-liked by her students. Going easy on the students, giving fewer assignments, or leaving out more difficult course work are typically actions that will lead to a teacher being well-liked by students, while giving out accurate (often lower) marks, giving out more difficult assignments and being strict about deadlines are all things that can lead to a teacher being disliked. However, doing these latter actions (tempered with other virtues, such as justice and compassion) will provide the students with a better education. It takes courage, then, for a teacher to risk being disliked by the students in order to do the right thing by them and provide them with a proper education. It also takes courage for a teacher to admit any mistakes she might have made. Teachers do sometimes make mistakes in marking, or might even make a factual error while presenting material to

the class. To admit to their students that they made an error is necessary for rectifying that error, but may take courage, as the teacher might fear appearing weak or incompetent to her students.

The next virtue I wish to consider is generosity, a virtue that involves being kind and caring towards another person, sometimes through giving something that has value to yourself to that other person in order to benefit her. The extremes in the case would be giving little or no kindness, or stinginess, and giving too much. Generosity might play a role for a teacher in certain cases: if a teacher puts in extra time outside her job description to help students out, it could be said that she is being generous by “donating” time to her students that she could otherwise spend on her own work or on leisure. The generous teacher will also tend to give the student the benefit of the doubt when giving assignments and marks. For example, if a generous teacher becomes aware that a student has had unexpected difficulties that week, she may decide that the student should be allowed a few extra days for her paper, even though this may pose an inconvenience to the teacher. A generous teacher may become aware that a student is facing some emotional turmoil in his life, and then act to assist that student, perhaps by giving him some leeway with due dates or by agreeing to meet with him to go over some course material. Generosity may also be exercised towards an entire class; the class average on an assignment may be very low, and so the teacher might generously decide to raise the marks. Justice will probably be involved in such a decision in that the teacher may decide that a whole class full of bad marks might mean that she didn't teach the material

properly or that she made the assignment too difficult.

There are a couple of things that I would like to take note of at this point. The first is the interconnectedness of the virtues. Rather than each virtue being solely responsible for moral action in a certain set of cases, the actual situation is that several virtues working together will be necessary for moral action in any given type of case. It is obvious, then, that possession of a single virtue or couple of virtues is not sufficient for being a virtuous agent. In fact, it has often been argued, notably by Aristotle, that a person must be in possession of all virtues in order to be truly virtuous. I don't want to address the arguments for and against this position here, but the position needs to be mentioned. It may not be clear when, exactly, a person ceases to be "virtuous-in-training" as it were and becomes truly virtuous, but it doesn't need to be made clear for the purposes of this discussion. An agent may exercise virtues, and hence avoid blameworthiness for ignorance, whether she can properly be said to possess those virtues or not.<sup>7</sup>

The above paragraph leads to the second thing that I would like to take note of: the way in which lack of a single virtue might affect moral action. Although all the virtues are interconnected, it does not necessarily follow that lack of one single virtue will mean that the agent performs no morally right actions. The other virtues may be enough to ensure that the agent still acts in the way that the virtuous person would in like circumstances. For example, a teacher may lack compassion for her students but still be just and thoughtful. This teacher is still likely to make the

---

<sup>7</sup>

Refer to my discussion of Hursthouse's definition of "right action," above, pages 31-32.

morally right choice in, for example, raising the overall average in a test that the entire class scored poorly on (in a case, we assume, where it is obvious that even the most dedicated and intelligent students who were keeping up in class scored poorly). She may not feel compassion for the students or desire to alleviate their feelings of stress and shame regarding the test, but her thoughtfulness will make her aware that they all will now have difficulty in getting a good mark in the course despite their level of work and understanding of lectures, her honesty and courage (in conjunction with her reason) will lead her to conclude that she might have made the test too difficult, and her desire for justice will make her decide that she ought not to punish her students for her own error. Compassion in this case might have provided the teacher with an extra motivation to change the marks, but the result would be the same.

However, in the case given above, if we remove different virtues from the list of those possessed by the teacher, the end result might be different. If we remove the virtue of thoughtfulness, then the teacher might not be fully aware of the situation her students are in, and so would have nothing to feel compassionate about and would be unlikely to see that perhaps she ought to act to change anything. If we remove honesty, then, even if the teacher is thoughtful and compassionate, she might not be willing to admit the possibility that she made an error and hence would not see that there is an injustice to be rectified. If we remove both thoughtfulness *and* honesty then it is even more unlikely that the teacher would recognize the situation as one in which the students have been put in a bind due to her own error.



My point here is that, although it is not the case that every single virtue is needed in every case of moral action, we can nonetheless quickly begin to see how lack of one or two virtues can lead to an action different from what we define as “right action” under virtue theory. Further, since I wish to discuss consistent moral action over a variety of cases, it doesn't matter if there are some virtues that are not strictly necessary in certain cases: each virtue will be necessary in some case somewhere, and hence each virtue is necessary for consistent moral action.

#### **Chapter 4: Virtue and Relevant Alternatives**

I now want to move on to the central problem of my thesis: when we “should have known better,” or when we are morally blameworthy for our own ignorance. As I indicated earlier, I wish to take a virtue ethical approach to this problem, and argue that the cases in which we should have known better are cases in which we did not believe justifiably according to virtue ethics; we did not believe what the virtuous agent would characteristically believe under like circumstances. Another way to put this is that we did not believe according to the necessary virtues. The question for the relevant alternatives issue, then, is what are the necessary virtues? I have said that the necessary virtues will be both the moral virtues and the investigative virtues. Now, I need to give an account of what these investigative virtues might be.

Before looking in more depth at what kinds of virtues are needed to properly identify and rule out RA's, I first would like to provide an account of the different kinds of RA cases that can be encountered. The possible cases in which an agent faces the possibility that she “should know better” are infinite, but a discussion of some of the kinds of cases where she does face it will facilitate my discussion of how the virtues relate to RA's. There are three factors regarding RA's that are important in identifying the different kinds of cases. One factor is the agent's situation with respect to the information that is relevant to the case. One possibility here is that the agent has yet to investigate the case, meaning that she still needs to

ask questions, look information up, run an experiment or go through some other process in order to identify the possible RA's and rule them out. A second possibility is that the agent believes (provisionally at least) that she already possesses the relevant information "in her head." In these sorts of cases, she must first be able to determine, from the information she possess, which alternatives are relevant, and then she must try to determine whether or not there may be alternatives she has not even considered. Both of these possibilities constitute a kind of investigation or inquiry, and so investigative virtues will be needed for both.

The second factor is the type of information that the agent is claiming to know. Some of the different kinds of information that an agent might claim to know are: mundane inferences about daily life (such as "There is a tree in my yard,"); factual information about the state of the world ("There is a war going on in the Middle East," or "The Earth is the third planet from the Sun,"); facts or theories about some event that has happened in the past ("Julius Caesar was stabbed by Roman senators to prevent him from becoming emperor," or "My grandmother was born in Budapest,") and theories about how the world works ("All matter is made up of atoms," or "Humans evolved from more primitive life-forms,"). Once again, some form of investigation has a role to play in all of these kinds of knowledge claims.

There also appears to be a third factor in determining the kind of situation an agent may find herself in: an agent might be assessing a knowledge claim that is the result of her own investigations and deliberations, or she might be assessing a claim

that is the result of another person's (or several other persons') investigations and deliberations. However, it is important to remember here that virtually no knowledge claims will be based entirely on either the agent's own deliberations or the deliberations of others. Arguably, statements such as "There is a tree in the yard," do not require any input from another person.<sup>8</sup> However, other kinds of knowledge claims do rely on others. A scientist who investigates and develops a theory herself still must rely on information and calculations gathered and made by others. Generally speaking, when we deliberate on an issue or investigate a question, we rely on facts learned from other persons, methods we use that we learned from others and so on. On the converse side, when we assess a knowledge claim made by some other or others, it is we ourselves who must assess the quality of their investigations and deliberations, and so this RA situation is not completely reliant on the work of other persons. Both cases are examples of kinds of inquiries.

Different combinations of these factors will lead to different kinds of situations with respect to the ruling out of relevant alternatives. Some examples of these different kinds of situations, then, can be identified. I will provide three. One type of case involves an agent investigating what happened in some past event. Examples of this kind of situation are: a historian trying to determine what events led to the abandoning of a village six hundred years ago or a detective trying to determine who committed a crime. A second kind of case is one in which the agent is asked to mediate between one or more people who are in disagreement. Examples

---

<sup>8</sup> It might be argued that the agent is relying on others for knowledge of the correct usage of the terms involved.

of this situation are: a jury attempting to decide whether or not the accused is guilty or a parent trying to determine which of her children broke a light fixture. In these cases, the mediating agent must determine whether the alternatives being presented are exhaustive of the RA's. A third kind of situation is that of confirming a hypothesis, either one's own or another person's. An example of this situation is a scientist attempting to corroborate a colleague's theory.

### *The Virtues Required to Avoid Morally Blameworthy Ignorance*

In order to be able to avoid situations in which she is morally blameworthy for her ignorance, an agent needs to possess certain virtues. First of all, the agent will need to possess many moral virtues, such as courage, generosity, justice and compassion. Although I think the role that some of these virtues have to play is fairly self-explanatory, I will illustrate the role that moral virtues can have in inquiries by discussing two of these virtues: courage and generosity.

Moral courage is characterized as the willingness to put oneself at risk of some sort in order to do what is morally right, often to benefit others. Proper consideration of propositions and arguments can put the epistemic agent at risk. For example, consider a detective who is in the process of investigating a series of murders which she believes were committed by a mob boss. The detective is attempting to justify her belief that it was the mob boss who is responsible for the murders. Investigation to discover more information about the case is needed in order for the detective to determine what the relevant alternatives to the mob boss

ordering the murders are before she can rule out those alternatives. However, if the detective is correct in her belief, then the very action of investigating the murders will catch the attention of the mob boss, who may decide to arrange the death of the detective who is close to the truth. So, in order to rule out the relevant alternatives in this case, the detective requires sufficient moral courage to put herself at potential risk. If the detective lacked this moral courage or if she possessed the vice of cowardice, then she might decide instead to make her accusation quickly, before her belief in the boss' guilt was justified, in order to ensure that the boss was arrested, or even to accuse someone else altogether.

The next moral virtue to consider is that of generosity, which involves being kindly disposed to others, and is sometimes exercised through giving to those others. This does not at first glance seem to have much to do with epistemic investigation. However, there will be ways in which generosity might be needed in order to identify all RA's. There may be situations in which, in order to continue an epistemic investigation into possible RA's, a personal sacrifice is required. For example, consider a scientist, Dr. Jones, who is attempting to discover which theory regarding a biological mechanism is correct. If Dr. Jones can identify all the relevant alternatives and in so doing come to a justified belief about which theory is best, then there is a possibility that she will be able to use this theory to develop a vaccine for a disease. In order to finish her research, Dr. Jones has devoted all of her time to working on the project and has turned down a higher-paying job in order to finish her work. In this case, in order to identify all the alternatives on the way to

forming a justified belief regarding which theory is correct, the scientist has sacrificed something of value to herself, namely, a better-paying job.<sup>9</sup>

In this case, possession of the virtue of generosity will make the scientist willing to give up a significant amount of money to continue her investigation. If the scientist is not generous, then she might decide that the slim chance of being able to come up with a vaccine is not worth continuing the investigation and turning down the better salary. She might conclude this even if she possessed the virtues of curiosity and perseverance, which will be discussed in more depth shortly. In such a case, the scientist's inclinations to seek knowledge and to continue to pursue knowledge until she has reached justified belief would be overridden by the desire to have a better job. The case of vice in this case is clearer: if Dr. Jones possessed the vice of stinginess or greed, then she would almost certainly drop her research and take the money. Generosity affects the agent's ability to identify RA's, then, by changing the relative values the agent places on investigation and finances, and hence determining which sorts of propositions she will be disposed to consider.

I want to note something more about the Dr. Jones example and how it relates to RA's. It may seem that my argument does more to show that moral virtues are needed to do the morally right action than to show that moral virtues are needed

---

<sup>9</sup> I am assuming a) that the scientist is not working under the assumption that being successful in her research will lead to some monetary award and b) that she wishes to finish her research for the purpose of preventing deaths from the disease in question, and not out of a desire for personal glory.

to identify RA's. It is true that the examples I have used (those of the detective and of Dr. Jones) involve a person doing what is arguably the morally right thing because of the influence of moral virtues. However, it also the case that the morally right thing to do in these cases is to identify (and rule out) the relevant alternatives. In my examples, if the agent were to decide not to identify and rule out the RA's in question and if, as a result of this, a negative result were to take place, then the agent would hold some degree of moral responsibility for that result. In other words, these cases fit the template I gave at the beginning of my thesis.

I have asked when a person can be considered morally blameworthy for their ignorance, and the examples I am using in my current discussion involve such cases. For example, consider Dr. Jones. Suppose that the correct theory regarding the biological mechanism is the one that leads to a vaccine, but that Dr. Jones has chosen to abandon her research of the mechanism in order to take a better-paying job. If, in a period of time after which the vaccine could have been developed and used as a result of her research, several people die, then Dr. Jones has failed to do all that she might have done to avoid those deaths could be considered blameworthy for them. If another scientist, many years after Dr. Jones worked on this mechanism, were to discover the vaccine and also to discover that a scientist had been working on this project but had ceased her research, then this scientist might say something like "That Dr. Jones should have known better than to cease her work on the mechanism. People died because of her failure to finish this research." In the Dr. Jones examples, it was a moral failure that led her to choose to focus her epistemic



energies elsewhere despite her awareness of the possibility of a vaccine.

### *The Investigative Virtues*

Over and above moral virtues, the agent wishing to follow the proper process of investigation will also have to possess a number of virtues that are not moral virtues in the Aristotelian sense, even though they bear many structural similarities to moral virtues. Like moral virtues, these other virtues are states of character that are acquired through habituation and that we can formulate v-rules to help with the habituation process. However, while moral virtues have to do with the direction of emotions, appetites and desires, these other virtues do not. These non-moral virtues can be characterized by the fact that they are primarily pragmatic and have a crucial role to play in the process of developing knowledge claims. For the purposes of this thesis, I will refer to these virtues as “investigative virtues,” although the name is not crucial and, as I mentioned above, many of them appear on some philosophers' recent lists of “intellectual virtues.” The agent who possessed all of these virtues could be considered to be “practically wise” in the Aristotelian sense.

A list of the investigative virtues might look something like this:

1) Virtues of capacity:

~ ability to properly understand arguments, theories, etc.<sup>10</sup>

-adaptability of intellect

-ability to represent different points of view and to give a coherent explanation of the facts

---

<sup>10</sup> This subcategory of virtue is more properly thought of as the disposition to exercise the relevant ability, as the ability itself could perhaps be seen as being a skill rather than a virtue.

-synoptic grasp of different domains; understanding of  
persons, problems, theories

~ imagination

2) Virtues of objectivity:

~ fairness and impartiality when evaluating others' arguments  
-open-mindedness in collecting and evaluating information  
and arguments

-epistemic integrity

~ epistemic empathy

~ humility

3) Virtues of procedure:

~ conscientiousness

-thoroughness; perseverance; care

-recognition of salience

-recognition of reliable authority

~ curiosity

-intellectual sobriety

The investigative virtues that I would like to focus on are conscientiousness, imagination, curiosity, empathy, humility and fairness/impartiality. These virtues are not among those Aristotle discusses, but are key for the process of inquiry and for avoiding cases in which one “should have known better.” The investigative virtues are related to moral blameworthiness because they dispose an agent to think and act in certain ways with respect to carrying out investigations that may have moral consequences.

The first virtue I will discuss is conscientiousness. The conscientious agent is someone who makes sure that she thoroughly and carefully seeks out and considers evidence. Further, she is capable of recognizing which facts are salient and which are not, as well as when a supposed authority is legitimate. To recognize the importance of conscientiousness to investigation, consider the following

example. Consider the case of Dr. Jones, the scientist who is attempting to discover which theory regarding a biological mechanism is correct. If Dr. Jones can complete her investigation and in so doing come to a justified belief about which theory is best, then there is a possibility that she will be able to use this theory to develop a vaccine for a disease. If, however, Dr. Jones fails to be fully conscientious, the results of her investigation might be very different. If she were, for instance, to fail to take enough care with her interpretations of her experimental results or if she were to ignore certain salient factors, she might come up with an inaccurate theory about the mechanism and fail to recognize the optimal theory, which might be used to develop a vaccine. Here, Dr. Jones' lack of conscientiousness will have led her to form a morally blameworthy belief.

Another important characteristic of the conscientious person is that she will want to make certain that the manner in which she undergoes any investigation is not disruptive or harmful to others. This will include ensuring that she is not neglecting other persons or her duties in order to complete the investigation. For example, suppose a graduate student is attempting to finish a major paper and devotes the majority of her time to it. However, by spending most of her time in the library, in her office and so on, she ignores her responsibilities to her friends, family or roommates. She forgets her father's 60<sup>th</sup> birthday, forgets to come help her friend move, and fails to do her share of the chores by forgetting to buy essential groceries without informing her roommates that she has not bought them. Furthermore, she neglects her own health and deprives herself of sleep. A person who acted in this

manner would not be exhibiting the virtue of conscientiousness, because she has allowed her epistemic pursuits to negatively affect other persons as well many aspects of her own life. There is another, perhaps more direct, way in which an agent must be mindful of others when performing inquiries. She must make sure that she does not insult another person, cause them physical harm or violate their other rights, such as privacy, when conducting an investigation. Conscientiousness, then, involves making sure that the right amount of effort and care is put into forming knowledge claims in a manner that is not damaging to the agent herself or to others.

Consider next the virtue of curiosity. The curious person will want to find out answers to questions that she encounters, increase her knowledge of numerous fields or issues, and learn about things that she does not yet understand. Like the virtue of conscientiousness discussed above, and like many of Aristotle's virtues, the virtue of curiosity involves finding a mean between extremes. The virtuously curious person will be motivated to inquire into many questions, but will also exercise intellectual sobriety: she will restrain and focus her curiosity. She will not spend the same amount of time and effort researching every single question or issue. Rather, she will tailor her level of curiosity to each potential inquiry. For example, during an election the virtuously curious agent will not expend much of her curiosity on the question of which political candidate most closely shares her musical tastes, but will expend a lot of her curiosity on the issue of which candidate most closely shares her views on health care policy.

The lack of curiosity can cause an agent to fail in an attempt to properly complete an investigation. Once again, consider the case of Dr. Jones, the scientist occupied with studying a biological mechanism. Her curiosity about the nature of the biological mechanism and its potential medical applications will lead her to expend more of her resources on the issue, to understand it completely, and to stick with her investigation even if she faces some adversity (such as failed experiments). However, if she were to lack this curiosity, she may not be as concerned to understand the mechanism fully and to apprise herself of all of its applications, and this might lead to her abandon her investigation or cut it short after a handful of failed experiments. The curious character is a highly motivated one, but one that is motivated in the right sort of way to explore interesting questions.

Another investigative virtue that is of crucial importance is that of empathy towards others. This may look more like an Aristotelian moral virtue, but it is distinct. Empathy might involve the capacity to feel and will almost certainly cause emotions to arise in many cases, but it is not concerned with regulating those emotions. Rather, it involves a certain kind of cognitive attitude towards the situation of others. The person with the virtue of empathy will be able to recognize what other people are experiencing and what their challenges or problems are and to feel what those others are feeling. Empathy can also be epistemic; the epistemically empathetic agent will be better able to understand the arguments others are trying to make and better appreciate the value of their point of view. Hence, she will be better able to evaluate the epistemic contributions of others and will be better able to

decide which arguments are worth her attention.

So, empathy will help the agent to appreciate and understand others' feelings and points of view. This virtue is very important to inquiry in that it allows the agent to properly understand and weigh arguments and input coming from other persons, and thus to include important information in her investigation.

Furthermore, the virtue of empathy may lead the agent to pursue certain lines of inquiry she may not have otherwise pursued for the reason that the result of such lines of inquiry may be beneficial to others. In the case of Dr. Jones, if she were to possess empathy for those persons suffering from the disease she could potentially find a cure for (and their families), she would be more disposed to continue her investigation into the biological mechanism in the best possible way.

The next investigative virtue to discuss is humility. The virtue of humility concerns the agent's estimation of her own abilities. Like other many other virtues, humility represents a mean between two extremes. The humble agent will recognize which areas she is particularly skilled in, and those situations in which she is likeliest to be correct. She will tend to trust her own judgment more in those areas, rather than indulging in crippling self-doubt. On the other hand, she will also recognize her limitations, and be aware of those areas in which she knows little or has little expertise. In these areas, she will be less likely to rely solely on her own judgment. The agent who possesses too much or too little humility may be hindered in her inquiries.

Take the case of Dr. Jones. When studying the biological mechanism that

she is concerned with, she must possess the proper estimation of her abilities as a scientist. If she were to estimate her abilities too highly, she might decide not to confer with colleagues to get their opinions on interpretations of her experimental results, and she might even reject or ignore differing results from others from a belief that her skills were superior to theirs. If, however, she has too low an opinion of her own abilities, she might ignore unique ideas that she has out of a belief that they couldn't possibly be correct, she might rely too much on her other colleagues and hence ignore situations where she has potentially done something right that others have done wrong, or she might even be inclined to abandon her work. In any of these situations, her lack of the virtue of proper humility may interfere with her investigation; it may lead her to ignore important information or arguments which would lead her to be aware of new alternatives that are relevant.

A further virtue is that of fairness/impartiality. This virtue involves the ability to assess different ideas, arguments and theories without the influence of bias. When evaluating others' positions, the impartial agent must evaluate these arguments based solely on their reasonableness and not on any opinion she has of the persons putting forward those arguments. Furthermore, when collecting and evaluating information and arguments she must consider those from all sources, regardless of her liking or disliking those sources. She also must not allow her own personal preference for one theory (or argument, etc.) to lead her to unfairly favour that theory over others. Furthermore, the fair and impartial agent will also be a person of epistemic integrity: she will make every effort to avoid self-deception and

will not allow herself to be convinced to accept or reject an argument or some evidence because of any non-epistemic consideration.

Fairness and impartiality function similarly to humility in their influence on epistemic inquiry. Lack of fairness and impartiality may lead the agent to ignore certain salient propositions, arguments and so on, in favour of others that may or may not be salient. She may judge one proposition to be more salient than it is, and may judge yet another to be less salient than it is. Without fairness, the agent's investigations may be skewed one way or the other. Unlike some of the virtues mentioned earlier, fairness does not seem to be a mean between extremes. It is hard to see how someone could be *too* fair or *too* impartial in the process of investigation. While in other moral situations it may be allowable or necessary to be partial (for example, a parent should perhaps be partial to her own children and give them more of her time and effort than she does, say, her coworkers), in investigations impartiality is typically required. The agent should only view one proposition as stronger than another if epistemic considerations indicate that it *is*; in this sort of case, it would not be right to say that the agent is "partial" to this proposition in the sense that she is biased towards it.

The final investigative virtue, which I have left until last because it is different than the other investigative virtues discussed, is imagination. Imagination may seem to be something that is supererogatory; we do not typically expect people to be imaginative in order to be considered practically wise or morally good. However, a closer look at imagination will show us how important it can be for



inquiry. Much of the writing that has been done about imagination is in the context of imagination's role in aesthetic observation or appreciation. Emily Brady, discussing imagination in "Imagination and the Aesthetic Appreciation of Nature," lists what she considers to be four modes of imaginative activity with respect to nature. Imagination, Brady argues, can be exploratory, projective, ampliative or revelatory (143). The exploratory imagination examines the objects of nature in front of the agent and supplies any missing pieces. The projective adds to the scene or object, or perhaps replaces pieces of it with other images or overlaps a projected image. The ampliative imagination may not involve images at all, says Brady, but involves "going beyond" what is given, interpreting the scene in a new way involving images or emotions or other associations (143). The revelatory imagination stretches what the agent sees to its limits, until she sees some greater truth. Brady gives an example of the revelatory mode at work in a person looking at a glacier and seeing the "tremendous power of the earth" (143). These four modes of activity must be tempered by disinterestedness, which Brady feels is necessary to prevent the agent from indulging in personal fantasy, and must be developed through practice until imagination becomes a virtue (145).

Brady's description of the modes of imagination as it is used for appreciation of nature can be applied to imagination as it is used for other tasks as well, including the task of identifying relevant alternatives. The exploratory mode, first of all, can be quite useful. Whether or not the agent has the information needed in her head already or whether she needs to do further investigation first, it is safe to assume that

she will not always be able to obtain absolutely all of the information that might be relevant. At some point, the agent will have to make a decision based on what information she has. However, if the agent possesses the virtue of imaginativeness, the exploratory mode will be useful in possibly “filling-in” some of the gaps in her information, making it more likely that she will have enough information to properly identify RA's.

For example, take the historian who is attempting to piece together the events of a past battle. She has information regarding the relative positions of the opposing armies at the beginning of the battle, as well as their numbers and armaments. She also knows some of the movements of the armies, and the outcome of the battle. However, she is missing some important information about the middle of the battle in which the army that was stronger and seemed certain to win suddenly faltered and ended up the losers. She has come up with several alternative explanations of what happened, but wants to ensure that she is not missing some further alternatives. The exploratory imagination will provide her with some of these alternate scenarios by providing possibilities for the missing information; perhaps the terrain favored the weaker army, perhaps there was a miscommunication between generals of the stronger army, and so on. The projective imagination will function similarly to the exploratory. The projective mode will allow the agent to project ahead to information she has not been able to obtain, and to integrate pieces of information that do not at first seem related but that, when put together, may provide additional insights.

The ampliative and revelatory modes will act in similar ways to each other. Both of these modes will assist the agent by allowing her to follow the implications of the information that she has along their natural course. The implications that the imagination will be uniquely able to discover will be those that are not immediately obvious to the average but unimaginative observer. For instance, an imaginative scientist will be able to come up with unique alternate theories to explain a phenomenon by “seeing” a possible implication of an accepted theory or of some body of information that other scientists do not see. This imagination can serve a scientist even if the alternative suggested by the imagination is ultimately incorrect. Theoretically, by first identifying *all* the alternatives and then eliminating all but one it would be possible to find the truth about the phenomenon in question. Some potential alternatives, however, would take a great deal of imagination to discover.

Imagination of the ampliative or revelatory modes perhaps had a part to play, for example, in theories of heat transfer. Scientists started from the observation that heat seems to move through objects. Put a metal poker in a fire and first the end of the poker that is in the fire will grow hot, and then the middle parts of the poker and finally the end nearest the handle. Furthermore, heat can be observed to travel faster in some materials than others. Scientists searched for alternatives to explain this behavior. One alternative was the caloric theory, which stated that heat was a kind of invisible liquid that moved through objects. Surely some imagination was involved in the postulation of an invisible liquid moving through objects to make them hot. Even more imagination, perhaps, was needed for the theory of heat which

succeeded caloric theory, a theory which states that heat is the result of molecular movement in objects. The fact that many new theories come from the minds of only one or two scientists (despite the fact that many others may be studying the same phenomenon) lends support to the idea that imagination may be needed for these scientists to come up with their alternate theories.

## **Chapter 5: Moral and Investigative Virtues and Relevant Alternatives**

Now that we have seen how investigative virtues contribute to making the agent practically wise, we can ask ourselves how these virtues work together with the moral virtues in RA cases. One question that comes to mind is *how many* virtues the agent will have to act in accordance with in order to act with practical wisdom in RA cases. Would she need to act in accordance with all possible moral virtues and investigative virtues in every individual case in order to avoid moral blameworthiness for her ignorance? There are some problems with this possibility. First of all, the lists of virtues given are not exhaustive, nor were they intended to be. In fact, making a fully exhaustive list of virtues that could be agreed on by most people may not be feasible. So, it would not be possible for us to determine in any given case whether an agent acted in accordance with all virtues or not. Even if we could determine that the agent acted in accordance with all known virtues, the possibility would be open that she failed to act in accordance with some other virtue we neglected to consider.

However, this issue can easily be resolved from the perspective of virtue ethics. Remember that virtue ethics is very much concerned with the process of persons becoming virtuous, and that accounts of moral education are central to virtue theory. As I discussed in an earlier section, an act can be virtuous even if the agent taking the action fails to possess some virtues or even if the agent does not possess *any* virtues. What matters is, rather, that the agent acts in the way that a

virtuous person would characteristically act. In other words, the agent will act in accordance with virtues, even if she does not fully possess those virtues.

However, this doesn't seem to fully address the problem. Rather, the question has merely changed from whether an agent needs to possess all the virtues to whether an agent needs to act in accordance with all the virtues. Once again, though, virtue ethics has an answer. First we can ask why an agent would need to act in accordance with every single virtue. One plausible answer to this question is that the agent needs to act in accordance with all virtues because the loss of any one virtue will have an effect of a certain sort on her disposition to act. However, as I noted in chapter two, this is not the way that virtues operate. Rather, the effects of virtues overlap and work together to create the overall character of the agent. Absence of one virtue will not necessarily completely alter a character, and given any decision made by an agent we typically cannot point to a single virtue that led to that decision. Since, in real life, we deal with imperfect persons who are in the process of becoming virtuous, these redundancies in the effects of virtues are actually quite necessary. Without them, it would be more difficult for a non-ideal moral agent, for example, to perform morally good acts.

Given the nature of virtues, then, it is not necessary for an agent to act in accordance with all moral and investigative virtues at all times in order to avoid moral blameworthiness in RA cases. At the same time, in any given situation the agent will still have to behave in a manner consistent with a good number of virtues. As my cases given in the discussion of individual cases above indicate, the absence

of action in accordance with certain virtues can, in some cases, cause the agent to act in a blameworthy manner. What effect lack of a specific virtue will have will depend on the specific case. Another way to express the thought that an agent must act in a manner characteristic of a certain number of virtues in order to avoid moral blameworthiness, and that the number and nature of those virtues will differ from case to case, is to say that an agent must act in accordance with practical wisdom in each case.<sup>11</sup>

The critic of the virtue approach might, at this point, challenge whether or not arguing that an agent must act with practical wisdom in order to avoid moral blameworthiness tells us enough to actually be able to determine, in specific cases, whether or not an agent is blameworthy for their ignorance. While it is true that exact answers, applicable in all cases, are not possible from this conception of moral blameworthiness, it is nonetheless the case that it gives us quite a lot of information. In any given case, we can look at the lists of virtues we do possess to determine which will have an impact on the given case. Further, we can use v-rules and wisdom gained from our moral educations to determine how a virtuous person would be disposed to act in the situation. Finally, we can determine whether or not the agent in question in the case succeeded in acting as the virtuous person would.

I now intend to summarize my arguments, and to provide an answer to the question I posed in my introduction of when an agent is morally blameworthy for the results of her ignorance with a single definition.

---

<sup>11</sup> This is so because practical wisdom is the disposition to adjudicate amongst virtues and determine in what cases one must action in accordance with certain virtues.

*Def. A person is morally blameworthy for the negative results of her ignorance if and only if she failed to undertake her investigation into the issue at hand in accordance with those moral and investigative virtues needed in the case at hand.*

In many (but not all, as we shall see below) cases, undertaking an investigation in the same way that an ideally virtuous person will dispose the agent to act and believe as the virtuous person would.

*Moral and Investigative Virtues at Work: The Case of Alex and Jane*

To see just how much work virtue ethics can do for us, I will use concepts of moral virtues and investigative virtues and my definition as given above to revisit the case I set out in my introduction. Recall that the case I laid out was one in which a worker named Alex believed that his co-worker, named Jane, was behaving very irresponsibly, resulting in damage to the productivity of their entire office. After observing this behavior for three weeks, Alex decided to confront Jane about it, and expressed his anger and disappointment in her and threatened to take his grievances to their boss. However, this confrontation resulted in Jane running off in evident anguish, and Alex soon learned that the actual situation was that Jane's husband had been in the hospital after a serious accident and that her erratic behavior was due to this fact and not to laziness or irresponsibility on her part. Alex was subsequently confronted by another colleague, who accused him of being very much in the wrong in the case, and who informed him that everyone else was aware of the true



situation. The colleague concluded that Alex “should have known better”.

According to my discussions, Alex's colleague would be correct in saying that Alex is morally blameworthy for his ignorance of Jane's situation if and only if he failed to act in accordance with the necessary moral and investigative virtues. In order to determine whether or not this is so, we must first look at which moral and investigative virtues are most relevant in the Jane case, and determine how possession of those virtues would dispose a morally virtuous agent to act in this case. Some moral virtues that are clearly relevant here are courage, generosity, thoughtfulness, compassion and justice. Both courage and generosity will dispose the agent to be willing to give her time and effort and risk being wrong in order to determine what the real situation is. Thoughtfulness and compassion will make the agent sensitive to and aware of Jane's true situation and will dispose her to feel for Jane and want to prevent her from suffering further. Finally, the virtue of justice will dispose the agent to determine whether or not Jane is in the wrong and what the proper method of proceeding is.

Some investigative virtues that are relevant in the Jane case are curiosity, humility, imagination, conscientiousness, empathy and fairness. Curiosity and humility will dispose the agent to fully investigate the situation and to be willing to realize that her initial conclusions might be wrong. Imagination will help her to consider other possible explanations of Jane's behavior besides the initial one that comes to her mind. Conscientiousness will ensure that she goes about her investigations into Jane's situation thoroughly and with respect for Jane. Fairness

will dispose the agent against being biased because of any personal liking or disliking she might have for Jane and will dispose her to be fair to Jane by considering as many possibilities as she can. Finally, empathy will dispose the agent to be open to see Jane's side of things and will lead her to want to know the true situation in case there is something she can do to assist Jane. Notice that some of the moral and investigative virtues play a similar role. Thoughtfulness and courage will have an effect on the agent's actions similar to curiosity and humility, as both will dispose the agent to learn more about the situation and to be willing to admit any errors she might have made with her initial assessment.

So, how might someone in possession of all these virtues have handled the Jane case? How might she have used these virtues to identify relevant alternatives to the proposition that Jane had been selfish and lazy, and in becoming aware of these alternatives avoided confronting Jane and causing her anguish? First of all, humility would caution the virtuous agent that her initial conclusion, although it might seem sound, is not necessarily correct; she would be aware of the possibility that she could have made a mistake. Courage, then, would dispose the agent to accept the possibility that she might be wrong and make her able to go ahead with an investigation that could prove her wrong. Given her willingness to admit error in her initial conclusions, the agent might use her imagination to come up with relevant alternatives that could explain Jane's behavior. She might consider the possibilities that Jane is ill, that she has some family obligations keeping her busy in the evenings, that she suffers from insomnia and so on.

After considering these alternatives, the virtuous agent would want to assess which ones are relevant, and examine whether there might be even more alternatives she has not yet considered. Thoughtfulness, generosity, curiosity, and thoroughness would dispose the virtuous agent to explore the avenues of investigation open to her. Furthermore, compassion and empathy would dispose the agent to want to pursue the possibility that Jane is in a painful situation so that she can help to alleviate that situation, while fairness and justice would dispose the agent to make sure that she has given Jane fair consideration and explored all possible avenues of investigation.

One of these avenues would certainly be the agent's own observations of Jane's behavior. However, another avenue of investigation would be to talk to the co-workers whom Jane might have confided in and Jane herself. The agent could ask Jane if she is feeling well and if she needs a day off. This would give Jane an opportunity to tell her concerned co-worker what was happening, and at the very least the agent would show kindness towards Jane and avoid being accusatory. Furthermore, the virtuous agent could easily ask her colleagues, without gossiping or prying, if anyone else noticed that Jane seems tired or ill lately. This would give her co-workers an opening to discretely inform her of the truth. As a result of these investigations, the virtuous agent would have discovered, through considering relevant alternatives, that the knowledge claim she initially wanted to make, that Jane is lazy and selfish, is incorrect and hence she would avoid taking actions based on that knowledge claim.

At this point, we can turn to the question of whether or not, in the actual Jane

case, Alex can be held to be morally blameworthy for his ignorance of Jane's situation. To answer this question, we need only to determine whether or not he undertook his investigation in the way in which the virtuous agent would. To summarize, I argued that the virtuous agent would be willing to admit that her initial assessment of Jane is wrong, would use her imagination to consider the possibility that there are relevant alternatives to that assessment, and would be disposed to find those alternatives out of curiosity, as well as from a sense of fairness, justice and empathy towards Jane. Further, she would be thorough and conscientious in investigating her available sources of information.

Has Alex carried out his investigation in the way in which the virtuous agent would? I believe that the answer is clearly “no.” Alex seems to be disposed to trust his initial judgment, and not at all disposed to seek alternative answers. He seems to accept his own assessment uncritically and spends his time considering what to do with this “knowledge” rather than investigating the possibility that Jane's behavior is due to other factors. In some respects, Alex is behaving in accordance with moral virtues. He shows compassion for Jane in confronting her privately and giving her a chance to improve her work habits before taking his grievances to their boss. This act also shows some courage, as he is willing to face Jane without “backup” and is therefore willing to take the brunt of the anger that Jane might possibly show for the sake of giving her a second chance. Furthermore, his decision about the way to confront Jane shows that Alex has acted with some conscientiousness and creativity: rather than simply filing a complaint to a higher authority in the office or taking

some other act that could be more harmful to Jane, he has taken the time to come to a decision that he believes will do her no damage and that will protect her privacy from the rest of the office.

However, these virtues, which Alex has exercised after coming to his conclusion that Jane is lazy, are not enough to keep Alex from being morally blameworthy for his ignorance of Jane's true situation. Alex has failed to consider that there may be other explanations for her behavior, either out of an over-confidence in his own judgment or from a lack of thoughtfulness or conscientiousness. Perhaps he has some degree of arrogance about the strength of his own instincts and assessments of others, or perhaps he is acting with undue haste or from a laziness of thought, meaning that he is simply unwilling to take the trouble to consider alternatives. This complete dismissal of the possibility of other explanations essentially precludes the possibility that Alex will act from other relevant virtues, such as empathy or fairness. Alex might be disposed to judge Jane less harshly under different circumstances, but he is not giving himself the chance to exercise that disposition due to his complete closed-mindedness to the possibility that there could be different circumstances at work.

Notice that I left open the possibility that Alex's behavior could be due to lack of any of a number of virtues. I argued that his apparent assumption that his assessment of Jane's situation is correct could be due to a lack of humility or to a lack of thoughtfulness or conscientiousness. Perhaps given more information about Alex or about these particular circumstances I could answer the question of which of

these virtues Alex is most likely lacking, or if he is most likely lacking in all of them. However, this information is not actually relevant to our assessment of Alex's moral blameworthiness. What we need to be able to determine is whether or not Alex acted (in this case, undertook his investigation) in a way in which a virtuous agent characteristically would. The virtuous agent's character is formed by all of the virtues she possesses working together, and there will be a great deal of overlap in the way the individual virtues the ideal agent possesses affect her character. What matters is the cumulative effect on character and action, not the specifics of what each individual virtue contributed.

### *Other Cases*

We can apply the techniques utilized in the above case to many other situations in which we ask whether or not a person is morally blameworthy for the consequences of their ignorance. These cases will fit the template laid out in my introduction. They will all involve an agent reaching a conclusion, taking some action based on that conclusion, and then later discovering that this action had some morally bad consequence. In all these cases, the agent will have failed to become aware of some information, F, which would have undermined her initial conclusion and dissuaded her from taking the morally damaging action. The agent's failure to become aware of F, I have since argued, stems from a failure to act in accordance with moral and investigative virtues. I want to look at one such case, and determine what my thesis would have to say about the moral blameworthiness of the agents

involved.

One situation in which the question of moral blameworthiness for ignorance might arise is in the case of voters in a given country who cast their ballots for a particular candidate without having much knowledge of that candidate's background or policies. Suppose that a candidate for the position of head of state, Ms. Green, states that she intends to change the law in order to allow the government to wiretap the phone of any citizen for the sake of national security. She mentions this plan when questioned at public debates, and it appears on a list of intentions on her campaign website. In the end, Green wins the election, but many of those who voted for her are surprised when she introduces the wiretapping bill. As the bill is debated, critics argue that the bill could allow personal information to be leaked and for arrests of innocent people to become rampant. One person who voted for Green, Mary, is confronted by a friend, who argues that, since Mary voted for Green, it will be partially her fault if innocent people are arrested on the basis of the bill. Mary protests against this accusation, stating that she "Had no idea" Green would do such a thing. Mary's friend is unimpressed, and tells her "Well, you *should* have known."

What can we make of this case? We can begin to solve it by asking ourselves how the morally and investigatively virtuous voter would act during an election campaign. While she would not devote all of her time to following the election to the exclusion of other aspects of her life, she would surely devote a good portion of her time to it. For instance, the virtues of justice and fairness would dispose the agent to study the policies of all the major candidates in order to make

an informed choice. Thoughtfulness, compassion and generosity would lead her to consider many aspects of the candidates' platforms, not just those that affected the agent personally. Conscientiousness and curiosity would prompt her to ensure that she became aware of and understood the candidates' policies and platforms.

Imagination would dispose her to think through those policies and the sorts of effects they might have on the country.

Has Mary acted as the morally and investigatively virtuous agent would?

The answer appears to be that she has not. If Mary actually did not know that Green planned to introduce the wiretapping bill, then she was not acting with sufficient conscientiousness and curiosity, as the information was readily available to her.

Mary might be generous and compassionate and desire to vote for a candidate whose policies are best for the country overall and not just for her personal concerns.

However, her lack of other virtues will have made this impossible. Since Mary has not put in enough effort to make herself aware of the candidates' policies, she will not be able to ensure that she votes for the candidate who is likeliest to form policies that are better for the country overall. As well, we may assume that Mary is imaginative enough to have become aware of the negative consequences of wiretapping if she had known of the bill and used her imagination to consider the possible outcomes of the policies she was aware of. However, she did not behave imaginatively with respect to the wiretapping policy, because she was unaware of the policy due to her failure to behave in accordance with virtues such as conscientiousness. Mary's friend, then, is right in stating that Mary should have



known that the candidate had such a policy, because a virtuous person would have had that knowledge. So, presumably, if the wiretapping bill went through and the rights of persons in the country were infringed on, then Mary would also share some responsibility for that.

Nonetheless, it might seem ridiculous for us to consider Mary, and any others who voted for Green without knowledge of her desired wiretapping bill, to be morally vicious and blameworthy for her action of voting for Green. Instinct tells us that a voter who knew about the bill and had the imagination to know where it would lead would deserve a much greater degree of moral blameworthiness. My theory can account for both of these instincts. I will argue that a person's degree of moral blameworthiness is dependent on the degree to which they behave as the virtuous agent would.

Consider the case of an agent, Martin, who is conscientious enough to be aware of the bill and imaginative enough to know where it might lead but who votes for Green because he doesn't mind the thought of wiretapping being used on others, and assumes that he has nothing to worry about. In this case, Martin is making a knowledge claim to himself that Green is the best candidate for the country. This case fits my template, as Green's knowledge claim is set up against alternatives, including the alternative that Green would be a disastrous head of state. Further, this belief has led Martin to take an action, that of voting for Green, that has contributed to specific consequences, such as Green getting elected and putting forward her bill. Nonetheless, Martin has acted as the virtuous agent would in certain ways: in

becoming aware of the various policies of the candidates and in using his imagination to become aware of the potential outcomes of those policies.

However, Martin has acted quite differently from the virtuous agent in other ways: he differed with respect to how he considers others in the country, how he assesses the effects of arrests of innocents and how he views his own opinions. The virtuous agent, out of generosity, compassion, justice and other virtues, would consider it wrong to vote for a candidate who would wish to harm others by spying on them or falsely arresting them. Further, due to empathy she would take the arguments against wiretapping seriously and recognize why it could be a dangerous policy for the citizens. Courage and humility would, further, dispose the virtuous agent to recognize that any preference she might have had for Green was misguided. Since Martin's behavior diverges from the virtuous agent in all these ways, he has a fairly high degree of moral blameworthiness compared to Mary, who only differed from the virtuous agent with respect to how thoroughly she investigated the candidate, but acted very much the same as the virtuous agent in considering the situation of others and so on.

Finally, consider the case of an agent, Susan, who votes for Green after acting exactly like the virtuous agent in most respects except that she lacks the imagination to see what the wiretapping bill might lead to and lacks the epistemic empathy to fully appreciate the arguments against wiretapping. Susan has tried to use her imagination and understand these arguments, but has failed. In this case, Susan has acted very much like the virtuous agent. The only divergence is that she

is quite far from fully possessing the virtues of imagination and empathy, even though she has developed generosity, justice, humility, conscientiousness and so on such that she is closer to being fully virtuous in these areas. Susan will still possess a degree of moral blameworthiness, but a lower one than Mary.

The above discussion does imply that people may have some degree of moral blameworthiness in a great number of situations. This view may seem unnecessarily harsh or negative, but keep in mind that sometimes this degree of blameworthiness might be very small. There will also be many cases in which humans are ignorant but through no fault of their own. Suppose in the election case that Green never indicated by any of her speeches, actions or past voting records that she would introduce something like the wiretap bill. However, shortly after being elected she put forward such a bill. In this case, even a virtuous agent possessing conscientiousness, generosity and imagination would not have known that Green being elected would have such a result.

#### *Two Further Points of Clarification*

There are two points regarding the scope of my conclusions that need to be addressed at this point. First of all, it is important to keep in mind that persons can very often undertake investigations as a virtuous person would, and hence acquire knowledge, even if those persons lack certain virtues. This follows from my statement above that there is a great deal overlap amongst the virtues (meaning that different virtues may have a similar effect on the dispositions of agents) as well as

from my discussion of moral action, in which I argued that non-virtuous persons can still succeed in performing morally right acts. This applies to epistemology as well: a non-epistemically virtuous agent can still believe in an epistemically sound fashion.

What the agent will need to do in order to be able to consistently perform morally good acts or to consistently obtain knowledge will be to act in accordance with a number of relevant virtues and to make an effort to become more virtuous. This effort is made throughout the agent's life, and is not necessarily consciously made in every single situation that agent faces. In fact, an agent may not necessarily ever think in terms of virtues and vices. Nonetheless, an agent who does not think in these terms may still be successful in her attempts to become a morally good person and may still consistently act in accordance with the virtues.

Secondly, more needs to be said about cases in which a proposition is knowable or known by the ideally virtuous agent, but only through her ideal level of virtue. For example, suppose that there is some concept or theory that could only be conceived of by someone with an extraordinary imagination. The ideally virtuous agent would be aware of this theory, but the normal person almost certainly would not. Now supposing that ignorance of this theory leads an agent to perform an act that leads to a morally wrong result, even though she acted in accordance with virtue concerning all other aspects of her situation and even though she made a great effort to exercise her imagination. It seems unusually cruel to condemn this agent as morally blameworthy merely because she lacked an extraordinary level of

imagination.

This issue can be dealt with within the theory of moral and investigative virtues as I have laid it out. In the case of normal levels of virtue, a person carrying out an investigation in the way that a virtuous person will exercise her virtues in a similar way and hence will achieve similar information through those virtues. For example, she will act as the person of conscientiousness would act, and hence will investigate the case thoroughly and will become aware of new possibilities, the knowledge of which will affect her ultimate conclusion. No extraordinary level of conscientiousness is necessary here, so a normal agent, even if she does not fully possess the virtue of conscientiousness, will be able to act and believe as an ideally conscientious agent would.

However, if extraordinary conscientiousness were needed, then it would be perfectly feasible for the non-ideal agent to act with conscientiousness but not quite discover that extra information. Nonetheless, even though the information she finally reaches will be different, the agent still undertook her investigation in the same way that the virtuous person would, which was all the definition given above required. Another way to word this is that, in cases where the ideally virtuous agent's actions or beliefs stem from an extraordinary level of virtue, the moral agent does not have to *act* or *believe* as the virtuous agent would in order to avoid moral blameworthiness. Rather, it is sufficient that she exercises the relevant virtue to the extent she can, and that she undertakes her investigation as the ideally virtuous agent would.

Furthermore, the virtue theorist will never judge a person to be vicious or virtuous on the basis of a single act. Basic moral judgments in virtue theory are judgments of character, and character is judged by how an agent behaves over a period of time, not with respect to a single action. So, an agent may fail to carry out an investigation in the way that an ideally virtuous agent would in a specific situation, but if we judge that her actions and investigations, *characteristically*, are done in accordance with virtue, then we will not judge the agent to be vicious. She may be morally blameworthy in this one action, but this does not entail that she is not a virtuous person.

## Conclusion

I set out in this thesis to argue that an agent is morally blameworthy for the morally bad consequences of her ignorance if she has failed to undertake her investigation in accordance with the moral and investigative virtues. My investigation led to relevant alternatives literature, but I argued that the answers it could provide to the question of when an agent is responsible for her lack of knowledge were unsatisfying. The rules or criteria given in relevant alternatives literature were invariably vague regarding their limits of application; a significant failing given that one supposed advantage of rules or universally applicable criteria is that they can provide the agent with clear, indisputable answers in difficult situations. From this dissatisfaction with the answers provided by relevant alternatives literature, I moved towards a virtue approach.

My goal in taking this virtue approach has been to provide a method for answering a very interesting problem for ethics and epistemology. I have aimed to create a basic working model of this approach, and to add to the body of virtue literature. From the examples given in chapter five, it is clear how this theory can be applied to “real-world” situations. However, the focus is, as it should be from a virtue-centered approach, on character and the kind of person who can avoid moral blameworthiness. This kind of person will be morally good and practically wise. She will possess both moral and investigative virtues, and will use them successfully as she encounters questions of knowledge throughout her everyday life. With her moral and investigative virtues, she will be morally justified in believing that she

knows everything from day-to-day claims about the world she sees around her to more complex theories about the nature of the world.

Although the ideally virtuous investigator is a necessary fixture of the theory, she does not necessarily exist in the real world. In reality, people are imperfect moral and investigative agents, in the process of developing our virtues. My theory gives an account of what virtuous action means to real-life agents, and how we might come to determine what is the right thing to do or the right way to carry out our investigations. We are fallible, and will often err in our actions or beliefs. However, this does not have to mean that we are usually morally bad people. In fact, we often succeed at doing the right thing by acting in accordance with kindness, courage, justice and so on. When we do err, it is not usually in a very morally serious manner. Furthermore, being able to recognize failure and the reasons for it help us in furthering our moral education and avoiding making the same mistake in future situations.

This thesis has laid the groundwork for a virtue approach to relevant alternatives cases and a virtue-centered account of moral blameworthiness for ignorance. From here, more investigative virtues could be identified and discussed. One possible avenue of further investigation is how the significant moral and investigative virtues and degrees of moral blameworthiness may be affected by the relationship of the agent to those receiving her knowledge claims. Ultimately, this thesis has offered a possible reply to an interesting epistemic and moral problem. Rather than assessing a knower based on her ability to follow a specific procedure or



set of rules, I suggest evaluating her according to how well she has undergone her investigations in accordance with virtues. Whether or not an agent can be blamed for reaching a false conclusion at the end of an inquiry can be discovered by determining how well she has followed the moral and investigative virtues. This method of assessment is widely applicable, and fits with many of our intuitions about ignorance and blameworthiness. It is my belief that virtue account of moral responsibility for knowledge could have a significant role to play in philosophical literature.

## Bibliography

- Annas, Julia. "Being Virtuous and Doing the Right Thing." Pacific Division Presidential Address, *Proceedings of the American Philosophical Association*. 78.2 (Nov 2004), 61-75.
- Aristotle. *The Nicomachean Ethics*. Trans. J.A.K. Thomson, Revised Hugh Tredennick. London: Penguin Books, 2004.
- Bonjour, Lawrence. "Sosa on Knowledge, Justification and 'Aptness.'" *Knowledge, Belief, and Character: Readings in Contemporary Virtue Epistemology*. Ed. Guy Axtell. Lanham, Md: Rowman & Littlefield Publishers, 2000. 87-98.
- Brady, Michael and Duncan Pritchard. "Introduction." *Moral and Epistemic Virtues*. Eds. Brady, Michael and Duncan Pritchard. Oxford: Blackwell, 2003. 1-12.
- Code, Lorraine. "Towards a 'Responsibilist' Epistemology." *Philosophy and Phenomenological Research*. 45.1 (Sept 1984). 29-43.
- Cohen, S. "Skepticism, Relevance, and Relativity." In B. McLaughlin, *Dretske and his Critics*. Oxford: Basil Blackwell, 1991. 17-37.
- Dancy, Jonathan. "Supervenience, Virtues and Consequences." In *Knowledge, Belief, and Character: Readings in Contemporary Virtue Epistemology*. Ed. Guy Axtell. Lanham, Md: Rowman & Littlefield Publishers, 2000. 73-86.
- Das, Ramon. "Virtue Ethics and Right Action." *Australasian Journal of Philosophy*. 81.3 (Sept 2003), 324-339.
- DePaul, Michael and Linda Zagzebski, Eds. *Intellectual Virtue: Perspectives from*

*Ethics and Epistemology*. Oxford: Oxford University Press, 2003.

Dretske, F. "Epistemic Operators." *The Journal of Philosophy* 67.24 (1970):  
1007-1023.

Dretske, F. "The Pragmatic Dimension of Knowledge." *Philosophical  
Studies* 40 (1981): 363-378.

Driver, Julia. "The Conflation of Moral and Epistemic Virtue." In *Moral and  
Epistemic Virtues*. Eds. Brady, Michael and Duncan Pritchard. Oxford:  
Blackwell, 2003. 101-116.

Driver, Julia. "The Virtues and Human Nature." In *How Should One Live?: Essays  
on the Virtues*. Ed. Roger Crisp. Oxford: Clarendon Press, 1996. 111-129.

Eflin, Juli. "Epistemic Presuppositions and Their Consequences." In *Moral and  
Epistemic Virtues*. Eds. Brady, Michael and Duncan Pritchard. Oxford:  
Blackwell, 2003. 47-66.

Haack, Susan. "The Ethics of Belief Reconsidered." In *Knowledge, Truth and  
Duty: Essays on Epistemic Justification, Responsibility and Virtue*. Ed.  
Mathias Steup. Oxford: Oxford University Press, 2001. 21-33.

Hursthouse, Rosalind. "Virtue Theory and Abortion." *Philosophy and Public  
Affairs*. 20.3 (Summer 1991): 223-246.

Hursthouse, Rosalind. "Normative Virtue Ethics." In *How Should One Live?:  
Essays on the Virtues*. Ed. Roger Crisp. Oxford: Clarendon Press, 1996.  
19-36.

Hursthouse, Rosalind. *On Virtue Ethics*. Oxford: Oxford University Press, 2001.

- Hursthouse, Rosalind. "Virtue Ethics and the Emotions." In *Virtue Ethics: A Critical Reader*. Ed. Daniel Statman. Washington D.C.: Georgetown University Press, 1997. 99-117.
- Lewis, D. "Elusive Knowledge" *Australasian Journal of Philosophy* 74 (1996): 549-567.
- Monmarquet, James. "An 'Internalist' Conception of Epistemic Virtue." In *Knowledge, Belief, and Character: Readings in Contemporary Virtue Epistemology*. Ed. Guy Axtell. Lanham, Md: Rowman & Littlefield Publishers, 2000. 135-148.
- Paul, Richard. "Critical Thinking, Moral Integrity, and Citizenship: Teaching for the Intellectual Virtues." In *Knowledge, Belief, and Character: Readings in Contemporary Virtue Epistemology*. Ed. Guy Axtell. Lanham, Md: Rowman & Littlefield Publishers, 2000. 163-176.
- Pincoffs, Edmund L. *Quandaries and Virtues: Against Reductivism in Ethics*. Kansas: University Press of Kansas, 1986.
- Rachels, James. "The Ethics of Virtue." Reprinted from *Elements of Moral Philosophy*. In *Vice and Virtue in Everyday Life: Introductory Readings in Ethics*. Eds. Christina Sommers and Fred Sommers. Toronto: Nelson, 2004. 241-255.
- Slote, Michael. "From Morality to Virtue." In *Virtue Ethics: A Critical Reader*. Ed. Daniel Statman. Washington D.C.: Georgetown University Press, 1997. 128-144.

Stine, G. "Skepticism, Relevant Alternatives, and Deductive Closure". *Philosophical Studies* 29(1976): 249-26.

Swanton, Christine. "A Virtue Ethical Account of Right Action." *Ethics*. 112.1 (Oct 2001), 32-52.

Wallace, James D. *Virtues and Vices*. Ithaca: Cornell University Press, 1978.

Zagzebski, Linda Trinkaus. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press, 1996.